NASA Contractor Report 189144

On-Board B-ISDN Fast Packet Switching Architectures

Phase I: Study

Faris Faris, Thomas Inukai, Fred Lee, Dilip Paul and Dong-Jye Shyy
Comsat Laboratories
Clarksburg, Maryland

January 1993

Prepared for
Lewis Research Center
Under Grant NAGW–4528

NASA
National Aeronautics and
Space Administration

FINAL REPORT

ON-BOARD B-ISDN FAST PACKET SWITCHING ARCHITECTURES

- PHASE I: STUDY -

Submitted to

Code EC
Attn: NRA-89-OSSA-14
National Aeronautics and Space Administration Headquarters
Washington, DC 20546

Work Performed Under Research Announcement NRA-89-OSSA-14 for
Satellite Communications Applications Research (SCAR) Program

January 31, 1992

Investigators: Faris Faris
Thomas Inukai (Principal Investigator)
Fred Lee
Dilip Paul
Dong-Jye Shyy

COMSAT LABORATORIES
22300 COMSAT DRIVE, CLARKSBURG, MD 20871

Executive Summary

The broadband integrated services digital network (B-ISDN) is an emerging telecommunications technology that will meet most of the telecommunications networking needs in the mid-1990s to early next century. The B-ISDN is characterized by very high-speed user access to the network at 155 Mbit/s and higher, the use of fiber optic transmission media for user access and networking, accommodation of various circuit and packet based services, and high-speed packet based switching systems (fast packet switching). Typical B-ISDN services include broadband video-telephony, broadband videoconference, high-speed image transmission, high volume file transfer, high-speed telefax, high definition TV (HDTV), and broadband videotext. CCITT is currently finalizing its recommendations on B-ISDN services, and telecommunications carriers/administrations and equipment manufacturers are preparing for network deployment and service offerings to the customers.

The terrestrial-based B-ISDN will likely have a significant impact on satellite communications, requiring the same type and quality of services to satellite network users. The satellite-based system is well positioned for providing B-ISDN service with its inherent capabilities of point-to-multipoint/broadcast transmission, virtually unlimited connectivity between any two points within a beam coverage, short deployment time of communications facility, flexible and dynamic reallocation of space segment capacity, and distance insensitive cost. On-board processing satellites, particularly in a multiple spot beam environment, will provide enhanced connectivity, better performance, optimized access/transmission link design, and most importantly lower user service cost.

A satellite B-ISDN system can provide direct (mesh) interconnection among users, interconnection between a user and a switching center, and interconnection between switching centers. Examples of satellite B-ISDN services include full motion video distribution, mesh private and business networks, high-speed trunking between two locations, emergency communications, thin-route networking, science data and communications networking among NASA centers and users, and supercomputer networking. Although the user interface rate for B-ISDN service is 155.52 Mbit/s (51.84 Mbit/s for SONET OC-1) or higher, its information rate can be very low (e.g. 64 kbit/s). Also, traffic may be extremely bursty, such as in LAN/MAN interconnection, or may be a steady flow of high-speed data as in video program distribution. Thus, a satellite network must be flexible to accommodate a wide range of transmission rates and different degrees of burstiness. In addition, a satellite network must support Asynchronous Transfer Mode (ATM) traffic, which is a packet-based transmission mechanism to provide connection-oriented as well as connectionless services at various

information rates. These new environments that are significantly different from the traditional circuit-oriented system must be taken into consideration in satellite system design.

The satellite can support high-speed transmission at a bit rate between 155 Mbit/s and 1.24 Gbit/s using an antenna size of between 2.4 m (155 Mbit/s at Ka-band) and 12 m (1.24 Gbit/s at Ku-band). These high-bit rates are primarily intended for use in trunking and high-speed circuit-switched applications and often demand efficient bandwidth utilization, low bit error rates (e.g. 10^{-11}), and high availability. A simpler on-board switching payload, such as SS-TDMA and SS-FDMA (in RF or baseband), may be used for beam interconnection.

The majority of future B-ISDN users will likely require bit rates between 1.544 Mbit/s and 155 Mbit/s and flexible interconnection among a large number of user terminals. In this environment, user earth station cost becomes of prime importance in designing a satellite system. Flexibility and efficient bandwidth utilization can be achieved with an on-board fast packet switch to provide an integrated circuit and packet-switched service. Uplink access and transmission can be optimized to a group of users according to their traffic requirements, and downlink may use a single carrier TDM transmission to maximize the use of spacecraft power. The fast packet switch can also provide integrated operation of B-ISDN and non-B-ISDN traffic, including N-ISDN, frame relay, X.25/X.75, LANs, MANs, and SMDS.

On-board baseband switching provides interconnection of user earth stations operating at different bit rates and access schemes. Circuit switching is most efficient for handling circuit-switched traffic in which assigned time/frequency slots are fully occupied by a constant flow of traffic. Circuit switching can also support transmission of packet-switched traffic, but its effectiveness diminishes as the connectivity of the network increases. Fast packet switching, on the other hand, not only efficiently accommodates packet-switched traffic with variable destinations but also supports circuit-switched traffic. However, some means of avoiding or minimizing traffic congestion must be incorporated in the network design. Unlike circuit switching, the on-board processor does not require control memory to route packets through a baseband switch. In some situations, a mixed switch configuration, consisting of a circuit switch and a packet switch, may result in an optimal on-board processor architecture. An example of such a system is a network which provides a high-speed trunking service as well as a packet-based data service for a large number of users. On-board interconnection is achieved by cross-strapping the two types of paylaods.

There are several architectural alternatives for a fast packet switch design. A conventional circuit switch, such as a common memory switch, a high-speed parallel bus, or a fiber optic ring, can be used for self-routing. There will be no internal contention problem with this type of switch. The total switch throughput is limited to a few gigabits/s due to a shared facility (a memory, a bus, or a ring). Multistage interconnection networks, such as banyan-based switches and self-routing crossbar switches, can provide higher throughputs. The contention problem with a multistage switch can be avoided using, for example, an output port reservation scheme or may be

minimized by increasing a switch speed or by providing multiple routing paths within a switch fabric.

The technology assessment based on a detailed design indicates that the power consumption of a multicast packet switch is about 9 watts per port using the current technology and is likely to be reduced to around 3 - 4 watts in the near future. This will translate into a total switching subsystem power requirement of 48 - 64 watts for a 16 x 16 fast packet switch. If all other processing hardware is included, the on-board baseband processor may require several hundred watts of power. This illustrates that the hardware complexity of an on-board fast packet switch is not significant compared with that of the total processing payload.

Table of Contents

Table of Contents (cont'd)

Table of Contents (cont'd)

Table of Contents (cont'd)

List of Illustrations

List of Illustrations (cont'd)

List of Illustrations (cont'd)

List of Illustrations (cont'd)

List of Illustrations (cont'd)

List of Illustrations (cont'd)

List of Illustrations (cont'd)

List of Tables

List of Tables (cont'd)

List of Tables (cont'd)

Acronyms and Abbreviations

AAL	ATM Adaptation Layer
ACK	Acknowledgement
ACTS	Advanced Communications Technology Satellite
ANSI	American National Standards Institute
APD	Avalanche Photodetector
APS	Automatic Protection Switching
ARQ	Automatic Repeat Request
ASIC	Application Specific Integrated Circuit
ATM	Asynchronous Transfer Mode
AU	Administrative Unit
B-ISDN	Broadband Integrated Services Digital Network
BER	Bit Error Rate
BIP	Bit Interleaved Parity
BPSK	Binary Phase-Shift Keying
BSM	Baseband Switch Matrix
BTP	Burst Time Plan
C/N	Carrier-to-Noise Ratio
CAD	Computer Aided Design
CAM	Computer Aided Manufacture
CBR	Constant Bit Rate
CCITT	International Telephone & Telegraph Consultative Committee
CCSDS	Consultative Committee for Space Data Systems
CDJ	Cell Delay Jitter
CIR	Cell Insertion Ratio
CLP	Cell Loss Priority
CLR	Cell Loss Ratio
CONUS	Continental United States
CP	Complete Partitioning
CRC	Cyclical Redundancy Check
CS	Complete Sharing

Acronyms and Abbreviations (cont'd)

CS	Convergence Sublayer
CTD	Cell Transfer Delay
CVCDU	Coded Virtual Channel Data Unit
DCE	Data Circuit-terminating Equipment
DE	Discard Eligibility
DEMUX	Demultiplexer
DLCI	Data Link Channel Identifier
DPCM	Differential Pulse-Code Modulation
DQDB	Distributed Queue Dual Bus
DSI	Digital Speech Interpolation
DTE	Data Terminal Equipment
E/A	Address Field Extension Bit
EIRP	Equivalent Isotropic Radiated Power
FCFS	First Come First Serve
FCS	Frame Check Sequence
FDDI	Fiber Distributed Data Interface
FDMA	Frequency Division Multiple Access
FEC	Forward Error Control
FIFO	First In First Out
FISU	Fill-In Status Unit
G/T	Gain-to-Temperature Ratio
GEO	Geosynchronous
GFC	Generic Flow Control
GOS	Grade of Service
HDLC	High-level Data Link Control
HDTV	High Definition Television
HEC	Header Error Control
HIPPI	High Performance Parallel Interface
HOL	Head of Line
HPA	High Power Amplifier

Acronyms and Abbreviations (cont'd)

IB	Internal Blocking
ISDN	Integrated Services Digital Network
ISL	Intersatellite Link
ISO	International Standards Organization
LAN	Local Area Network
LAPB	Balanced Link Access Procedure
LAPD	D-Channel Link Access Procedure
LI	Length Indication
LNA	Low Noise Amplifier
LSI	Large-Scale Integration
MAN	Metropolitan Area Network
MCD	Multicarrier Demodulator
MID	Multiplexing Identifier
MSM	Microwave Switch Matrix
MTP	Message Transfer Protocol
MUX	Multiplexer
N-ISDN	Narrowband Integrated Services Digital Network
NACK	Negative Acknowledgement
NCC	Network Control Center
NEA	Noise Equivalent Angle
NNI	Network Node Interface
OAM	Operation and Maintenance
OBNC	On-Board Network Controller
OBP	On-Board Baseband Processor
OC	Optical Carrier
OC	Output Contention
OH	Overhead
OOK	On-Off Keying
OSI	Open Systems Interconnection
PCM	Pulse-Code Modulation

Acronyms and Abbreviations (cont'd)

PDU	Protocol Data Unit
POH	Path Overhead
PSK	Phase Shift Keying
PT	Payload Type
QAM	Quaternary Amplitude Modulation
QFSK	Quaternary Frequency-Shift Keying
QOS	Quality of Service
QPPM	Quaternary Pulse Position Modulation
QPSK	Quaternary Phase-Shift Keying
RES	Reserved Field
RF	Radio Frequency
RS	Reed-Solomon
RX	Receiver
S/P	Serial to Parallel
SAPI	Service Access Point Identifier
SAR	Segmentation & Reassembly Sublayer
SDH	Synchronous Digital Hierarchy
SDU	Service Data Unit
SMDS	Switched Multimegabit Data Service
SN	Sequence Number
SOH	Section Overhead
SONET	Synchronous Optical Network
SS-FDMA	Satellite-Switched Frequency Division Multiple Access
SS-TDMA	Satellite-Switched Time Division Multiple Access
SS7	Signaling System Number 7
SSPA	Solid State Power Amplifier
SSTP	Switch State Time Plan
ST	Segment Type
STM	Synchronous Transport Module
STS	Synchronous Transport Signal

Acronyms and Abbreviations (cont'd)

SVL	Satellite Virtual Label
SVP	Satellite Virtual Packet
SYNC	Synchronization
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TEI	Terminal Endpoint Identifier
TST	Time Space Time
TX	Transmitter
UNI	User Network Interface
VBR	Variable Bit Rate
VC	Virtual Channel or Virtual Container
VCDU	Virtual Channel Data Unit
VCI	Virtual Channel Identifier
VPI	Virtual Path Identifier
VSAT	Very Small Aperture Terminal
WDM	Wavelength Division Multiplexing

Section 1

Introduction

The broadband ISDN (B-ISDN) extends the service integration concept to include future high speed, high volume network services as well as to accommodate narrowband ISDN in its infrastructure. The B-ISDN overcomes the communications bottleneck existing in the current network by providing high speed user access to the network and broadband inter-nodal communications links along with supporting future high speed communications services, such as broadband video telephony, high-speed video/audio information transmission, high-speed digital information transmission, and high-volume file transfer. At present, implementation of the B-ISDN is proceeding at a rapid pace in the United States, Europe, and Japan. It is believed that this emerging technology will have a significant impact on not only the business world but on people's daily lives as well.

The impact of B-ISDN services on satellite communications can be significant, requiring the same or similar quality of communications services provided by the terrestrial B-ISDN network. The purpose of this study is to investigate potential satellite applications, alternate satellite network architectures, system design issues, and on-board switching/processing options for satellite B-ISDN services. Special attention is focused on fast packet switching architectures and switch design details to exploit the feasibility of implementing a fast packet switch in the future advanced communications satellites. The sections of this report are organized as follows.

Section 1 addresses B-ISDN service characteristics by summarizing the user and network aspects of broadband services and by identifying potential services and their transmission characteristics. A brief overview of Asynchronous Transfer Mode (ATM) and Synchronous Optical Network (SONET) is also provided in this section.

Section 2 presents a survey of current development status in broadband services including B-ISDN equipment, HDTV codec, ATM switches, metropolitan area network (MAN) products, and broadband service offerings and deployment plans by various telecommunications entities.

Section 4 includes an identification of potential satellite applications of B-ISDN services and typical bit rate requirements, on-board processing options, a comparison of circuit and fast packet switching, and satellite network architecture alternatives. Three sample network architectures, i.e., high-speed circuit switched applications,

business/thin-route networks, and NASA science/networking applications, are also described in detail.

Section 5 address a number of system issues associated with satellite B-ISDN system design for utilizing satellite and ground resources most effectively. The system design considerations presented are ATM cell concentration, ATM cell packetization alternatives, bit error rate and quality of service, ATM cell header protection, and impact of satellite delay. Some of these issues are not only associated with B-ISDN but also applicable to on-board fast packet switching for non-B-ISDN traffic and satellite data communications in general.

Section 6 describes the design principles of different fast packet switching architectures for point-to-point and point-to-multipoint connections. Specific topics in this section include a basic principle, a classification of fast packet switches, multicast/broadcast operation, and priority control.

Section 7 presents detailed block diagram designs of fast packet switches. Five switch architectures are selected for detailed design: two point-to-point switching networks (i.e., sorted Banyan and self-routing crossbar) and three multicast switching networks (i.e., self-routing multicast Banyan, self-routing multicast crossbar, and point-to-point switch fabric with multicast output ports). This section also addresses fault tolerant design techniques for the fast packet switching subsystem.

The report includes four appendices on specific self-contained topics: (a) RF transmission system design, (b) optical transmission system design, (c) satellite B-ISDN traffic analysis, and (d) traffic simulation of fast packet switched networks. The use of optical links may be limited to special applications, such as broadband transmission of NASA science data between the satellite and certain geographic locations and intersatellite link transmission.

Section 2

B-ISDN Service Characteristics

Broadband communication and information services demand higher bandwidths and information bit rates than does a regular telephone. This section addresses the user and network aspects of broadband services which need to support both switched and non-switched ISDN services. The features and characteristics (peak bit rate, information density, holding time, and performance criteria) of each B-ISDN service are examined. The results will be used to develop network scenarios, input traffic patterns for numerical analyses, simulations, and alternative satellite network architectures in the following sections. Also included in this section is a brief description of the cell structure and header format of Asynchronous Transfer Mode (ATM) and the frame structure of Synchronous Digital Hierachy (SDH) and Synchronous Optical Network (SONET).

2.1 Classification of B-ISDN Services

Two main services have been categorized according to CCITT Recommendation I.211: interactive services and distribution services. The interactive services are divided into three classes of services: the conversational services, the messaging services, and the retrieval services. The distribution services are divided into distributed services without user individual presentation and distributed services with user individual presentation. The results are five groups of services: three interactive and two distributive. Figure 2-1 summarizes this classification.

2.1.1 Interactive Services: Transmission in Both Directions

Interactive services are a two-way exchange of information (which does not include signaling information).

2.1.1.1 Conversational Services: Sender and Receiver Simultaneously Present

Conversational services provide the means for bidirectional dialogue communication with real-time (no store-and-forward) end-to-end transfer of information from user to user or between user and service-provider host. The information is generated by the sending user(s) and is dedicated to one or more individual communication partners at

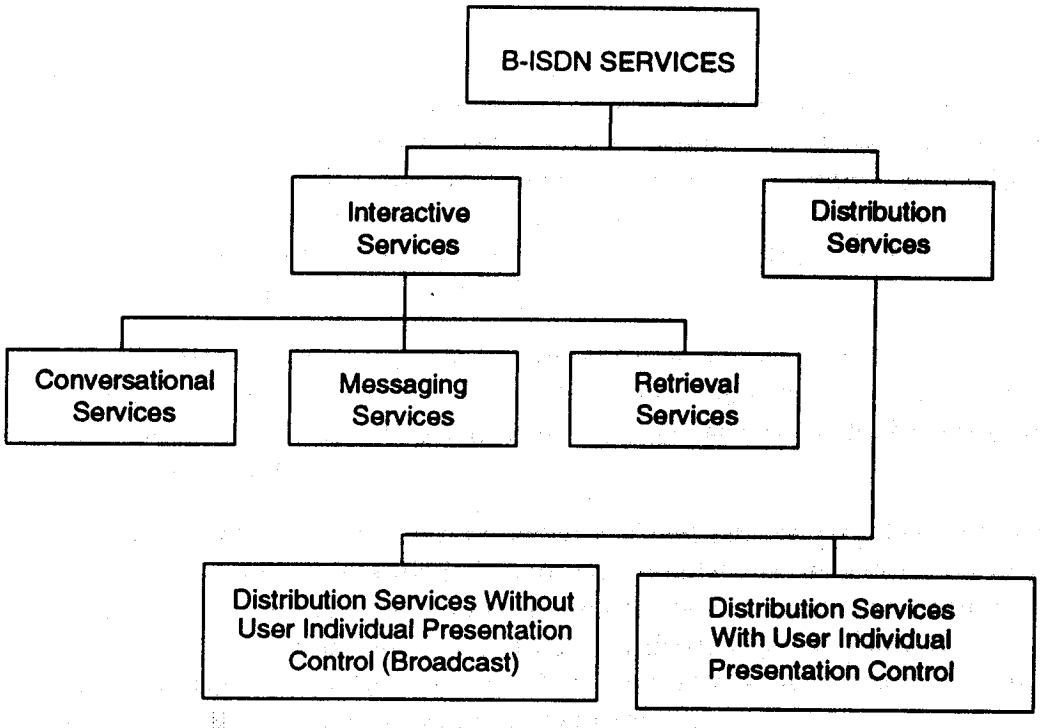

Figure 2-1. Classification of Broadband Services

the receiving side. The information in this category is not public information. The flow of the user information may be bidirectional symmetric, bidirectional asymmetric, or in some cases, unidirectional. A number of services and applications are provided as follows.

(a) Moving Pictures (Video) and Sound

(a.1) Broadband Video Telephony (Between Two Locations)

Video telephony means the telephone set includes a video transmit, receive, and display capability, and telephone calls will consist of voice and video information. Its application will occur when visual presentation is necessary or advantageous, for example in sales, consulting, instruction, and negotiation, and in the discussion of visual information, such as reports and charts.

Applications

- Tele-education (distance learning)
- Tele-shopping
- Tele-advertising

Possible Attributes

- Demand/reserved/permanent (establishment for communication)
- Point-to-point/multipoint (communication configuration)
- Bidirectional symmetric/bidirectional asymmetric (symmetry)

(a.2) Broadband Video Conference (Multipoint)

Video conference is used to exchange information among conference rooms, which can be point-to-point or point-to-multipoint. The difference of this service from the video telephony is the equipment used. For example, a point-to-point video conference may specify additional features such as facsimile and document transfer.

Applications

- Studio conference
- Research
- Tele-education and training
- Tele-shopping
- Tele-advertising

Possible Attributes

- Demand/reserved/permanent
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

(a.3) Video Surveillance

This service is only provided for a limited number of users; hence, it is not a distribution service. This service can be unidirectional, i.e., the user has no control over the video source.

Applications

- Building security
- Traffic monitoring

Possible Attributes

- Demand/reserved/permanent
- Point-to-point/multipoint
- Bidirectional symmetric/unidirectional

(a.4) Video/audio Information Transmission Service

This service is similar with the video telephony service. The difference is that in this service, a higher resolution and a higher quality video image is required.

Applications

- TV signal transfer
- Video/audio dialogue
- Contribution of information

Possible Attributes

- Demand/reserved/permanent
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

(b) Sound

(b.1) Multiple Sound-Program Signals

Applications

- Multilingual commentary channels
- Multiple program transfer

Possible Attributes

- Demand/reserved/permanent
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

(c) Data

(c.1) High-speed Unrestricted Digital Information Transmission Service

Applications

- High-speed data transfer
- Local area network (LAN) interconnection
- Computer-computer interconnection
- Metropolitan area network (MAN) applications
- Transfer of video and other information types
- Still-image transfer
- Multi-site interactive CAD/CAM
- Distributed processing

Possible Attributes

- Demand/reserved/permanent
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric
- Connection oriented/connectionless

(c.2) High-volume File Transfer Service

Applications

- Data file transfer
- Loading programs
- Changing database

Possible Attributes

- Demand
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

(c.3) High-Speed Teleaction

This provides user high speed responsive real time data.

Applications

- Real-time control
- Telemetry
- Alarms

Possible Attributes

- Demand
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

(d) Document

(d.1) High-Speed Telefax

Applications

- User-to-user transfer of text, images, drawings, etc.

Possible Attributes

- Demand
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

(d.2) High-Resolution Image Communication Service

Applications

- Professional images
- Medical images
- Remote games and game networks

Possible Attributes

- Demand

- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

(d.3) Document Communication Service

An electronic document can contain images, text, and data as well as commentaries in digitized form.

Applications

- User-to-user transfer of mixed documentation
- Remote printing

Possible Attributes

- Demand
- Point-to-point/multipoint
- Bidirectional symmetric/bidirectional asymmetric

Note that a mixed document is a document that may contain text, graphics, still and moving picture information, and voice annotation.

Characteristics of sample conversational services are provided in Table 2-1.

2.1.1.2 Messaging Services: Sending Information to a Database for Subsequent Retrieval

Messaging services offer user-to-user communication between individual users via storage units that have store-and-forward, mailbox, and/or message handling functions. This category does not require real-time transmission; therefore, both users do not have to be present at the same time, and the priority of this service is lower in the network. Although the end-to-end delay from sender to receiver might be long, the messaging services delay is only from the sender to the intermediate database. This delay should be less than the conversational services. The intermediate database is unable to recover the loss of information; hence, a low error bit rate is imposed for important applications. The messaging services provide only the send function, and the retrieval services provide the receive function. A number of services and applications are provided as follows.

Table 2-1. Characteristics of Sample Conversational Services [2-1]-[2-3]

SAMPLE SERVICES	BIT RATE (bit/s)	INFORMATION DENSITY	HOLDING TIME	PERFORMANCE CRITERIA
Alarm	10 - 100			
Telemetry	100 - 1K	low - 0.01	short	
Security monitoring	100 - 1K	low - 0.05	short	
Electronic fund transfer	300 - high	low - 1.0	short, medium	error rate
Inquiry/response	300 - 9.6K	< 0.1	short	delay
Voice	16K - 64K	0.4 - 0.5	medium	delay
Video telephony	64K - 45M	0.4 - 1	medium	delay
High Fidelity Audio	200K - 800K	0.5 - 0.8	medium	delay/error rate
Facsimile	1.2K - high	0.8 - 1.0	medium, large	error rate
Color facsimile	2M	0.8 - 1.0	medium, large	error rate
Bulk data transfer	1.2K - high	0.8 - 1.0	medium, large	error rate
CAD/CAM	10K - 100M	0.8 - 1.0	medium, large	error rate
File transfer	10K - 50M	0.8 - 1.0	medium, large	error rate
LAN interconnection	1M - 100M	low - 1.0	medium, large	error rate/delay
High-speed LAN	600M-1G	low - 1.0	medium, large	error rate/delay
MAN	155M - 600M	low - 1.0	medium, large	error rate/delay
Video transmission	30M - 45M	low - 1.0	medium, large	delay
Video conference	0.38M - 45M	low - 1.0	medium, large	delay

Note: Holding time is the mean service call duration: short - seconds to minute, medium - 1 minute to 30 minutes, and large - 30 minutes to hours.

(a) Moving Pictures (Video) and Sound

(a.1) Video Mail Service

Applications

- Electronic mailbox services for transfer of moving pictures and accompanying sound

Possible Attributes

- Demand
- Point-to-point/multipoint

- Bidirectional asymmetric/unidirectional

(b) Document

(b.1) Document Mail Service

Applications

- Electronic mailbox services for mixed documents

Possible Attributes

- Demand
- Point-to-point/multipoint
- Bidirectional asymmetric/unidirectional

Characteristics of sample messaging services are provided in Table 2-2.

Table 2-2. Characteristics of Sample Messaging Services [2-1] [2-4]

SAMPLE SERVICES	BIT RATE (bit/s)	INFORMATION DENSITY	HOLDING TIME	PERFORMANCE CRITERIA
Electronic mail	300 - 1M	low - 1.0	short, medium	
Voice mail	64K	0.4 - 0.5	medium	delay
Image mail	64K	0.8 - 1.0	medium,large	delay/error rate

Note: Short - seconds to minute, medium - 1 minute to 30 minutes, and large - 30 minutes to hours

2.1.1.3 Retrieval Services: Extraction of Information from a Database

Retrieval service is a complimentary of the messaging service. Retrieval services retrieve information stored in information centers that is provided for public use. The information is sent to the user on a demand only basis. The information can be retrieved on an individual basis. The time at which an information sequence starts is under the control of the user. Since the user is able to recover parts of the loss (using retransmission), the bit error rate requirement is generally less than the messaging services. However, the retrieval services require a shorter delay compared to messaging services. A number of services and applications are provided as follows.

(a) Text, data, graphics, sound, still images, moving pictures

(a.1) Broadband Videotex

The videotex is an interactive system designed to serve the needs of both home and business. The videotex providers maintain a variety of

databases. Information is provided in the form of pages of text and simple graphics. Broadband videotex is an enhancement of the existing videotex system. It includes text, graphics, sound, images, and short film.

Applications

- Videotex, including moving pictures
- Consultancy, ordering, and booking
- Remote education and training
- Telesoftware
- Tele-shopping
- Tele-advertising
- News retrieval

Possible Attributes

- Demand
- Point-to-point
- Bidirectional asymmetric

(a.2) Video Retrieval Service

Within this service, a user can order films or videos from a film/video library facility.

Applications

- Entertainment
- Remote education and training
- Telemedicine

Possible Attributes

- Demand/reserved
- Point-to-point/multipoint
- Bidirectional asymmetric

(a.3) High-Resolution Image Retrieval Service

Applications

- Entertainment
- Remote education and training
- Professional image communications
- Medical image communication

Possible Attributes

- Demand/reserved
- Point-to-point/multipoint
- Bidirectional asymmetric

(a.4) Document Retrieval Service

Applications

- Mixed documents retrieved from information centers, archives, etc.
- Electronic publishing
- Word processing

Possible Attributes

- Demand
- Point-to-point/multipoint
- Bidirectional asymmetric

(a.5) Data Retrieval Service

Possible Application

- Telesoftware and software engineering

Note: Mixed document is a document that may contain text, graphics, still and moving picture information, and voice annotation.

Characteristics of of sample retrieval services are provided in Table 2-3.

Table 2-3. Characteristics of Sample Retrieval Services [2-3][2-4]

SAMPLE SERVICES	BIT RATE (bit/s)	INFORMATION DENSITY	HOLDING TIME	PERFORMANCE CRITERIA
Audio (CD quality)	~1.4M	0.8 - 1.0	medium, large	delay/error rate
Video information retrieval	30M - 45M	low - 1.0	medium, large	delay
High-resolution X-ray film	50M - 100M	low - 1.0	medium, large	delay/error rate
Entertainment video	50M - 800M	0.8 - 1.0	medium, large	delay

Note: Short - seconds to minute, medium - 1 minute to 30 minutes, and large - 30 minutes to hours.

2.1.1.4 Signaling Aspect for Interactive Services

The requirements for the signaling in the interactive services are stated as follows:

- Signaling mechanisms should be service independent.
- Signaling mechanisms should support various services as mentioned above.
- Several call attributes need to be negotiated during call setup and possibly during a call. These attributes are quality-of-service parameters, bit rates for constant bit rate (CBR) and variable bit rate (VBR) services, and virtual channel identifier/virtual path identifier (VCI/VPI) and other asynchronous transfer mode (ATM) layer parameters.
- Signaling needs to support the interworking requirements.

2.1.2 Distribution Services: One-Way Transmission with No Backward Channel

Distribution services are one-way information transfer from a service provider to a user.

2.1.2.1 Distribution Services without User Individual Presentation Control

These services include broadcast services and provide a continuous flow of information distributed from a central source to an unlimited number of authorized receivers connected to the network. The user can access this flow of information without needing to determine at which instant the distribution of a string of information will be started. The user cannot control the start and order of the presentation of the broadcast information. Depending on the point of time of the user's access, the information will not be presented from its beginning. Present day broadcast TV and radio broadcast belong to this class. A number of services and applications are provided as follows.

(a) Video

(a.1) Existing quality TV Distribution Service (PAL, SECAM, NTSC)

Applications

- TV program distribution

Possible Attributes

- Demand(selection)/permanent
- Broadcast
- Bidirectional asymmetric/unidirectional

(a.2) Extended-Quality TV Distribution Service, Enhanced Definition TV Distribution Service, and High-Quality TV Service

Applications

- TV program distribution

Possible Attributes

- Demand (selection)/permanent
- Broadcast
- Bidirectional asymmetric/unidirectional

(a.3) High-Definition TV Distribution Service

Applications

- TV program distribution

Possible Attributes

- Demand (selection)/permanent
- Broadcast
- Bidirectional asymmetric/unidirectional

(a.4) Pay TV (Pay-Per-View, Pay-Per-Channel)

Applications

- TV program distribution

Possible Attributes

- Demand (selection)/permanent
- Broadcast/multipoint
- Bidirectional asymmetric/unidirectional

(b) Text, Graphics, Still Images

(b.1) Document Distribution Service

Applications

- Electronic newspaper
- Electronic publishing

Possible Attributes

- Demand (selection)/permanent
- Broadcast/multipoint
- Bidirectional asymmetric/unidirectional

(c) Data

(c.1) High-Speed Unrestricted Digital Information Distribution Service

Applications

- Distribution of unrestricted data

Possible Attributes

- Permanent
- Broadcast
- Unidirectional

(d) Moving Pictures and Sound

(d.1) Video information distribution service

Applications

- Distribution of video/audio signals

Possible Attributes

- Permanent
- Broadcast
- Unidirectional

Characteristics of sample distribution services without user individual presentation control are provided in Table 2-4.

Table 2-4. Characteristics of Sample Distribution Services without User Individual Presentation Control [2-3]

SAMPLE SERVICES	BIT RATE (bit/s)	INFORMATION DENSITY	HOLDING TIME	PERFORMANCE CRITERIA
Video distribution	30M - 45M	low - 1.0	medium, large	
NTSC, PAL, SECAM	20M - 45M	low - 1.0	medium, large	
EQTV	100M - 150M	low - 1.0	medium, large	error rate
HDTV	200M - 300M	low - 1.0	medium, large	error rate

Note: Short - seconds to minute, medium - 1 minute to 30 minutes, and large - 30 minutes to hours

2.1.2.2 Distribution Services with User Individual Presentation Control

These services distribute information from a central source to a large number of users. The information is provided as a sequence of information entities with cyclical repetition. The user can individually access the cyclical distributed information and can control the start and order of the presentation. Due to the cyclical repetition, the information entities selected by the user will always be presented from the beginning. A number of services and applications are provided as follows.

(a) Text, graphics, sound, still images

(a.1) Full channel broadcast videotex

Applications:

- Remote education and training
- Tele-advertising
- News retrieval
- Telesoftware

Attributes:

- Permanent
- Broadcast
- Unidirectional

Characteristics of sample distribution services with user individual presentation control are provided in Table 2-5.

Table 2-5. Characteristics of Sample Distribution Services with User Individual Presentation Control [2-4]

SAMPLE SERVICES	BIT RATE (bit/s)	INFORMATION DENSITY	HOLDING TIME	PERFORMANCE CRITERIA
Cable image	150 M	low	short	delay
Cable text	2 M	low	short	delay

Cable text uses a full digital broadband television channel for cyclical transmission of text, images, audio and video.

Note: Short: seconds to minute; medium: 1 minute to 30 minutes; large: 30 minutes to hours.

2.1.2.3 Signaling Aspects for Distribution Services

Signaling in distribution services must support frequent and simultaneous requests by a large number of users. Hence, the procedure to access the only meta-signaling must be defined. Meta-signaling procedures are used to establish, check, and release the point-to-point and selective broadcast signaling virtual channel connections. Meta-signaling has a standard virtual path identifier (VPI) and virtual channel identifier (VCI) value.

2.2 Network Aspects of Broadband Services

This subsection provides some of the important aspects when supporting and providing B-ISDN services.

2.2.1 Multimedia Aspect

The multimedia services in the B-ISDN need the following features:

- Flexibility for the user.
- Simplicity for the network operator.
- Ease of service interworking between N-ISDN and B-ISDN.
- Commonality of terminal and network components.

Based on the features above, it is also necessary to standardize a limiting set of information types.

To achieve high throughputs for real-time voice and video information, fast packet switching is needed for the satellite networks. To integrate all of these services, satellite links must be used more efficiently. Demand assignment multiple access (DAMA) is necessary for efficient satellite link transmission.

2.2.2 QOS Aspects (I.350)

In this subsection, the quality of service (QOS) requirements for ATM networks are discussed. QOS is negotiated during the call setup phase and possibly during a call. First the well known QOS subattributes that are not ATM specific are listed, followed by the ATM specific QOS subattributes.

Common subattributes are the following:

- BER of the cell information field (dependent on the transmission links)
- End-to-end transfer delay (dependent on the delay on the transmission medium)
- Connection setup delay
- Release delay

ATM specific sub-attributes are the following:

- Cell loss ratio (CLR): the ratio of the number of lost cells to the sum of the number of lost and successfully delivered cells.

- Cell insertion ratio (CIR): the ratio of the number of misinserted cells to the number of misinserted and successfully delivered cells.

- Delay jitter (dependent on delay variations in the network nodes)

The proposed quality of some services [2-5] is provided in Table 2-6.

Table 2-6. Proposed Quality of Services [2-5]

SERVICE	BER	CLR	CIR
Telephony	10^{-6}	10^{-3}	10^{-3}
Data transmission (2 Mbit/s - 10 Mbit/s)	10^{-8}	10^{-6}	10^{-6}
Video telephony/video conference	10^{-6}	10^{-8}	10^{-8}
Broadcast video	10^{-6}	10^{-8}	10^{-8}
HiFi Stereo	10^{-6}	10^{-7}	10^{-7}
User signaling/remote process control	10^{-5}	10^{-3}	10^{-3}

These values assume random error performance and they were obtained by experiments with existing video and voice codecs. Note that the BER is specified mainly for the link performance and the CLR and CIR are specified based on the terminal and network performance; hence, they are not necessarily correlated.

Except for data transmission, video telephony, video conference, and broadcast video, the satellite can provide satisfactory transmission services. For those high-requirement services, a larger high-power amplifier (HPA) at earth stations, a more robust coding scheme, on-board regeneration, and on-board coding/recoding are needed in the satellite network to perform a satisfactory job.

For high-quality audio and video, a better FEC is needed. Bit interleaving is also necessary for reducing bursty errors.

The schemes for packet multiplexing, bandwidth allocation, and congestion control need to be investigated to ensure fairness and efficient resource usage.

Traffic has tolerance difference to queueing delay, buffering, and bandwidth requirements. The following traffic types can be identified:

- Delay-sensitive high-bandwidth services: This type of services requires a fixed high bandwidth for the duration of a call and demands a real-time service. Examples: conference video, real-time image and document retrieval, and LAN interconnection.

- Delay-insensitive high bandwidth services: Examples: delay-tolerant document, image, and video delivery services.

- Delay-sensitive low-bandwidth statistically multiplexed services: Examples: packetized voice, interactive data, and inquiry-response.

2.2.3 Service Bit Rate Aspect

To support a different bit rate service, first the bit rate expression in the signaling from the user must be specified. This expression must also be simple. After the bit rate expression is received by the network, the network resources must be allocated according to the demand and the network resources have to be utilized efficiently.

2.2.3.1 Constant Bit Rate (CBR) Services

In these services, it is recommended that constant bit rates are negotiated at the call setup time and the necessary network resources are fully allocated for the duration of the call. Changing bit rates during the call is possible by using a signaling technique. A number of specific bit rates must be standardized.

2.2.3.2 Variable Bit Rate (VBR) Services

Variable bit rates can be expressed by several parameters. Candidates for the parameters are a peak bit rate measured over a short period of time (maximum peak bandwidth) and a mean bit rate measured over a longer period of time (average bandwidth and periodicity for data burstiness description). These traffic parameters are negotiated at call setup time. After the negotiation, the network resources are provided to meet the requirements for these parameters for the duration of a call. Changing bit rates during the call is possible by using a signaling technique.

2.2.4 Service Timing/Synchronization

The service timing requirements are based on end-to-end service information and facilities available from the network. Two methods can be used to provide timing as discussed below.

2.2.4.1 End-to-End Methods

Some services require end-to-end service timing methods. Some examples of end-to-end methods are provided as follows.

- The receiver writes the received information field in an elastic buffer with a recovered clock and reads it out with a local clock. The filling level of the elastic buffer controls the frequency difference between the transmitter local clock and the receiver local clock.

- The transmitter uses a (periodic) synchronization pattern in the information field to synchronize the local clock on the receiver side.

- The transmitter writes an explicit time indication in the information field. The receiver uses the time indication to synchronize the local clock.

2.2.4.2 Network Methods

In these methods, the network provides timing and synchronization information to the user. One example is provided as follows.

- The network provides time-stamped cells, i.e., the network writes an explicit time indication in the information field. The receiver uses the time indication to synchronize the local clock.

2.2.5 Maximum Service Bit Rate Supported by the 155.52 Mbit/s Interface

The transfer capability available with the 155.52 Mbit/s interface is 149.76 Mbit/s. This transfer capacity can be easily derived using SDH STM-1 frame structure. Note that ATM 155.52 Mbit/s interface and SDH STM-1 frame support the same transfer capability. There are 270 bytes in STM-1 frame and there are 4 bytes of overhead (section overhead, administrative unit-4 pointer, and path overhead); hence, the transfer capability is 155.52 Mbit/s * 260/270 = 149.76 Mbit/s. Since the cell format allocates 5 bytes for the header and 48 bytes for the information, the maximum information transfer rate is 135.63 Mbit/s. This information capacity is used to accommodate different kinds of cells, which include the following:

- cell information (48 bytes per cell)
- adaptation layer overhead
- signaling cell overhead
- OAM cell overhead
- idle cells
- time stamped cells.

2.2.6 Maximum Service Bit Rate Supported by the 622.08 Mbit/s Interface

The maximum service bit rate supported by the 622.08 Mbit/s interface is currently being studied by CCITT.

2.2.7 Connectionless Data Service

A connectionless data service uses connectionless data transfer protocol to support data transfer between users. Since ATM is a connection oriented technique, there are two ways of supporting a connectionless service using the B-ISDN.

a. Indirectly via a B-ISDN connection oriented service: A permanent, reserved, or on demand connection of the ATM layer is used between B-ISDN interfaces. The connectionless service and adaptation layer functions are implemented outside the B-ISDN. The B-ISDN has no constraints on the connectionless service.

b. Directly via a B-ISDN connectionless service: The B-SIDN connectionless service function terminates the connectionless protocols and routes cells to a destination user according to VPI/VCI information in user cells.

2.2.8 Service Interworking Aspect

A service available from an N-ISDN interface should also be available from B-ISDN. The service interworking capabilities provided between N-ISDN and other existing networks should remain in B-ISDN.

2.3 B-ISDN Performance

B-ISDN performance adopts a layered model, which is categorized into higher layer performance, AAL layer performance, ATM layer performance, and physical layer performance based on the Draft Recommendation I.35B. In this section, only the ATM layer performance issues will be discussed. ATM layer performance is measured by observing the flow of ATM cells and signaling messages. Two classes of ATM layer performance issues can be defined. One is the connection processing performance parameters and the other is the ATM performance parameters. The connection processing performance in ATM networks is based on the connection control signals of the network-user interface. Since this study focuses on cell performance instead of signaling performance, only the ATM performance parameters will be discussed.

- **Cell Loss Ratio (CLR).** CLR is the ratio of the number of lost cells to the sum of the number of lost and successfully delivered cells.
- **Cell (Mis)Insertion Rate.** Cell (mis)inserted rate is the number of (mis)inserted cells within a specified time interval.
- **Cell Error Ratio.** Cell error ratio is the ratio of errored cells to the number of successfully delivered cells.
- **Cell Transfer Delay (CTD).** CTD is the time interval between the occurrence of two corresponding cell reference events at two customer reference equipments.
- **Mean Cell Transfer Delay.** Mean cell transfer delay is the arithmetic average of a specified number of cell transfer delays.

- **Cell Delay Variation**. Cell delay variation is the difference between a single observation of cell transfer delay and the mean cell transfer delay on the same connection.

- **Cell Transfer Capacity.** Cell transfer capacity is the maximum possible number of successfully delivered cell outcomes occurring over a specified ATM connection during a unit of time.

Some ATM performance parameters may be adjusted for a specified ATM connection to support a negotiated QOS. For example, using the cell loss priority (CLP) bit in the ATM cell header allows one to choose between two values of CLR. In the 5-byte ATM cell header, there is 1 bit available for cell loss priority. A cell with CLP value 0 (CLP is not set) has a higher priority than a cell with CLP 1 (CLP is set). When network is congested, the cell with CLP value 0 can pass the network and the cell with CLP 1 will be dropped.

2.4 B-ISDN Frame Structures

This section describes ATM cell and SDH/SONET frame structures as currently being defined by CCITT.

2.4.1 B-ISDN/ATM

Broadband ISDN (B-ISDN) with its flexibility in bandwidth allocation for various traffic types, channel structures, and synchronization rate is a promising technique to accommodate current and future diverse services and traffic.

The functions of B-ISDN can be divided into three planes: a user plane,a control plane, and a management plane (see Figure 2-2). The user plane handles the user information

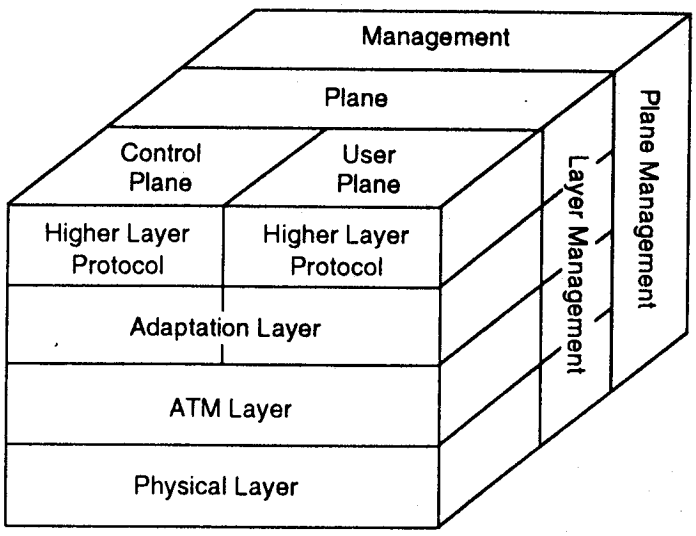

Figure 2-2. B-ISDN Protocol Reference Model

flow and supports flow control, error control, verification and retransmission functions. The control plane handles the signaling flow. The management plane consists of two parts: plane management functions and layer management functions. The plane management functions handle system management and provide coordination among all the planes. The layer management functions handle the operation, administration, and maintenance (OAM) information flow.

The user plane, control plane, and the layer management functions adopt a layered architecture model base on the principles of the ISO's OSI layered model. The layered model is divided into the higher layered protocols, the adaptation layer (AAL), the asynchronous transfer mode (ATM) layer, and the physical layer. The functions provided by each layer are summarized in Figure 2-3.

AAL performs cells assembly/disassembly and individual service dependent functions. In order to reduce the number of AAL services, CCITT has defined four classes of AAL service and all the applications are mapped into one of these four services.

The four classes of services are defined according to the parameters: time relationship between source and destination, bit rate (constant or variable), and connection mode (connection oriented or connectionless). These four classes are:

- Class 1: circuit emulation
- Class 2: variable bit rate video
- Class 3: connection oriented data
- Class 4: Connectionless data transfer

ATM has been proposed to be the optimal transfer mode to provide the B-ISDN services. ATM is a connection-oriented technique. ATM is a low-delay and packet statistical multiplexing scheme in which user information is divided into fixed-length units called cells and within each cell there is a cell header.

ATM is service-independent and performs cell transfer and switching for various B-ISDN services. Transfer capacity is assigned in demand. Signaling and user information are carried on separate channels. The cell header represents the functions supported by the ATM layer. The ATM cell consists of a 5-byte header and a 48-byte information field (cell payload). The cell formats are different at user network interface (UNI) and network node interface (NNI).

UNI Cell Header Format

The cell header at UNI consists of the following fields (see Figure 2-4):

- Generic Flow Control (GFC). There are four bits available for this field.
- Routing Field: Virtual Path Identifier/Virtual Channel Identifier (VPI/VCI). There are 24 bits available for routing: 8 bits for VPI and 16 bits for VCI.

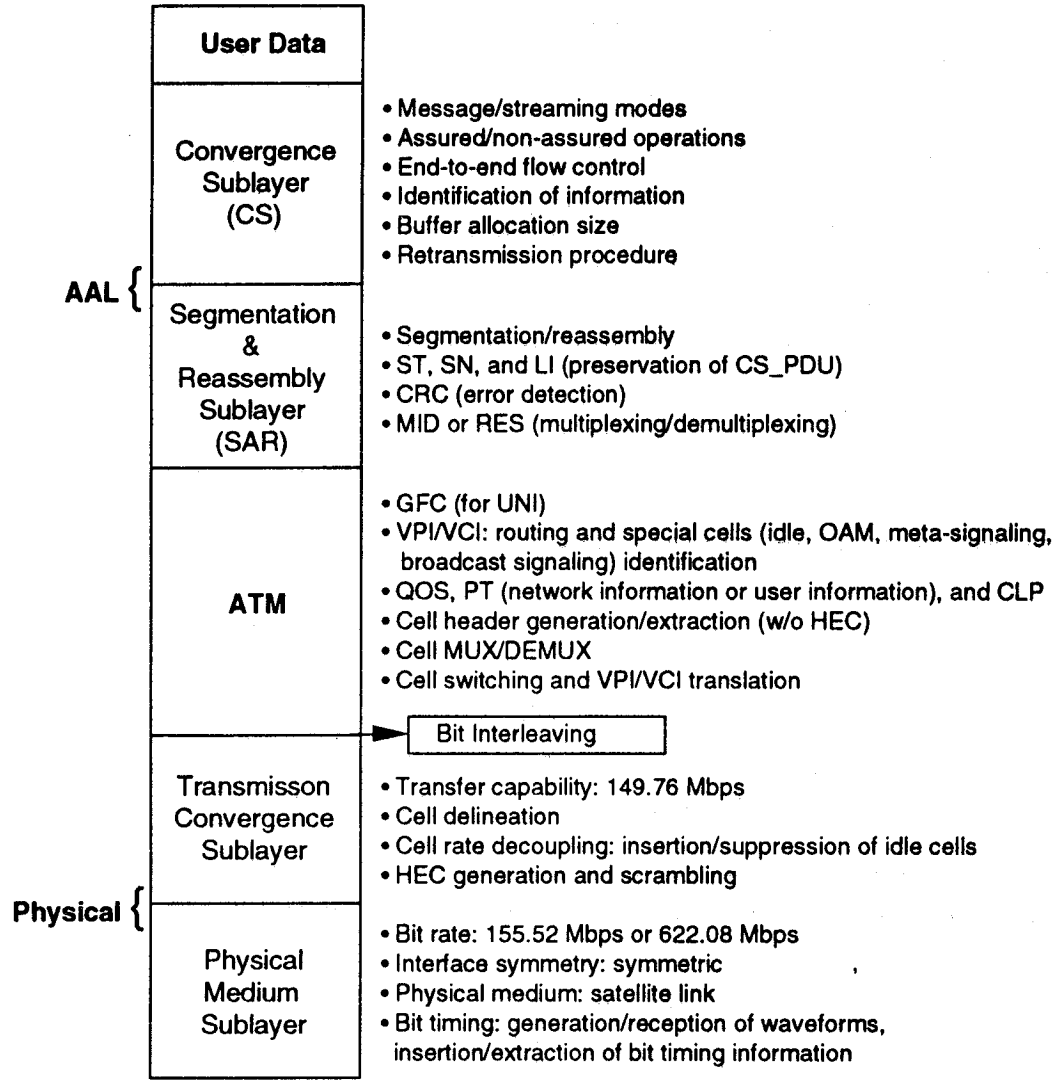

Figure 2-3. Function Requirements for Different Layers

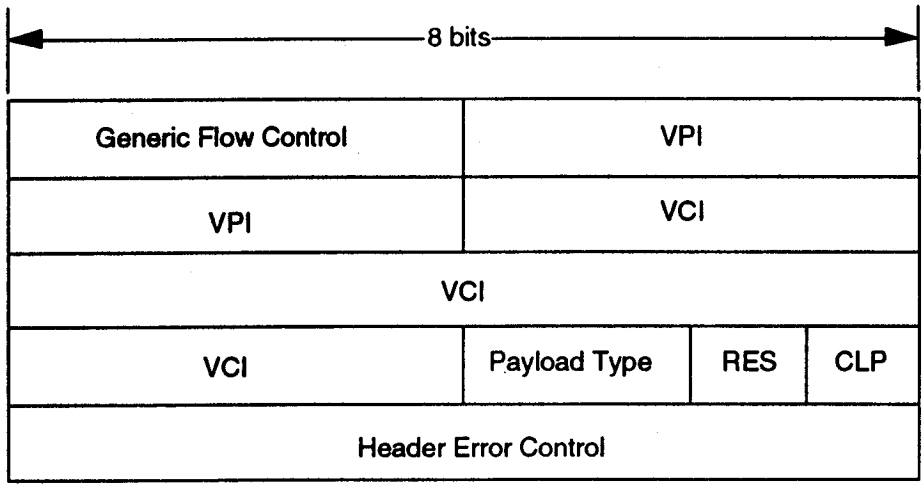

Figure 2-4. UNI ATM Cell Format

- Payload Type (PT). There are 2 bits which are used to indicate whether the cell information field contains user information or network information.
- Cell Loss Priority (CLP). There is 1 bit available for cell loss priority. A cell with CLP value 0 (CLP is not set) has a higher priority than a cell with CLP 1 (CLP is set).
- Header Error Control (HEC). There are 8 bits available for header error control.
- Reserved Field (RES). There is 1 bit available for reserved field.

NNI Cell Header Format

The cell header at NNI consists of the following fields (Figure 2-5):

- Routing Field: Virtual Path Identifier/Virtual Channel Identifier (VPI/VCI) - There are 28 bits available for routing: 12 bits for VPI and 16 bits for VCI.
- Payload Type (PT) - There are 2 bits which are used to indicate whether the cell information field contains user information or network information.
- Cell Loss Priority (CLP) - There is 1 bit available for cell loss priority. A cell with CLP value 0 (CLP is not set) has a higher priority than a cell with CLP 1 (CLP is set).
- Header Error Control (HEC) - There are 8 bits available for header error control.
- Reserved Field (RES) - There is 1 bit available for reserved field.

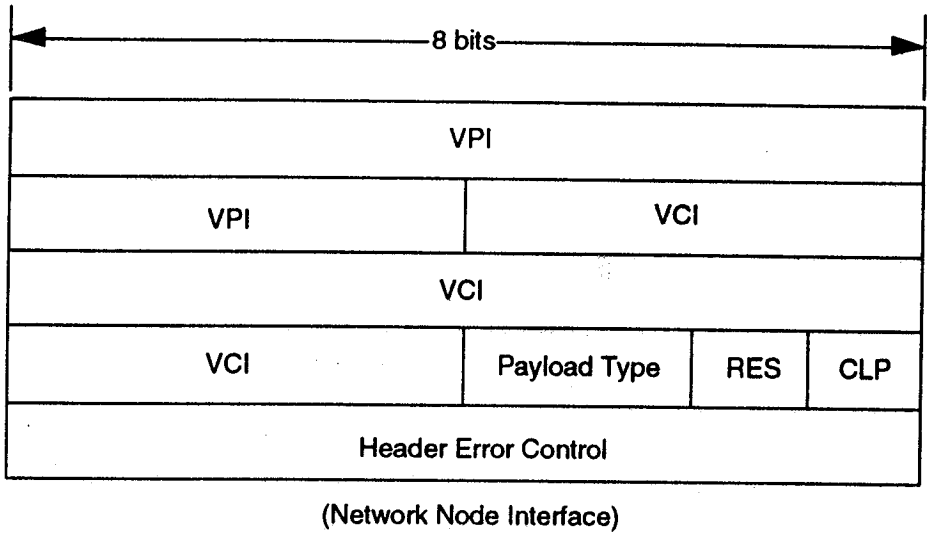

Figure 2-5. NNI ATM Cell Format

2.4.2 SDH/SONET

Two synchronous digital hierarchy (SDH) rates have been defined for ATM in CCITT G.707-G.709: 155.52 Mbit/s and 622.08 Mbit/s. The basic building block of SDH is synchronous transport module level 1 (STM-1) signal of which the bit rate is 155.52 Mbit/s. The STM-1 signal is shown in Figure 2-6. The STM-1 signal in SDH is

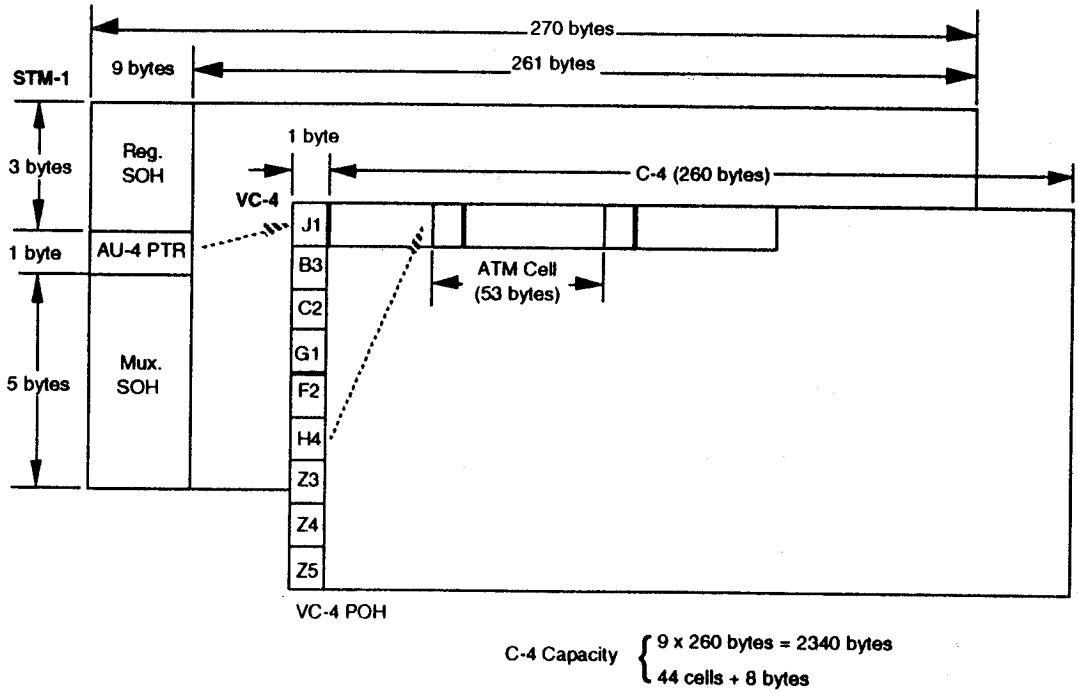

Figure 2-6. STM-1 Frame Format

equivalent to synchronous transport signal level 3 concatenated (STS-3c) signal in SONET. The STM-4 signal is obtained by byte interleaving four STM-1 signals. The bit rate for the STM-4 signal is 622.08 Mbit/s.

The STM-1 payload may consist of circuit slots or ATM cells. The multiplexing scheme for the ATM cell case is introduced below as an example (see Figure 2-7). The ATM cells will first be mapped into a C-4 container. Then the path overhead (POH) will be aligned with the C-4 container to form the VC-4 virtual container. The VC-4 is combined with the administrative unit pointer to form the administrative unit (AU-4). The STM-1 frame is formed by multiplexing the section overhead (SOH) and AU-4.

Two kinds of overheads are used for supervisory and maintenance operation in SDH: the path overhead (POH) and the section overhead (SOH). The SOH and AU pointer format are shown in Figure 2-8.

The SDH overhead processing functions include the following:

- framing
- performance monitoring
- end-to-end performance monitoring
- data communication channels
- user channels
- path-user channel
- automatic protection switching
- orderwire

The path overhead and AU pointer functions are listed in Tables 2-7 and 2-8. The section overhead functions are listed in Table 2-9.

SONET, originally proposed by Bellcore, is a standard optical interface. The basic building block and first level of the SONET signal hierarchy is called the synchronous transport signal level 1 (STS-1). The STS-1 has a bit rate of 51.84 Mbit/s. The frame structure is shown in Figure 2-9. The entire frame is transmitted every 125 μs and each byte represents a 64-kbit/s DS0 channel. The STS-1 building block consists of transport overhead (section overhead and line overhead), path overhead, and the information payload. There are 9-byte section overhead, 18-byte line overhead, and 9-byte path overhead.

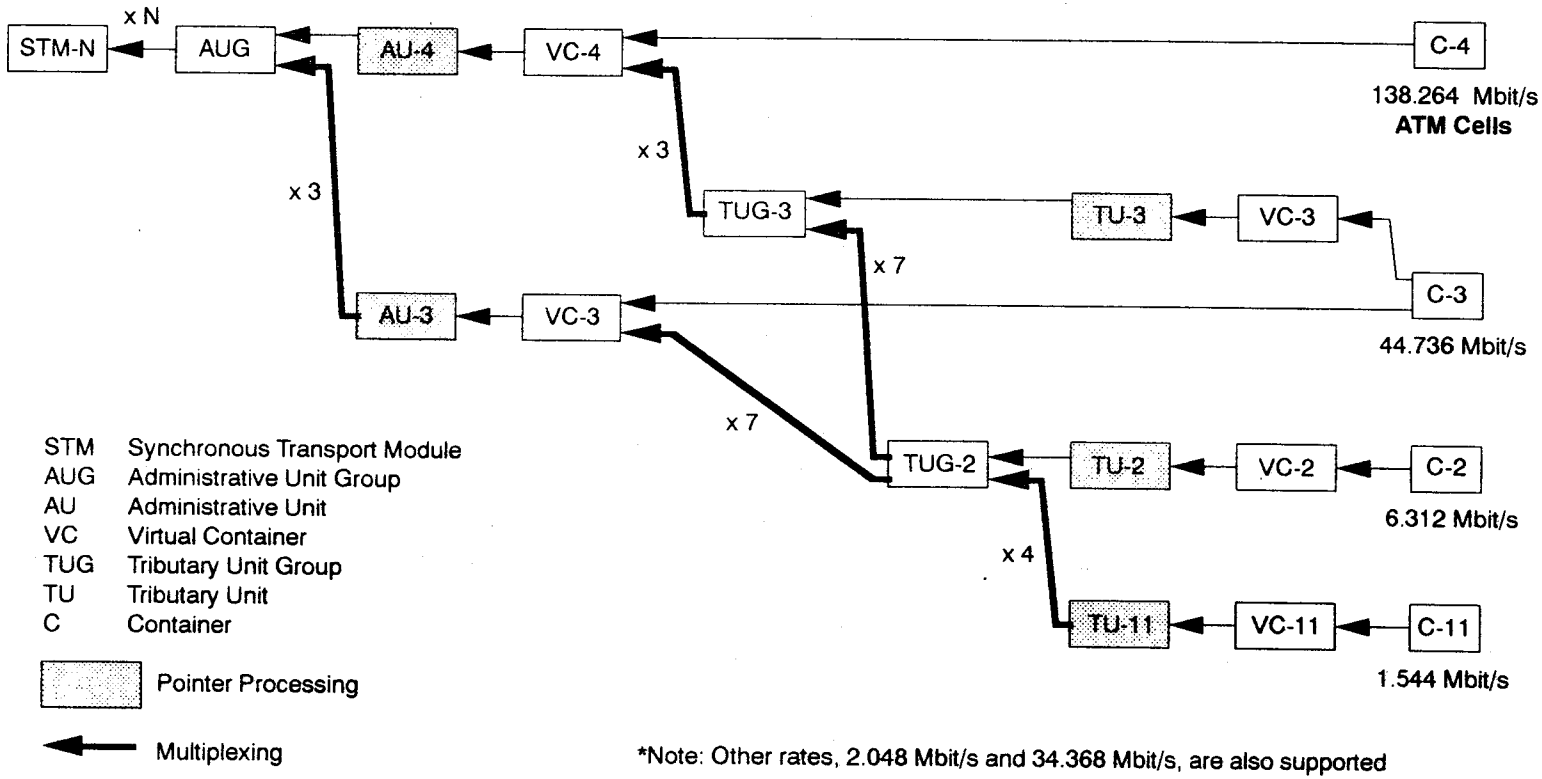

Figure 2-7. SDH Multiplexing Structure

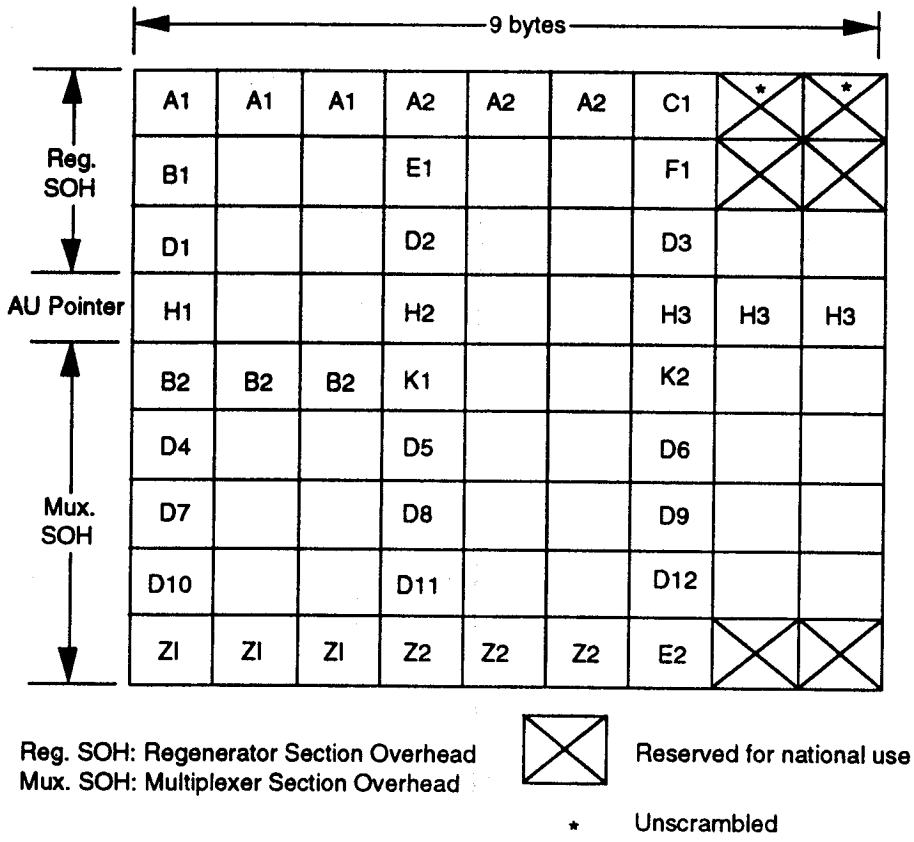

Figure 2-8. Section Overhead and AU Pointer Format

Table 2-7. Path Overhead Functions

ELEMENT	ABBRV.	DESCRIPTION
VC-4 Path Trace	J1	Verifies the VC-4 path connection
Path BIP-8	B3	This byte contains the even parity of all bytes of the previous VC-4
Signal Label	C2	Indicates the composition of the VC-4
Path Status	G1	Contains the path termination status and performance
Path User Channel	F2	Contains user communication between path elements
Multiframe Indicator	H4	Generalized multiframe indicator with payload specific application. Bits 3-8 are for cell offset indicator
Spare	Z3 - Z5	Reserved for future use

Table 2-8. AU Pointer Functions

ELEMENT	ABBRV.	DESCRIPTION
Pointer Value	H1, H2	Indicates the offset between the pointer and the first byte of the VC-4 POH
Pointer Action	H3	Three bytes are for frequency justification

Table 2-9. Section Overhead Functions

ELEMENT	ABBRV.	DESCRIPTION
Framing	A1, A2	Provides framing
STM identifier	C1	Contains value of the multi-column, interleave depth coordinate
Data Com. Channel(DCC)	D1-D12	Provides a 192 and a 576 Kbit/s data communication channel
Orderwire	E1, E2	Provides voice orderwires
User Channel	F1	Reserved for user purposes
BIP-8	B1	Provides regenerator section error monitoring
BIP-Nx24	B2	Provides multiplex section error monitoring
APS Channel	K1, K2	Automatic protection switching channel
Spare	Z1, Z2	Spare

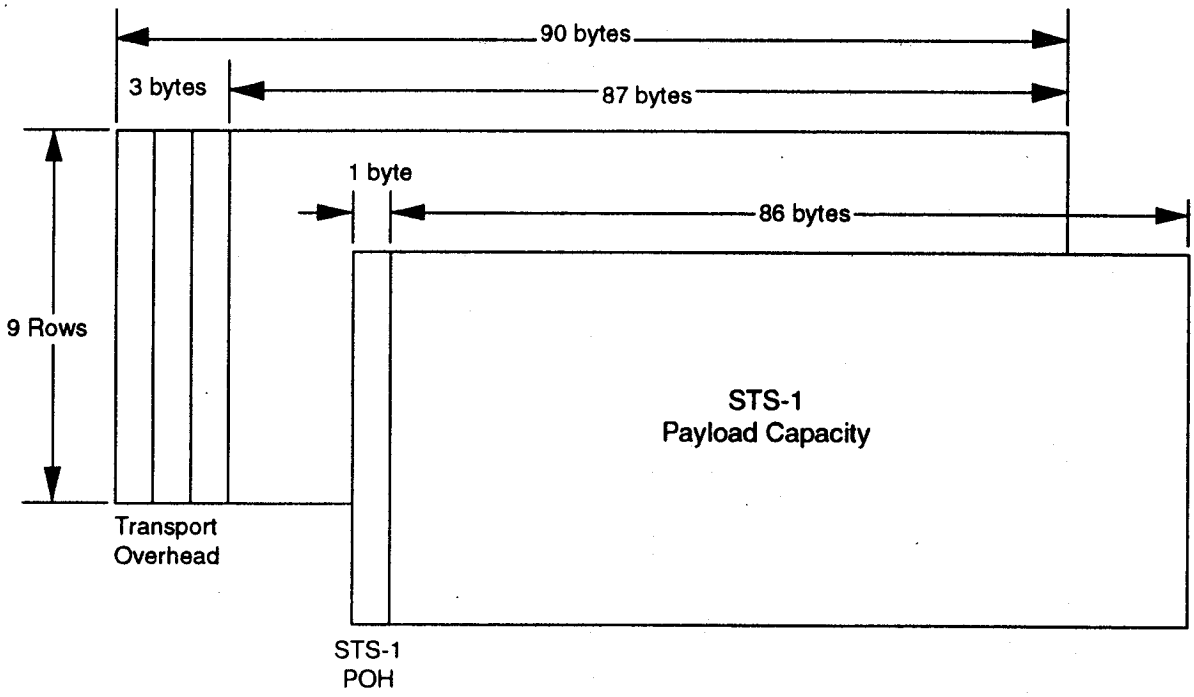

Figure 2-9. STS-1 Frame Format

An STS-N (level N) signal contains N copies of the an STS-1 signal. The STS-Nc (concatenated) signal only contains one copy of the path overhead. The 155.52 Mbit/s signal is carried by an STS-3c frame. An STS-3c signal provides a payload of 149.76 Mbit./s. The STS-3c transport overhead structure is shown in Figure 2-10. The overhead functions have been introduced in Tables 1, 2, and 3. The 622.08 Mbit./s signal can be carried either STS-12 or STS-12c. The STS-12 is constructed by multiplexing (byte interleaving) four STS-3c frames. Four sets of the path overhead are independently remained in the STS-12 to get access to different STS-3c frames. For the STS-12c only one set of path overhead is existed.

The transmission rates of B-ISDN (ATM/SDH) specified by CCITT and SONET specified by ANSI are listed in Table 2-10.

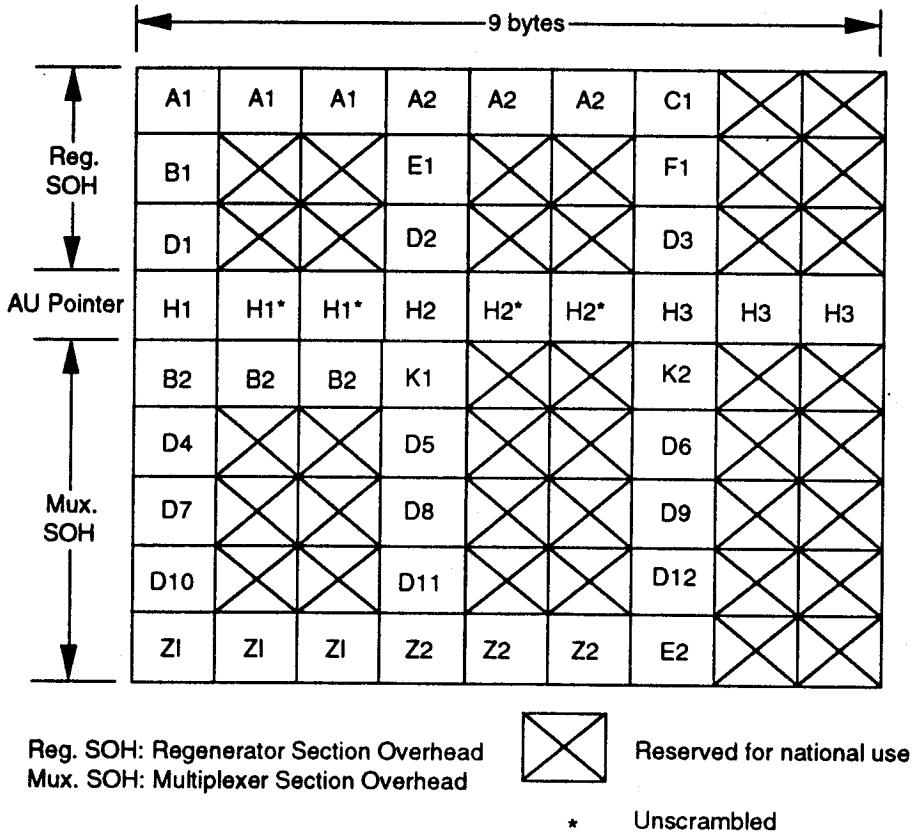

Figure 2-10. STS-3c Transport Overhead Format

Table 2-10. B-ISDN and SONET Hierarchy and Bit Rates

STANDARD	INFORMATION STRUCTURE	BIT RATE (Mbit/s)
B-ISDN (ATM & SDH)	ATM/STM-1	155.52
	ATM/STM-4	622.02
	STM-16	2488.32
SONET	STS-1/OC-1	51.84
	STS-3/OC-3	155.52
	STS-9/OC-9	466.56
	STS-12/OC-12	622.08
	STS-18/OC-18	933.12
	STS-24/OC-24	1244.16
	STS-36/OC-36	1866.24
	STS-48/OC-48	2488.32

2.5 Reference

[2-1] J. C. Kohli, D. S. Biring, and G. L. Raya, "Emerging Broadband Packet-Switch Technology in Integrated Information Networks," IEEE Network, pp. 37-51, Nov. 1988.

[2-2] W. E. Falconer and J. A. Hooke, "Telecommunications Services in the Next Decade," Proceedings of The IEEE, vol. 74, no. 9, pp. 1246-1261, Sep. 1986.

[2-3] F. Kishino, K. Tokikuni, Fumio Kondo, and K. Takemoto, "Overview of Broadband Telecommunications Services - Trial Experienced in Japan," IEEE GLOBECOM, pp. 1990-1994, 1987.

[2-4] D. J. Wright and M. To, "Telecommunication Applications of the 1990s and their Transport Requirements," IEEE Network Magazine, pp. 34-40, March 1990.

[2-5] CCITT Study Group XVIII - Report R4 1989.

Section 3

Current Development Status

3.1 Introduction

Two approaches to provide B-ISDN multirate and multimedia services have been proposed and are being developed. The first approach is based on the ATM technology, while the second uses the synchronous optical network (SONET) technology. Since SONET has been successfully defined and developed by Bellcore, the U.S. industry is apt to use both SONET and ATM technologies to implement the B-ISDN. The first phase will use SONET to provide B-ISDN services, and then SONET can be used as a stepping stone for the ATM technology, since it can accommodate both circuit slots and ATM cells. It is envisaged that SONET and ATM will coexist in the future to provide B-ISDN in North America. European countries will provide B-ISDN services using the synchronous digital hierarchy (SDH) and ATM technology. Before fully implementing B-ISDN services, switched multimegabit data service (SMDS) is currently available to support a high performance, connectionless, packet switched data service. SMDS is implemented using current technology based on metropolitan area networks (MANs). The distributed queue dual bus (DQDB) adopted by IEEE 802.6 is the underlying protocol for the implementation of the SMDS. In this section, the current status of the B-ISDN equipment development for SONET, ATM, DQDB, and SMDS is surveyed. Network deployment plans and services offering plans including SMDS and B-ISDN are also addressed.

3.2 SONET Equipment

The basic building block and first level of the SONET signal hierarchy is called the synchronous transport signal-level 1 (STS-1). The STS-1 has a bit rate of 51.84 Mbit/s. As shown in Figure 2-9, the STS-1 frame structure is described as nine rows of 90 bytes. The entire frame is transmitted every 125 μs. The optical carrier-level 1 (OC-1) is obtained from the STS-1 signal after STS-1 has been scrambled and converted from electrical to optical. During the conversion, no extra bandwidth is added to the signal; hence, an OC-1 and an STS-1 have the same rate. An STS-N (level N) signal contains N copies of an STS-1 signal.

3.2.1 SONET Analyzer

ANDO Corporation has developed the AP9460 SONET Analyzer. It supports both STS-1 and STS-3 electrical and OC-1 and OC-3 optical interfaces. It can measure the alarm and maintenance signals and test the PATH, LINE, and SECTION overheads as well as the pointer and BER.

ANRITSU Corporation has developed the MP1560A SONET Analyzer. Besides the features in the ANDO AP9460 SONET Analyzer, the MP1560A handles the various standard mapping among SONET (ANSI), SDH (CCITT), and Japan. The difference among these standards is the multiplexing and mapping method. It has an add/drop capability for DS1 and data communication channels, and it is capable of measuring delay jitter.

3.2.2 SONET Products

Adaptive Corp. has announced a SONET based switch called the SONET Transmission Manager (STM). The first version of the STM will support up to 18 DS-3 trunks, and the total capacity is over 1 Gbit/s. The later versions will support more DS3 ports, a SONET interface, and fast packet switching technology based on ATM.

Alcatel Network Systems has developed a SONET product family including multiplexer, fiber transport, and access systems. The multiplexer products consist of a TM-50 Terminal Multiplexer and an ADM-50 Add/Drop Multiplexer, both operating at 50 Mbit/s. The TM-50 is a DS1-to-STS-1 terminal multiplexer, which grooms and transports DS1 signals within the STS-1 or OC-1 frames. The ADM-50 Add/Drop Multiplexer is a DS1-to-STS-1 SONET add/drop multiplexer. There are several fiber optic transport products. The FTS-150 operates at the OC-3 line rate and multiplexes any combination of up to 3 DS3 or STS-1 signals for transport. The FTS-600 operates at the OC-12 line rate and multiplexes any combination of up to 12 DS3 or STS-1 signals. The FTS-2400 system operates at the OC-48 line rate. The VT/ADM-150 offers DS1 to STS-3 or OC-3 multiplexing and virtual tributary cross-connection. The ADM-600 provides DS3 and STS-1 to OC-12 multiplexing and cross-connection. The LCX-50/150 digital loop carrier cross-connect access system assembles 672 lines of loop traffic into OC-1 or OC-3 line rates. The LCX-600 system assembles loop traffic into the OC-12 line rate. The access products can support either a hub network or a feeder distribution network. The Alcatel also supplies the network management system, which comprises network element management, mediator, and node/network supervisor products.

AT&T announced its 2000 SONET product line, which comprises the DACS III-2000 and DACS IV-2000 digital cross-connect systems, the DACScan-2000 network controller, the DDM-2000 multiplexer, and the FT-2000 lightwave system. The two SONET-compatible cross-connect systems are the DACS III-2000 and DACS IV-2000. DACS III-2000 is a software-based, high-capacity, digital cross-connect system providing efficient automated and remote provisioning and test access. DACS III-2000

can support up to 960 DS3 terminations, which will be expanded to 1,920 DS3 terminations. This cross-connect is a nonblocking switch. DACS IV-2000 provides both electronic cross-connect and multiplex functions for DS3 and DS1 channels. It can support up to 248 DS3 channels, and can be configured for different combination of DS3 and DS1 interfaces. The DACScan-2000 provides interactive enhanced network management capability to the end user, which allows users to maintain network elements down to the DS1 level. The DDM-2000 network multiplexer offers two types of multiplexing: virtual tributary and M13. The FT-2000 lightwave system operates at the OC-12 line rate, which will be upgraded to the OC-48 line rate in 1991.

AT&T Network Systems has announced five frequency-control components which recover timing from information transported in digital systems. These include the TRU-200C Clock Recovery and Data Retiming Module, three Retiming Surface Acoustic Wave (SAW) Filters, and the 157-type VHF Voltage-Controlled Crystal Oscillator, all of which are SONET compatible.

Fujitsu America provides the FLM 600 which operates at the OC-12 line rate and is a fiber loop multiplexer. The FLM 600 can support 12 DS3 signals (8,064 channels). Three versions are available for this multiplexer: terminal, hub, and add/drop multiplexer. The FLM 50/150 are SONET optical transmission systems. The FLM 50 can support up to 28 DS1 signals (672 channels). The FLM 150 can support OC-3 transport of 84 DS1 signals (2,016 channels), 2 independent OC-1 transports with 28 DS1 lines each, or an OC-3/OC-1 hub with DS1 add/drop.

Hekimian Laboratories Inc. provides T::DAX wideband digital cross-connect with SONET capability. This cross-connect supports the STS-1 interface. It provides 1-to-1 transport of DS1 signals, 3-to-1 cross-connect of DS3 and DS1 signals, and virtual tributary 1.5 cross-connect. This system has the advantage of being upgraded to higher SONET speeds without any change to the switch matrix.

LICOM Inc. offers the IMX30 which is an intelligent multiplexer. This multiplexer, operating over digital microwave or optical fiber at the T3 rate, integrates add/drop, DS0 and DS1 cross-connection, embedded maintenance provisioning, and diagnostic functions. The LICOM's Integrated Management and Control System (LIMACS) is designed to enhance the network versatility and performance of the IMX30 product line. LIMACS can collect and summarize the alarm and performance data for the IMX30 systems and digital facilities in report forms.

NEC America has developed the IMT-150 intelligent multiplexer transport system, the ITS-600FL, and the ITS-2400 intelligent transmission systems. The IMT-150 can support over 2,000 subscriber circuits. The ITS-600FL supports the OC-12 line rate, which transports 12 DS3 or 4 OC-3. The ITS-2400 transports up to 48 DS3 or 16 OC-3 over an OC-48 line rate. This company plans to develop the intelligent network management system, the intelligent subscriber carrier, and the wideband and broadband cross-connect systems.

Northern Telecom's FiberWorld S/DMS product family, based on frame relay technology, comprises the S/DMS TransportNode, S/DMS AccessNode, and S/DMS

SuperNode. The components of the S/DMS TransportNode family include the OC-48 and OC-12 transport terminals, SONET transport radio terminal, and broadband manager. The OC-12 and OC-48 transport terminals are both bandwidth managers. The SONET transport radio terminal is a 512 quadrature amplitude modulation (QAM) digital microwave system with a maximum 6 STS-1 signals per channel. The broadband management system is a nonblocking cross-connect system. It supports both optical and electrical interfaces and handles OC-3, OC-12, OC-48, DS1, and DS3 interfaces. The S/DMS AccessNode is a multifunctional system that handles delivery services among FiberWorld products and supports broadband, wideband, and narrowband services. It handles both fiber and copper loops. The S/DMS SuperNode is the broadband switching component, which consists of the service processor layer, signaling and connectivity layer, and physical access layer. The SuperNode also provides OA&M features.

Optilink Corp. has introduced the TR-303 digital loop carrier (DLC) system which features integral SONET single mode fiber transport at OC-1 (51.84 Mbit/s) or OC-3 (155.52 Mbit/s). This system can serve from 200 to 2,000 lines.

Reliance Comm/Tec provides a SONET optical interface which delivers OC-3 to 672 channels.

Rockwell International has announced a family of SONET products. The ROC-3 SONET lightwave transmission system operates at the OC-3 rate and will be upgraded to OC-12 (622.08 Mbit/s) in 1991. The ROC-3 has DS1 add/drop functionality and diagnostics, loopback, and performance monitoring capabilities. The two other SONET products are the RDX-33 broadband digital cross-connect system (B-DCS) and the RDX-31 wideband digital cross-connect (W-DCS). These cross-connect systems can interface with the current asynchronous networks and the synchronous optical network without any changes of the switching fabric. The RDX-33 B-DCS supports up to 2,048 DS3/STS-1 signals. The RDX-31 W-DCS supports 256 DS3, 8192 DS1, or an equivalent combination.

Seiscor Technologies has the SONET-compatible intelligent digital loop carrier system, the FiberTraq. The input can be at the DS1 line rate using facility interface cards or at the OC-3 line rate using integrated SONET interface.

Telco Systems has upgraded its 828AF multiplexer with a SONET-compliant S-828AF which operates at 45 Mbit/s. This multiplexer has DS1, DS2 and DS3 transport capability. The S-600 OC-12 high speed terminal will be introduced by the end of 1991, which can be upgraded to S-2400 OC-48 in the future.

Telecom Solutions provides the digital clock distributor for timing reference within SONET.

Telllabs announced several TITAN Digital Cross-Connect Systems. This cross-connect family embraces DS3-to-DS1, DS1-to-DS1, and DS3-to-DS3 cross-connect functions, acting like a gateway from the asynchronous network to the synchronous optical

network. It also provides the intelligent network management systems, which can be integrated with the cross-connect system.

Timeplex introduces the TX3 SuperHub, which is a DS3-to-DS3/DS3-to-DS1 transport and cross-connect system. It can handle up to 10 DS3 rate signals or a combination of DS1 and DS3 signals not exceeding the equivalent of 360 DS1 signals. Its element management system (EMS) provides network management capabilities including real-time monitoring and control of system alarms, configuration control, troubleshooting, detailed diagnostics at the module level, and statistical reporting.

TRANSWITCH has introduced several DS3/SONET chips, evaluation systems, and boards. The DS3F evaluation board is for evaluating the DS3 Framer device, which frames a DS3 rate signal into a DS3 format signal in either M13 or C-bit format. The DS3RT evaluation board is for evaluating the DS3RT line interface device, which consists of a DS3 signal transmitter and receiver. The DS3M/SOT-1 evaluation system is used for evaluating the DS3 Mapper (which maps a DS-3 signal to STS-1 signal) and SOT-1 (which is a SONET STS-1 overhead terminator). The SONET 3:1 Multiplexer (SM3) chip can multiplex three STS-1 serial signals into one STS-3 signal.

VITESSE Semiconductor Corporation has announced several SONET chips. The VS 8010 series ICs are a set of high speed GaAs multiplexers and demultiplexers for SONET applications. This series comprisesthree chips: the VS8011 multiplexer, the VS8012 demultiplexer, and the frame recovery chip. The VS8021/VS8022 are high-speed SONET interface devices that can handle serial data at rate up to 2.5 Gbit/s. The VS8021 contains an 8:1 multiplexer and a self-positioning timer. The VS8022 contains an 1:8 demultiplexer and SONET frame recovery circuitry. The VS8021/VS8022 are used for STS-3 through STS-48 SONET applications.

3.3 HDTV Codec

HDTV is expected to play an important role in B-ISDN service. A survey of HDTV codecs is provided below.

BellCore has announced two HDTV codecs. The first one operating at 622 Mbit/s is based on the pulse code modulation (PCM) technique. The second one operating at 140 Mbit/s divides the frequency band into six different sub-bands. At each sub-band, DPCM or PCM can be applied to reduce the data rate.

Compression Laboratories, Inc. (CLI) plans to build a variable rate HDTV codec at 140 Mbit/s, 90 Mbit/s, and 45 Mbit/s. The rate reduction technique will possibly use T-DPCM, DCT, and motion compensation.

General Instrument (GI) has announced a 15 Mbit/s codec to accommodate the U.S. Advanced Television (ATV) Standard. The rate reduction algorithm is also based on T-DPCM, DCT, and motion compensation.

KDD/Canon has introduced a 120/140 Mbit/s HDTV codec. This codec uses a sub-Nyquist sampling technique for 2:1 data rate reduction, and uses source coding,

adaptive differential pulse code modulation (ADPCM), for an additional data rate reduction. In 1989, COMSAT, INTELSAT, AT&T, and KDD conducted an international field trial using this codec between the U.S. and Japan through an INTELSAT satellite.

The Telettra HDTV codec supports two European 34 Mbit/s rate channels. The encoding scheme uses a discrete cosine transform (DCT) after temporal differential pulse code modulation (T-DPCM). Telettra will introduce a Codec supporting two DS3 interfaces in 1991.

3.4 ATM Switches

ATM switches use fast packet switching technology to route ATM cells to output lines according to the VPI/VCI address field in the cell header. Different terms have been used for referring to the fast packet switching technology. The United States uses fast packet switching, Europe uses asynchronous time division switching, CCITT uses ATM switching, and AT&T uses wideband packet technology.

AT&T Bell Laboratories has researched several fast packet switches. The first one is the Starlite Packet Switch, which consists of a Batcher's Sorting network, a Trap network and a Banyan network. The Banyan network is an unbuffered switch and, because of the sorting network, there is no internal blocking. The output contention problem of the switch is solved by the Trap network, which recirculates the duplicate address packets to the reentry input ports. The second fast packet switch is a buffered banyan network. There are two buffers for each input within one switching element and the buffer size is one packet. The internal blocking and output contention problems are tackled using this buffering approach. The switch speed is higher than the trunk speed to further reduce the blocking problems within the switch.

AT&T Network Systems has developed the BNS-2000, a cell relay switch based on the ATM protocol. This switch is designed to be used with SMDS, X.25, asynchronous, synchronous, ISDN, frame relay, and other broadband services.

Bellcore has developed a Batcher-Banyan chip set which contains a 32 X 32 banyan chip and a 32 X 32 batcher chip using CMOS technology. The Batcher and Banyan chips have been tested at bit rates of 170 Mbit/s. Based on the above chip set, Bellcore has also built a 32 X 32 prototype fast packet switch that uses the ring reservations scheme to resolve the output contention. Bellcore has also used the 32 X 32 Batcher-Banyan chip set as a basic module to build a 256 X 256 sorted-banyan-based switch with a total capacity of 35 Gbit/s.

Bell South has built an ATM switch prototype. The ATM switch design was originally developed at Bell Core.

Fujitsu Limited has developed a B-ISDN experimental system. The Fetex-150 switch is constructed using multiple stages of self-routing switch modules. This switch uses ATM technology to provide a transmission pipeline of 80 Gbit/s and is able to support thousands of high-speed broadband terminals. The internal blocking and output

contention problems are resolved using the 3-stage configuration (multiple paths between each input and output pair) and a buffered self-routing switch module. The self-routing switching module consists of ATM switching large-scale integrated chips (LSICs). The ATM switching LSICs were developed in Fujitsu Labs. The logic technology used in this IC is Bi-CMOS logic gate array with ECL interface. This system offers different features with different modules. One module called the broadband signal path subsystems offers ATM and STM using different switching fabrics and supports SONET and B-ISDN. SMDS is supported with a separate terminal adaptor.

GTE has developed a broadband circuit switch capable of handling full-motion video and other services over telephone lines.

IBM has developed a fast packet switch prototype capable of supporting more than 1 million packets. This switch supports a 45 Mbit/s line speed and optical interface. Routing is based on the source routing method, where the packet carries end-to-end path information. In this sense, the virtual channel number will not be translated at the intermediate nodes.

NTT Corporation has announced an R&D project for developing the ATM node system and the ATM link system. According to the announcement, the ATM node system executes call processing and transfers the cells. This system can convert transferring nodes necessary for interworking with the existing networks and provide supplementary services. The ATM link system consists of subscriber loops and trunk lines, and can multiplex cells into the transmission line.

Siemens Public Switching Systems, Inc. has announced an ATM switch for transmission of video images as well as text, data, and speech. The line speed is 140 Mbit/s. The switch uses the CMOS logic technology. Commercial ATM switches are expected to be available in early 1991.

Washington University at Saint Louis will develop a broadband packet network using fast packet switching technology. This program will demonstrate real-time applications such as HDTV, interactive classroom learning, and distributed teleconferencing. The program will last for the next three years under the support of Southwestern Bell Telephone Company (SWBT), SBC Technology Resources Inc (STR), and NEC America, Inc. (NECAM).

3.5 MAN DQDB

A metropolitan area network (MAN) is defined as a standardized, discrete, high-speed network for voice, data, and video that provides LAN/LAN or LAN/WAN connections for public or private communication systems on noncontiguous properties within a metropolitan distance range. The QPSX Communications Pty Ltd of Western Australia has developed queued packet and synchronous switch (QPSX) prototypes that use a distributed queues dual bus (DQDB) protocol to support the requirements of a public or a private MAN. This DQDB has been adopted as the IEEE 802.6 standard, which is the interface between users and MAN. The DQDB protocol is a media access control (MAC)

packet switching scheme, which uses fixed length data segments to access the slots on the bus. This fixed length data segment has the same concept as the ATM cell, although the formats of the two are different. The DQDB is planned to migrate to and support ATM.

3.5.1 DQDP Products

The Alcatel MAN/QPSX consists of a set of equipment clusters based on the DQDB technology. The bus is operated at the DS3 rate and in slot mode. This MAN supports OSI management functions such as configuration, fault detection and correction, performance monitoring, accounting, and network security. The MAN subnetwork can be configured as an open bus or a self-healing looped bus. Alcatel has signed a licence agreement with QPSX of Australia to manufacture and supply the MAN systems. AT&T Network Systems and Siemens have also signed a similar agreement to manufacture the QPSX products.

3.5.2 SMDS Products

Advanced Computer Communications, Digital Link, Ungermann-Bass, Verilink, and Wellfleet have announced an SMDS interface specification in May 1991. The proposed specification for LAN-to-SMDS connectivity allows routes to connect directly to SMDS networks through data service unit and channel service unit products. The standard is based on the high-level data link control (HDLC) protocol.

Advanced Computer Communications has announced an enhancement of its series 4000 products supporting SMDS.

Cisco Systems has announced a router interface that allows Cisco route/bridge products to link remote LANs over SMDS at the DS1 rate (1.5 Mbit/s). A version for the DS3 rate is under development with NEC America, Inc.

Digital Link manufactures data service units (DSUs) and channel service units (CSUs). Using the new DSU/CSU products and routers made by Wellfleet, users can interface to SMDS networks.

3.6. B-ISDN Field Trial

Nynex Corp.'s New England Telephone Co. will use the Fujitsu Ltd. Fetex-150 broadband switch to perform an SMDS trial in Boston. The SMDS will interconnect LANs by providing broadband transmission over public packet-switched networks at rates 45 Mbit/s. This Nynex trial was the first commercial application of the ATM version of the Fetex-150 switch.

GTE and Bellsouth announced a joint experiment broadband telecommunication project in North Carolina that will feature transmission speeds of up to 2.4 Gbit/s and an access rate of 622 Mbit/s using ATM switching and transmission technology. Bellsouth has

selected Fujitsu Ltd. Fetex-150 switch for the broadband service trial, operated at rates up to 622 Mbit/s. Bellsouth will provide an ATM switch and synchronous optical network transmission equipment, while GTE will provide broadband switching technology based on its high-speed circuit switch.

Southwestern Bell Telephone will conduct a trial of SMDS service using the AT&T 2000 series equipment to interconnect LANs.

Pacific Bell and Stanford University, Apple Computer, Tandem Computer, Sun Microsystems, and Pacific Gas and Electric will conduct a field trial of the SMDS service at the end of 1990. Pacific Bell plans to offer the SMDS service in 1992.

NTT has successfully tested its experimental B-ISDN system at the Musashino Research and Development Center. This experimental system consists of an ATM switching system, an ATM ring system, ATM terminals, cell assembly and disassembly devices, and a network simulator. The ATM switching system accommodates 256 lines with a 156 Mbit/s line speed. NTT is planning to have public B-ISDN service in operation by 1995.

AT&T has performed lab tests to measure the performance of SONET network equipments. The lab tests are a preliminary step to test its Network Systems 2000 product line of SONET equipments with the Experimental University Network II (XUNET-II). The members of XUNET-II include the University of California at Berkeley, the University of Illinois and the University of Wisconsin, and Illinois Bell, Wisconsin Bell, Bell Atlantic and Pacific Bell. The XUNET-II trial will examine the ability of SONET-compatible equipment to carry digital applications such as medical imaging and multimedia digital libraries at OC-3 level.

SDH is currently being tested in Spain. Two commercial trials are scheduled for the United Kingdom in the third quarter of 1991. Switzerland has ordered an SDH link for installation between Lausanne and Geneva to be ready for Telecom 91 in October, and Germany has placed orders for its U2000 digital cross-connect project for SDH experiments.

Alcatel and Deutsche Bundespost Telekom of Germany have plans to trial the QPSX MAN by the end of 1990. The MAN will consist of several network nodes and one main network management system. Initially, it will operate at the rate of 34 Mbit/s, and eventually it will be upgraded to an international transmission rate of 140 Mbit/s. This technology will be migrated to support the ATM service for B-ISDN.

Siemens Communication Systems Inc. has announced that it will provide technical support in a Bell Atlantic MAN trial using switched multimegabit data service (SMDS) at Temple University. This trial will use fiber optic cables in the Bell of Pennsylvania network and MAN switching equipment provided by QPSX systems Inc. A tariffed MAN offering from Bell Atlantic could appear by late 1991.

3.7 Broadband Services Offering Plan

Broadband services are led by two directions. One is the demand of the market and the other is the investment in fiber and satellite networks for new demands.

The B-ISDN standard will be finalized at the CCITT 1992 meeting. Southern Bell supported by Northern Telecom has tested the residential broadband service at Heathrow, Orlando, Florida. The service includes 4 switched 107 Mbit/s video channels, one basic rate (2B+D) 144 kbit/s ISDN channels, and 2 64-kbit/s Plain Old Telephone System (POTS) channels. This service was provided in July 1989.

The research and development in advanced communication for Europe (RACE) project will ratify an architecture for integrated broadband communication (IBC) in 1991. The IBC will be commercially available in 1995. Specifically, the Atmospheric project (which means ATM and hybrid) will develop a system concept for broadband switching and multiplexing using the ATM technique. The Atmospheric project is expected to last five years.

ELLEMTEL (Televerket's and Ericsson's joint-owned development company) is conducting research and experiments for B-ISDN/IBC, where IBC stands for integrated broadband communications. The information transfer mode will be based on ATM. This scenario is to offer SMDS first through MAN switching systems. Next an ATM switch will be introduced, which will be operated at 155.52 Mbit/s. Finally, the B-ISDN/IBC will be developed to support both ATM and SDH. Experiments have been conducted to study and implement an optical customer access and ATM switch. The switching element is based on a ring topology. The interconnection network topology uses a Clos network.

NTT is planning to have public B-ISDN service in operation by 1995.

3.8 Broadband Network Deployment Plan

AT&T Network Systems has announced a high-speed, fast packet switch in October 1990 to support broadband services. The broadband switching system, called the BNS-2000, includes customer premises-based bridges and routers as well as integrated network management. The trunk speed for the switch can be up to 155 Mbit/s. This switch has been used by Nynex, Bellsouth, Southwestern Bell, and Pacific Telesis for internal switched multimegabit data service (SMDS) trials. There are two associated fiber access products with this BNS-2000 switch: the BRT-2000 access node for business services and the SLC-2000 digital loop carrier system for mixed and residential developments. These two fiber access products and the BNS-2000 switch are all part of the Service Net-2000 family.

Northern Telecom has announced a broadband digital architecture called FiberWorld. This all-fiber network comprises SONET-based access, transport, and switching products, which has self-healing and remote provisioning capability. Self-healing uses

digital cross-connect machines for traffic rerouting in the SONET transport network to achieve facility restoration after loss of the physical transmission facility between two sites. Several carriers: MCI, Bellsouth, Bell Atlantic, Contel and United Telephone have started testing the FiberWorld products.

DBP Telekom has announced the installation of Network 2000 next year, with commercial service due in mid-1992. The digital cross-connects, which can switch channels from 2 to 155 Mbit/s, are the first stage of implementing the SDH network across Germany.

Section 4

Satellite Network Architecture Alternatives

The satellite has inherent capabilities of providing point-to-multipoint and broadcast connection, interconnection between any two points within beam coverage areas, implementation of transmission facilities with short notice, instantaneous capacity reallocation of the entire network, and distance-insensitive cost. The use of multiple spot beams and on-board rate conversion/switching will provide additional flexibility, a larger capacity, and lower user terminal cost. Although the major portion of B-ISDN traffic will be carried by terrestrial networks, communications satellites will play a vital role for complementing and sometimes competing with terrestrial-based services. Satellites are also indispensable for future NASA missions, such as data collection, data relay, data distribution, and space station communications. This section includes an identification of potential satellite applications of B-ISDN services and typical bit rate requirements.

There are several architectural options for implementing a satellite B-ISDN system, ranging from a conventional point-to-point trunking service at a bit rate of 155 Mbit/s or higher to a flexible on-board fast packet switching service with dynamic allocation of system capacity. The major factors influencing architecture selection include a bit rate, system capacity, connectivity, traffic characteristics, and a user profile. Satellite network architectures can be classified into a circuit switched system and a fast packet switched system according to their payload switching characteristics. This section presents a comparison of the two types of systems and a high level description of several alternate satellite network architectures. Sample architectures for specific network applications are also included.

4.1 Potential Satellite B-ISDN Services

The potential users of satellite B-ISDN may virtually include every sector of industry, such as manufacturing, transportation, utilities, banking/financial, insurance, broadcasting, oil exploration, government agencies, universities, international organizations, and military. In fact, any current and future telecommunications users may be regarded as the potential users of satellite B-ISDN. Usage may range from a secondary role, such as cable restoration or backup for terrestrial links, to major backbone networks for private/business users and TV program distribution. A variety of B-ISDN services and their typical bit rates are described in Section 2. This section

addresses potential satellite networks for B-ISDN services. Table 4-1 shows broadband satellite network services, bit rate ranges, burstiness of traffic, and the degrees of user interaction.

Table 4-1. Satellite B-ISDN Networks

SATELLITE NETWORK SERVICE	BIT RATE RANGE (Mbit/s)	BURSTINESS	USER INTERACTION
Full Motion Video Service	8 - 140	Low	Low to Medium
Science Data and Communications Network	0.064 - 300	Low to High	Low to High
Supercomputer Networking	45 - 1600	Medium to High	Low
Private/Business Networks	.064 - 100	High	High
Trunking	51 - 1200	Low	Low
Emergency Communications	0.064 - 45	Low to Medium	Low to Medium
Thin-Route Networks	0.064 - 10	Low to High	Low to High

As seen in the table, a wide range of bit rate is expected in B-ISDN services. The B-ISDN operates at a bit rate of 51 Mbit/s (SONET STS-1) or higher. However, an actual information rate can be lower, and a satellite network may be designed independent of actual B-ISDN user interface rates. Burstiness of traffic may be characterized by the peak-to-average ratio of traffic flow; it is low for a continuous flow of traffic, such as video program distribution, and high for interactive data communications. The burstiness of trunking service is rated low due to its nature of leased circuit operation, but an actual traffic flow can be quite bursty. The degree of user interaction indicates a dynamic network behavior in terms of user connectivity and capacity allocation. Low interaction implies semi-static network operation with traffic reconfiguration at a slow rate, and high interaction requires dynamic network reconfiguration and flexible capacity allocation to meet the user's traffic needs. Each network service listed in the above table is further exploited in the following.

A satellite network can provide B-ISDN services through one of three interface configurations: (a) direct (mesh) interconnection between users through user-network interfaces (UNI-UNI), (b) interconnection between a user and a switching center through a user-network interface and a network-node interface (UNI-NNI), and (c) interconnection between switching centers (NNI-NNI). Figure 4-1 depicts these options.

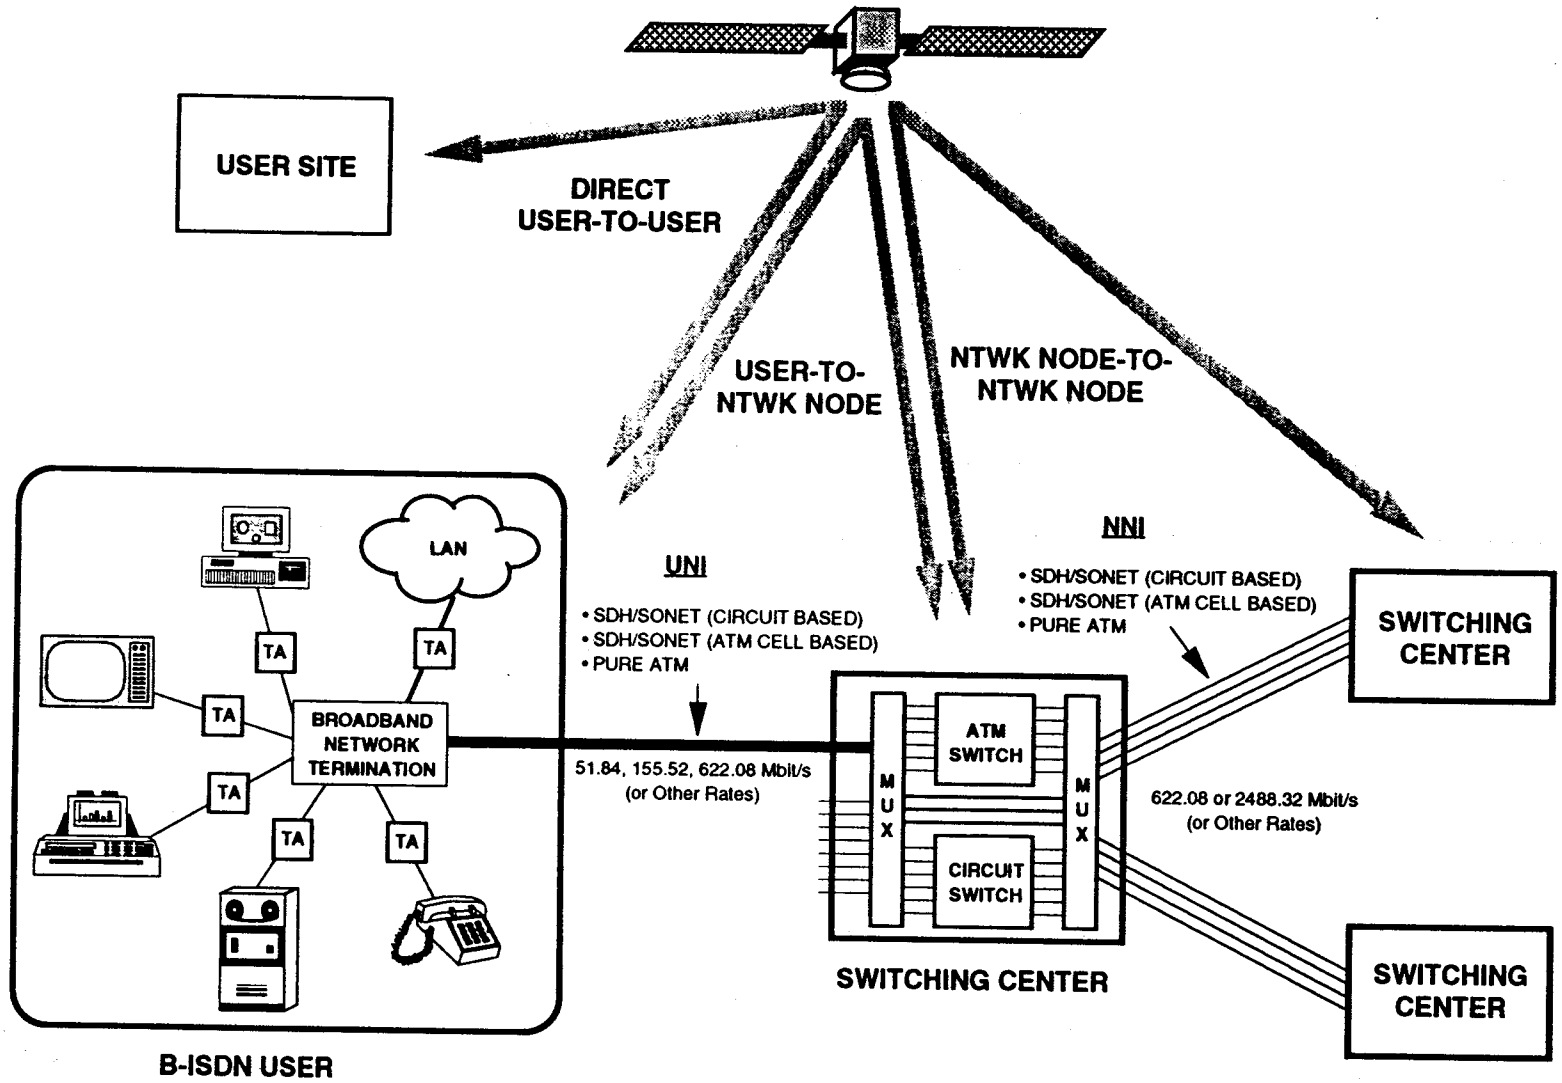

Figure 4-1. A satellite B-ISDN can provide direct interconnection between users, switching centers, and a user and a switching center

4.1.1 Full Motion Video Service

The broadcast capability of a satellite is well suited for providing cost-effective distribution of video programs to local broadcast stations, cable TV operators, and movie theaters. Currently, the satellite is the major transmission medium for video program distribution, live video broadcast of national news events, emergency news, sports events, and religious programs. As the trend for globalization of TV program distribution and news gathering activities intensifies, a full motion video service includes a large number of overseas locations. Present video transmission is primarily analog, but will eventually utilize digital technology for video processing and transmission. The transmission bit rate varies according to the coding technique used and the picture quality desired. Table 4-2 shows bit rate requirements for different classes of picture quality.

Table 4-2. Bit Rates for Full-Motion Picture Coding [4-1]

SERVICE QUALITY OF MOVING PICTURES	UNCOMPRESSED BIT RATE (Mbit/s)	COMPRESSED BIT RATE (Mbit/s)	QUALITY CLASS
Picture Quality A	~ 1,000	92 - 200+	High Definition TV Quality
Picture Quality B	216	30/45 - 140	Studio-TV Quality
Picture Quality C		8 - 34	Present TV Quality (PAL, SECAM, NTSC)
Picture Quality D		0.384 - 1.92	Reduced Resolution Video Quality (Spatial and Temporal)
Picture Quality E		0.064	Highly Reduced Resolution Video Quality (Spatial and Temporal)

High definition television (HDTV) requires a baseband signal of between 20 and 30 MHz and an uncompressed bit rate of between 600 Mbit/s and 1 Gbit/s. With digital video compression, the rate may be reduced to around 140 Mbit/s. Studio quality and current broadcast quality (e.g., PAL and NTSC) pictures may require transmission bit rates of 60 - 70 Mbit/s and 30 - 45 Mbit/s, respectively. Intensive on-going development efforts of video coding techniques will likely result in lower bit rate requirements for the transmission of high quality video. Video coding techniques and exact bit rates are currently being studied by national and international standard organizations as well as by manufacturers.

Video program distribution is typically point-to-multipoint such that the same program is distributed to a large number of local broadcast stations or cable TV operators. Some live national TV programs involve multiple studios and field sites to interactively broadcast the program from different locations. Efficient space segment utilization may be achieved by dynamically reallocating the same uplink capacity to different feederlink

stations. Traffic flow will be virtually constant for the duration of a video connection. Variable bit rate video is currently being investigated by a number of research organizations and may be more efficiently supported by a fast packet switch based transmission system.

4.1.2 Science Data and Communications Network

A broadband satellite network may find application for supporting future NASA space missions, such as realtime distribution of space experiment data (experiment monitoring, control, and observation data) from the space station, space shuttle, and low orbit science spacecraft to NASA centers, experimenter locations, and science data users [4-2]. The same satellite may also provide high-speed telecommunications services among various NASA centers and NASA network users for video teleconference, document/file transfer, exchange of science information, database access, data archives, and supercomputer access/resource sharing. The satellite will provide communications links to a variety of low orbit spacecraft/shuttle/space stations through intersatellite links, and to NASA centers and user locations through both high capacity spot beams and wider fixed beams covering the CONUS. A sample satellite network design for NASA science and communications applications is presented in Section 4.5.

A wide range of transmission bit rates and traffic characteristics are expected in the NASA network. For example, full motion video transmission from a space station and image transmission from a space telescope may require a bit rate of 45 Mbit/s to 300 Mbit/s with relatively slow return links (i.e., asymmetric circuit switched traffic). Some science experiment data, on the other hand, are at lower speeds (64 kbit/s - 1.544 Mbit/s), but require realtime distribution to a large number of users (low bit rate multicast connection). Communications among NASA centers may include interactive voice/data transmission at a bit rate of 64 kbit/s (voice and low speed data), video teleconferencing at 384 kbit/s, full motion video transmission at 45 Mbit/s, and supercomputer access and high speed image retrieval at over 100 Mbit/s. Thus, the NASA science and communications network must support both wideband circuit switched connection and bursty interactive communications services.

The NASA network uses the standard protocols specified by the Consultative Committee for Space Data Systems (CCSDS) for space/ground data communications for space science application [4-3]. Proper interface must be provided at the ground station between the space science community's CCSDS protocols and the B-ISDN protocols. Telescience implies the capability to feed commands to spacecraft for control of on-board functions and to receive appropriate responses, and it requires protocol crossing between CCSDS and B-ISDN in near realtime.

4.1.3 Supercomputer Networking

Supercomputers have been extensively used to perform research in such areas as aerodynamics and fluid dynamics, weather modeling and forecasting, molecular modeling for human DNA, nuclear weapons, and fusion technology. The major

supercomputer facilities in the U.S. include the NASA Ames Research Center, the Los Alamos National Laboratory, the National Center for Atmospheric Research, and Lawrence Livermore National Laboratories. Some of these facilities dedicate several supercomputers on a single task to perform parallel processing, and the number of workstation users ranges from a few hundred to a few thousand. A computing facility is often available to researchers in remote locations, such as universities and various national research organizations, and even to foreign research institutes.

A telecommunications network for supercomputers is used for sharing of geographically distributed computational resources and local visualization of results produced by remote application programs. The transmission requirement varies from a low rate (64 kbit/s to 1.5 Mbit/s) for the display and analysis of results at the user sites to an ultra high rate (over 1 Gbit/s) for a visual display in 3D graphics. Scientific visualization using high performance workstations (with 2D or 3D graphics) is critical (and essential in some applications) for getting most benefits from supercomputing. Future supercomputer networks will likely be required to meet the ANSI's interface standards for the High Performance Parallel Interface (HPPI), which operates at 800 Mbit/s or 1.6 Gbit/s.

A network configuration is typically a star — a large number of workstations connected to a central processing facility. User interaction is minimal, and the major data flow is from a supercomputer site to a user location. The traffic flow can be bursty (e.g. transmission of an image) or almost continuous (e.g. display of 3D simulation). A satellite system must support dynamic reconfiguration of a network to allow user access from various remote locations.

4.1.4 Private/Business Networks

Private and business networks may be regarded as an extension of current VSAT networks, providing a high speed communications service to a large number of users at a bit rate of 10 - 100 Mbit/s. A typical user terminal will employ an antenna size of between 1.2 m and 2.4 m installed on the customer's premise and will provide low-cost wideband service to closed user groups. The use of high power spot beams along with on-board baseband switching allows the users to directly communicate with other users without going through a hub. A variety of user interfaces, not necessarily limited to B-ISDN, include a B-ISDN interface at 51 Mbit/s (STS-1) or 155 Mbit/s (STS-3 or ATM), basic rate or primary rate ISDN, frame relay, X.25/X.75, local area networks (LANs) operating at 10 - 100 Mbit/s, and metropolitan area networks (MANs) (155 Mbit/s). Although the user interface rates can be high, the majority of users have relatively low information rates, ranging from 64 kbit/s to several megabit/s. Some user sites, such as corporate data processing centers, may require somewhat higher bit rates.

Direct user-to-user communications include the interconnection of high-speed LANs, voice, data, image, and video communications. Video telephony may be used, for example, for sales, consulting, instruction, negotiation, and a discussion of visual information, such as reports, charts, and advertising layouts. A video teleconference, with such features as facsimile, document transfer, and an electronic blackboard, would

allow the participants to discuss business matters without leaving their workplace. Other applications include computer aided simulation, evaluation, and optimization procedures, program downloading, file transfer, color facsimile, the transmission of newspaper pages, tele-education for employee training, and telemedicine. The use of data and image compression techniques can reduce the required transmission rate by a factor of 3 to 15 and would allow most of these applications to be supported by a transmission rate of no more than 2 Mbit/s.

The majority of users in private/business networks will require relatively low bit rates (in the range of 64 kbit/s to 2 Mbit/s) and will have highly bursty traffic with constant interactions with other network users. The user traffic will be primarily packet based, such as frame relay, ATM, X.25/X.75, datagram, LANs (e.g. Ethernet, token ring, and FDDI), and MANs (e.g. SMDS), although some require circuit switched connections (e.g. voice and video teleconference). In this environment, the satellite network must be flexible to dynamically accommodate changes in user traffic demand. On-board fast packet switching not only provides the desired flexibility but also can accommodate circuit switched connection through a common switch fabric.

4.1.5 Trunking

A trunking network provides a high-speed interconnection between large network user sites, such as international switching centers through Gateway stations, corporate headquarters and their major affiliated offices or data processing centers, and telephone switching centers of the main land and islands. The primary application of satellite will be limited to long-haul transmission in which satellite links provide lower cost than terrestrial fiber optic networks, areas or regions where alternate wideband transmission facilities do not exist, back-up transmission links to backbone networks, and cable restoration between major switching centers.

A trunk satellite link can provide a point-to-point connection at a B-ISDN bit rate of 155.52 Mbit/s (STS-3), 622.08 Mbit/s (STS-12), or 1,244.16 Mbit/s (STS-24). The network should also be capable of supporting fractional as well as multiplexed B-ISDN rates at n x 51.84 Mbit/s and n x 155.52 Mbit/s. Transmission links may be set up on a semi-permanent basis or may be dynamically reconfigurable as connectivity and/or traffic volume changes. Network reconfiguration will be less frequent compared with a call-by-call based demand assigned network. The user may require concentration of traffic channels to efficiently utilize the allocated or leased capacity.

4.1.6 Emergency Communications

There have been numerous local and regional disasters every year in the world, often resulting in a fatal toll and also directly affecting tens of thousands of people. Large earthquakes, hurricanes, and floods often disable communications links, isolating the region from the rest of the world. The first step in disaster recovery operation is to set up a communications link to report damage assessment and to request medical aid, food, clothes, supplies, and proper machinery and tools to rescue people. Satellite

communications is vital for emergency rescue operation. Using a fly-away terminal, a transportable earth station or a mobile earth station, a communications link can be set up within hours after the occurrence of the event.

Depending on the area and type of disaster and the accessibility to the area, communications requirements range from a single voice channel to full motion video along with several voice channels and data terminals. Because of the critical nature of communications, dedicated satellite channels will be initially allocated for the rescue operation. Along with emergency communications, a temporary switching center to provide a small number of voice channels to unaffected regional switching centers may be installed until terrestrial communications facilities are partially restored. During disaster recovery operation, satellite channels can be dynamically allocated according to the traffic volume and connectivity required.

4.1.7 Thin-Route Networks

A thin-route network provides low-to-medium speed communications links between remote areas and major cities, in which other communications links do not exist or cannot cost-effectively provide desired link performance. Examples include islands, rural areas, branch offices of corporations or government agencies in developing countries, and isolated remote areas. There is some commonality between a thin-route network and private/business or trunking networks, but a thin-route network is characterized by low speed requirements and potentially shared space segment operation by many customers.

A thin-route network may carry a mix of traffic types: voice only users, voice and data users, low-speed data users, and users with voice, data, and video teleconferencing. A small earth station equipped with an unmanned switching system may provide telephone channels to a remote area. A satellite network may also be used to extend a terrestrial-based corporate network to a major branch office where a good transmission link is not available.

4.2 Comparison of Circuit and Fast Packet Switched Systems

A satellite B-ISDN system may be implemented using high-power transparent transponders, satellite-switched TDMA or FDMA, or an on-board baseband processor (OBP). The use of on-board baseband switching/processing technology provides better link margins, flexible interconnection of users in different spot beams, and optimized design of uplink and downlink transmission systems based on user traffic requirements. This section addresses two alternate switching architectures for on-board processing satellite systems.

The switching architectures considered herein are circuit switching and destination-directed fast packet switching, where the former has been used in the Advanced Communications Technology Satellite (ACTS) [4-4] and European Space Agency (ESA) time-space-time (TST) prototype development [4-5], and the latter has been proposed for

future advanced on-board processing satellites [4-6]. These remarkably different switching concepts are analyzed to assess flexibility for supporting circuit and packet switched traffic, impact of traffic reconfiguration, and switch implementation complexity. Table 4-3 summarizes the results of the analysis, and detailed discussions are given in the following.

Table 4-3. Comparison of Circuit and Fast Packet Switching Architectures

SWITCHING ARCHITECTURE	CIRCUIT-SWITCHED TRAFFIC	PACKET-SWITCHED TRAFFIC	TRAFFIC RECONFIGURATION
CIRCUIT SWITCHING	• Efficient Bandwidth Utilization	• Very Inefficient Bandwidth Utilization • Inflexible Connectivity	• Reprogramming of On-Board Switch Control Memories • Reconfiguration of Earth Station Time/Frequency Plans for Each Circuit Setup • Difficult to Implement Autonomous Private Networks
FAST PACKET SWITCHING	• Can Accommodate Circuit-Switched Traffic • Somewhat Higher Overhead Due to Packet Headers	• On-Board Congestion May Occur	• Self-Routing • Does not Require Control Memory for Routing • Reconfiguration of Earth Station Time/Frequency Plans for Major Traffic Changes • Easy to Implement Autonomous Private Networks

4.2.1 Circuit Switching

Circuit switching employs a connection-oriented switching mechanism to set up a path from an uplink time slot to a designated downlink slot. A circuit switching concept is illustrated in Figure 4-2. An on-board path setup must be controlled from a central location, such as a ground network control center (NCC) or an on-board network controller (OBNC), to avoid conflicts. The path, once set up, will remain indefinitely or until it is reconfigured. Multicast connection is provided by mapping the same uplink time slot to multiple time slots on different downlink carriers.

The user earth station will be allocated a fixed traffic capacity for transmission to a specific destination earth station(s). The allocated time slots are exclusively used by the user to communicate with the corresponding destination user. Upon termination of a call, the allocated time slots may be reused, without reconfiguring an on-board switch connection, to transmit traffic to a different destination earth station which is operating on the same downlink carrier. Of course, necessary signaling messages must be exchanged for a new call setup.

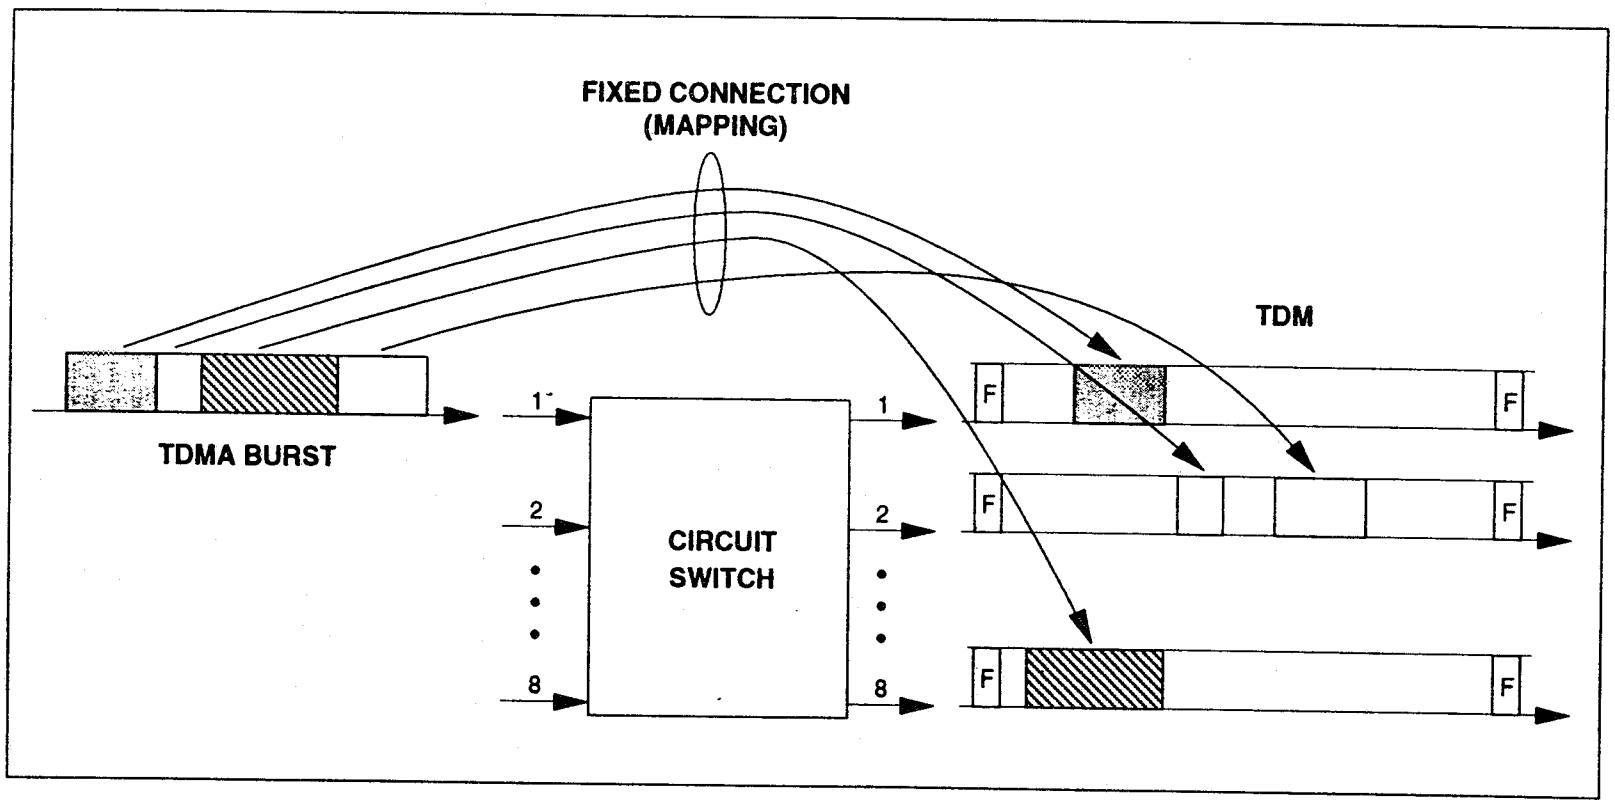

Figure 4-2. Circuit Switching Concept

Circuit switching is most efficient for circuit switched traffic in which assigned time slots are fully occupied by a constant flow of traffic. A typical TDMA system can provide a frame efficiency of 90 to 95 percent. Circuit switching can also support transmission of packet switched traffic, but its effectiveness diminishes as the connectivity of a network increases. For example, to provide simultaneous virtual circuit connections from one terminal to n beams, the network must provide n separate channels of dedicated circuit connections. Since packet traffic is bursty in nature, these channels will be extremely underutilized most of the time. Thus, practical accommodation of packet switched traffic may be limited to connections to a small number of destination beams. However, if the volume of packet switched traffic is relatively small, double-hop connection through a data hub (or a control center) may be used as currently pursued for a European system [4-7].

In circuit switching, an on-board routing path must be preconfigured prior to transmission of traffic channels. OBP reconfiguration is necessary for changing the mapping between uplink and downlink time slots, burst positions and burst lengths for TDMA uplink, and carrier bit rates and frequency slots for FDMA uplink. Each call setup will likely require reconfiguration of the OBP. Reconfiguration may be performed by the NCC, such as in the ACTS system, or by the OBNC. The former requires constant interaction with the OBNC through a satellite link. In either case, the OBNC must be capable of processing all call setup and teardown requests generated by the network. For a large network, this will impose a severe processing load on the OBNC processor.

A satellite bandwidth is often divided into smaller frequency bands, which are leased to various user groups. It is up to each user group how to use the leased capacity. Since a circuit switched OBP must be centrally controlled to avoid conflicts of assigning the same uplink and downlink time (or frequency) slots to different user groups, the NCC or OBNC must perform additional processing functions. Thus, a control subsystem must manage not only individual call connections but also subnetwork resource allocation.

There are numerous switch configurations for circuit switching. For a low to medium capacity system (up to 1.5 - 2 Gbit/s), candidate switch structures are a common memory switch, a parallel bus with distributed output or input memories, and a high speed fiber optic ring. For high-capacity application with a total capacity of multi-gigabit/s, a time-space-time (TST) or a fiber optic ring with wavelength division multiplexing may be used. An on-board circuit switch requires control memories to store routing path information. The amount of control memory, as seen in the ACTS system, may not be insignificant, and its implementation requires a special consideration to minimize a bit flip problem.

4.2.2 Fast Packet Switching

Fast packet switching, also known as destination directed packet switching, is a connectionless switching mechanism, in which data packets are routed from uplink to down-link beams according to the routing information contained in the packet header.

The concept is similar to fast packet switching for Asynchronous Transfer Mode (ATM) in broadband ISDN. Its main characteristic is self-routing without NCC or OBNC control. The on-board fast packet switching concept is illustrated in Figure 4-3. Since no fixed mapping exists between uplink slots and downlink slots, the packets generated from the same data source and transmitted in the same uplink slots may be routed to different positions in the downlink frame.

The user terminal is allocated a number of uplink (time or frequency) slots for transmission of data packets to different destination beams (or carriers) and is not restricted in use to each time slot for a particular connection. The number of slots allocated to a user terminal may be dynamically adjusted according to the user's traffic volume. One potential problem with fast packet switching is an on-board buffer overflow caused by a large amount of traffic flow to a particular downlink beam which exceeds the beam capacity. To minimize packet loss at the satellite, some form of flow control must be provided.

The congestion problem for fast packet switching can be completely eliminated for circuit-switched traffic by allocating a desired capacity (on a call-by-call basis) for both uplink and downlink carriers. For packet-switched traffic, flow/congestion control techniques include the following: (a) a dynamic allocation of fixed capacity from a transmit earth station to each downlink carrier, (b) call admission control at the earth station, (c) on-board capacity allocation based on current queue status, (d) feedback control, and (e) a combination of these. The goal of flow/congestion control is to achieve a certain packet loss ratio within a satellite system such that no significant degradation results in the end-to-end performance.

Fast packet switching is particularly suited for packet switched traffic. A single uplink time slot can carry data packets to all destination beams, provided that the total packet traffic volume from a user terminal does not exceed the slot capacity. Circuit switched traffic can also be supported by packetizing a continuous circuit switched bit stream and appending a packet header identifying a destination downlink carrier, a destination earth station, and a circuit connection. Alternately, these functions can be performed on board prior to packet routing in coordination with the user terminals and the OBNC.

Fast packet switching has a lower transmission frame efficiency, in the range of 80 to 85 percent, due to packet overheads. However, an actual throughput may be higher due to an inherent flexibility built into the architecture. Unlike the case of circuit switched OBP, assigned uplink slots can be shared, without redesignation, for different types of services (voice, virtual circuits, frame relay, datagram, and ATM). The use of voice activation will provide additional savings in transmission capacity. A 15-percent increase in channel utilization yields a comparable throughput efficiency to circuit switching. In practice, a throughput increase can be as high as 100 to 150 percent.

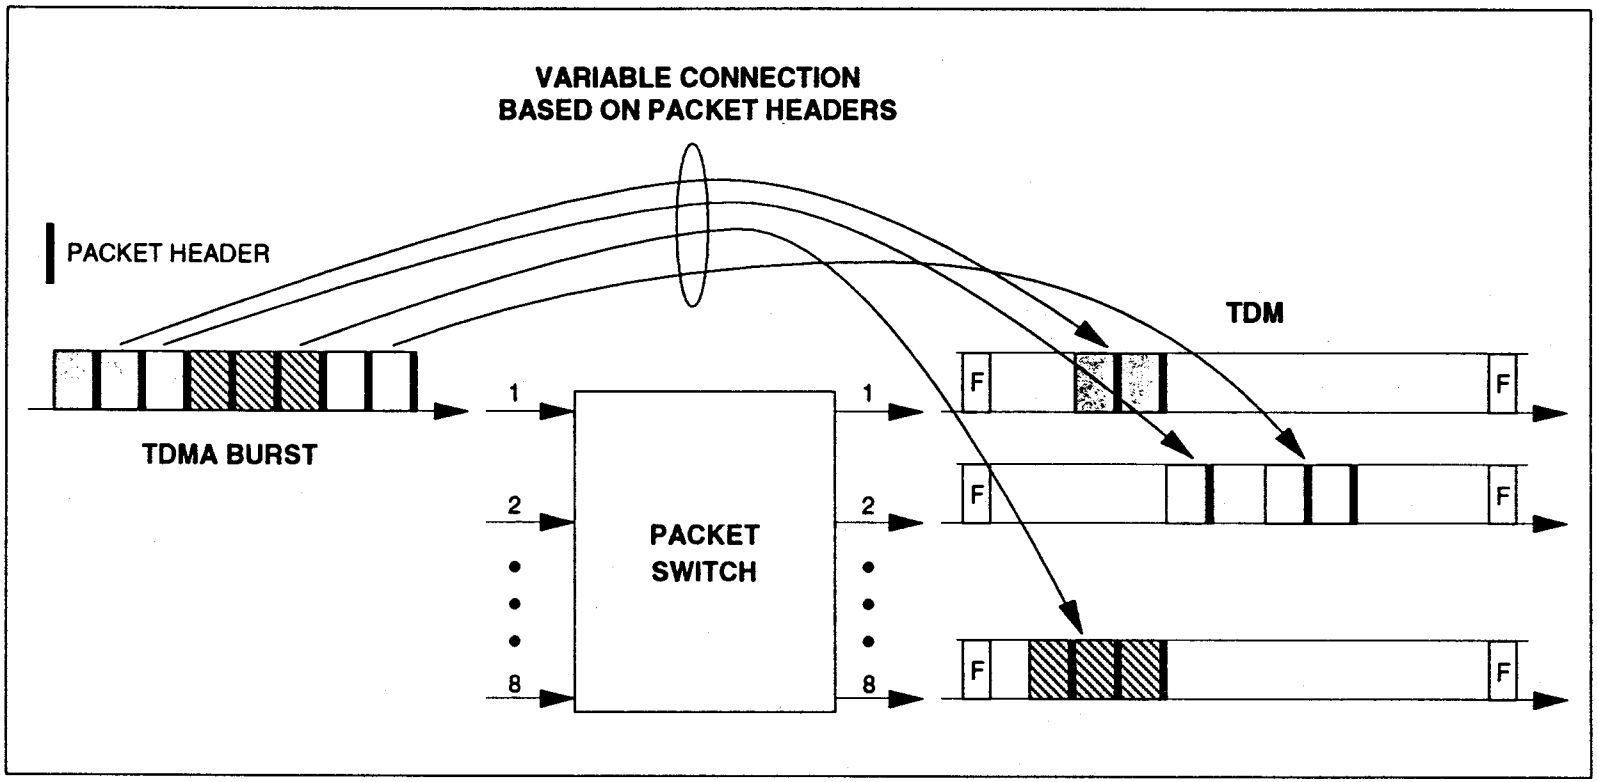

Figure 4-3. Fast Packet Switching Concept

Fast packet switching allows the use of allocated uplink slots for transmission to any downlink beam without reconfiguring on-board switch connection. OBP reconfiguration will be necessary for changing burst positions for TDMA uplink, and carrier bit rates and frequency slots for FDMA uplink. Reconfiguration will be in the form of reprogramming of on-board demodulators and timing generators but not the routing paths through the fast packet switch.

Distributed network control is possible by individual user groups. Each user group has its own NCC which manages assigned uplink and downlink resources such that the total amount of traffic within the user group does not exceed the allocated capacity. The NCC may perform demand assignment processing to allocate uplink slots to member earth stations and control the amount of traffic flow to each downlink beam. Alternately, the NCC functions can be implemented by the OBNC.

There are several alternatives for fast packet switch structures. A conventional circuit switch, such as a common memory switch, a high-speed parallel bus, or a fiber optic ring, can be used for self-routing. There will be no internal contention problem with this type of switch. The total switch throughput is limited to about 2 Gbit/s due to a shared facility (a memory, a bus, or a ring). Multistage space switches, such as banyan-based switches and self-routing crossbar switches, can provide higher throughputs. The contention problem with a multistage space switch can be avoided using, for example, an output port reservation scheme or may be minimized by increasing a switch speed or by providing multiple routing paths within a switch fabric.

A self-routing switch does not require a control memory to set up a routing path, and hence there will be no misrouting of data packets by a switch fabric. However, errors in a packet header may cause the packet to be routed to an incorrect output port. Therefore, headers must be protected from transmission errors. A simple error-correction-and-detection code usually provides sufficient error protection such that the performance degradation due to header errors is insignificant. If a transmission link is FEC encoded using a convolutional code, bit interleaving of header bits may be necessary along with header error correction coding.

4.2.3 Summary of Comparison

A general conclusion on the selection of circuit or packet switching can be drawn based on the discussions in the above sections. Circuit switching is advantageous under the following conditions:

a. The major portion of the network traffic is circuit switched, or most of the network users require circuit switched connection, and

b. The network does not require frequent traffic reconfiguration.

Fast packet switching may be an attractive option for a satellite network carrying only circuit switched traffic if frequent traffic reconfiguration is required. The use of a self-routing switch will eliminate the necessity of path finder processing. In some

situations, a mixed switch configuration, consisting of a circuit and packet switches, may provide an optimal on-board processor architecture. An example of such a system is a network which provides a high-speed trunking service as well as a packet-based data service for a large number of users. On-board interconnection is provided by cross-strapping the two types of payloads.

4.3 Circuit Switched Systems

Circuit switched systems considered in this section are exclusively for B-ISDN services and assume a minimum bit rate of 51 Mbit/s. The maximum (information) bit rate supported by the space segment depends on several factors, such as frequency band, link availability, satellite G/T and EIRP, and earth station equipment (in particular, an HPA and antenna size). According to the link analysis given in Appendix A, the Ka-band space segment with narrow spot beams can support an information rate of over 1 Gbit/s. The Ku-band link, on the other hand, can provide satellite interconnection at up to 622 Mbit/s using a 500 MHz bandwidth.

High-speed transmission necessitates the use of narrow spot beams to reduce earth station RF requirements while providing a large system capacity. Spot beam operation generally accompanies an interconnection problem among beam coverage areas. If connectivity requirements are limited, for example, for East Coast to West Coast or East Coast to Europe, a simple fixed beam satellite (with some steerable beams if flexibility is desired) will be most cost effective. To improve link performance, a system may employ on-board regeneration, FEC coding, and/or uplink power control.

Flexible beam interconnection can be implemented by transponder channelization and on-board interconnection of beam frequency bands or by some form of switching at the satellite. The first approach generally results in a simpler payload design, but reduces the maximum bit rate supported by a single carrier. Also, payload complexity significantly increases for a larger number of beams. The second approach will provide more flexibility in terms of beam interconnection and user bit rates. On-board interconnection may be implemented by satellite-switched TDMA (SS-TDMA) using either an RF or a baseband switch matrix, satellite-switched FDMA (SS-FDMA), or on-board baseband cross-connect. These system architectures are described in detail in the following subsections.

4.3.1 SS-TDMA System

The SS-TDMA system provides interconnection of uplink and downlink beams by the use of a microwave (MSM) or a baseband switch matrix (BSM). Examples of SS-TDMA systems are INTELSAT VI (120 Mbit/s, RF), ITALSAT (147 Mbit/s, baseband), and ACTS HBR (TBD Mbit/s, RF). Figure 4-4 illustrates a system concept. A TDMA traffic burst consists of a preamble (carrier and bit timing recovery pattern and a unique word), a signaling channel, and a number of subbursts. A signaling channel contains station status, subburst information, voice and/or TTY orderwire channels, and other messages necessary for setting up and maintaining a satellite link. A subburst is a

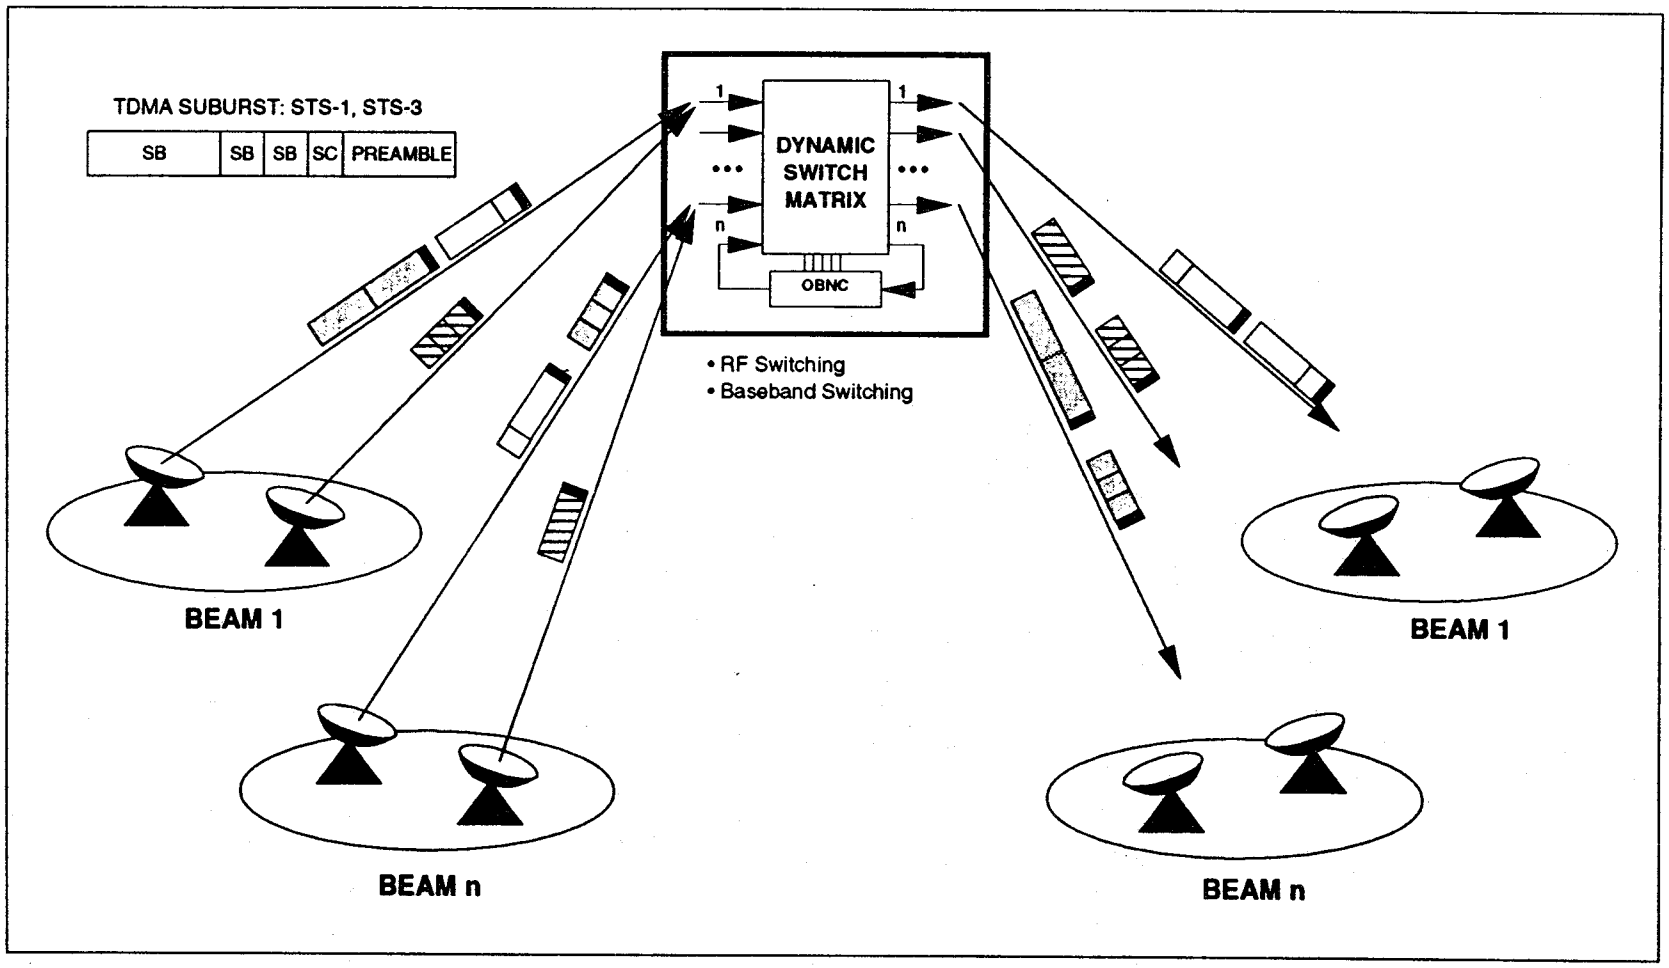

Figure 4-4. SS-TDMA Network

traffic channel destined to a specific destination earth station(s) and includes a SONET frame, such as STS-1 or STS-3, or an ATM stream.

A reference burst is transmitted from a ground reference station or an on-board network controller. It provides reference timing to traffic stations in the network. The reference burst may also contain burst time plan (BTP) messages in response to traffic stations' channel requests. Another type of burst, called a synchronization burst, may be needed for traffic stations to perform initial transmit acquisition and subsequent synchronization with a reference burst or on-board frame timing. The synchronization burst is often used for orderwire communications with a network control center to request allocation and deallocation of satellite capacity.

The SS-TDMA system provides full beam connectivity with possibly the simplest on-board hardware. A 622-Mbit/s SS-TDMA system can accommodate four active traffic stations per beam, each carrying 155 Mbit/s of traffic or 12 traffic stations per beam at 51 Mbit/s. The earth station traffic capacity varies from as low as 51 Mbit/s to a maximum burst rate without changing its basic equipment configuration. This desirable feature from a flexibility viewpoint can be a drawback to users with a small amount of traffic. In this situation, multiple bit rates may be used by interleaving different burst rates within the same frame.

The use of hopping beams allows efficient utilization of satellite power by dynamically steering the beams to areas where traffic exists. The hopping beam SS-TDMA thus provides additional flexibility. Increased beam connectivity, however, yields an inefficient capacity utilization due to a traffic scheduling problem for an on-board switch matrix and a hopping beam scan sequence. For example, a system consisting of n hopping beams with m active dwell areas per beam will require at least nm^2 dwell connection states to provide full connectivity. This value does not take into consideration the scheduling efficiency and will be about n times larger for optimal scheduling. This deficiency can be overcome with more complex on-board hardware, i.e., storing TDMA bursts on board. For a small number of active dwell areas, however, hopping beam SS-TDMA is very efficient.

One unique feature of an SS-TDMA system is its ability to reconfigure a network on short notice and without loss of traffic. Dynamic network reconfiguration is easily achieved by a coordinated time plan change between user earth stations and an on-board switch. Computation of a burst time plan (BTP) or a satellite switch state time plan (SSTP) can be complex if a network design imposes a certain scheduling constraint, such as multidestination digital speech interpolation (DSI) in the INTELSAT VI system. However, a broadband SS-TDMA system usually does not require sophisticated networking, and traffic scheduling can be made relatively simple. An SS-TDMA scheduling algorithm such as given in [4-8] can thus be implemented by the OBNC.

4.3.2 SS-FDMA System

In the SS-FDMA system, a beam bandwidth is divided into a number of frequency slots, where each slot is occupied by a single high-speed TDM carrier. The number of FDMA carriers in a spot beam varies according to the earth station traffic requirements within the beam. Connectivity among spot beams is provided on board the satellite using either a static microwave switch or a baseband switch. Figure 4-5 depicts a system concept. The advantage of this system over SS-TDMA is a lower user earth station cost. For example, an earth station with a small amount of traffic (e.g. one 51-Mbit/s carrier) does not require a large antenna and/or a costly HPA. Also, earth station baseband equipment will be simpler than that of TDMA. SS-FDMA, however, suffers from the following shortcomings: (a) multiple carrier transmission requires earth station HPA power backoff, (b) it cannot operate with hopping beams, (c) flexible carrier bit rate assignment is difficult and complicates a payload design, and (d) dynamic network reconfiguration without loss of data is difficult. Overall, the SS-FDMA system is attractive for a relatively simple network configuration.

Orderwire communications between user terminals and the OBNC may be provided by a low-speed random access channel for uplink and a low-speed broadcast channel for downlink. The random access channel is used to send request messages for capacity allocation/deallocation and station status to the OBNC and is shared by all the earth stations in the beam. The downlink broadcast channel contains frequency assignment messages and station control information to individual earth stations.

The SS-FDMA system cannot be used in conjunction with hopping beams and hence will impose constraints on beam connectivity. A large number of narrow fixed spot beams will be costly and wasteful in terms of spacecraft power. On the other hand, wide spot beams will result in a smaller system capacity and more expensive earth station equipment. Possibly, the best alternative is to use a small number of fixed or steerable narrow spot beams along with some wider beams. The narrow spot beams are steered to the areas of traffic concentration, and the wider spot beams are used to pick up traffic from the areas not covered by the narrow spot beams. An on-board static switch provides connectivity between the two types of beams. This type of beam arrangement has been used often in the INTELSAT system.

Narrow spot beams will yield a usable frequency bandwidth of up to 160 MHz in each beam in the Ku-band (or 320 MHz with the use of dual polarization) and up to 330 MHz in the Ka-band (over 1-GHz band). These bandwidths should be able to support four to eight 51-Mbit/s carriers or a single carrier at 220 to 470 Mbit/s using QPSK modulation. The Ka-band beam can be configured to provide any of the following combination of carriers: (a) eight 51-Mbit/s, (b) one 155-Mbit/s and six 51-Mbit/s, (c) two 155-Mbit/s and three 51-Mbit/s, or (d) three 155-Mbit/s carriers. The most flexible SS-FDMA system supports all these combinations, but will result in a complex payload design. Restricting the configuration options to one or two will significantly simplify the payload design.

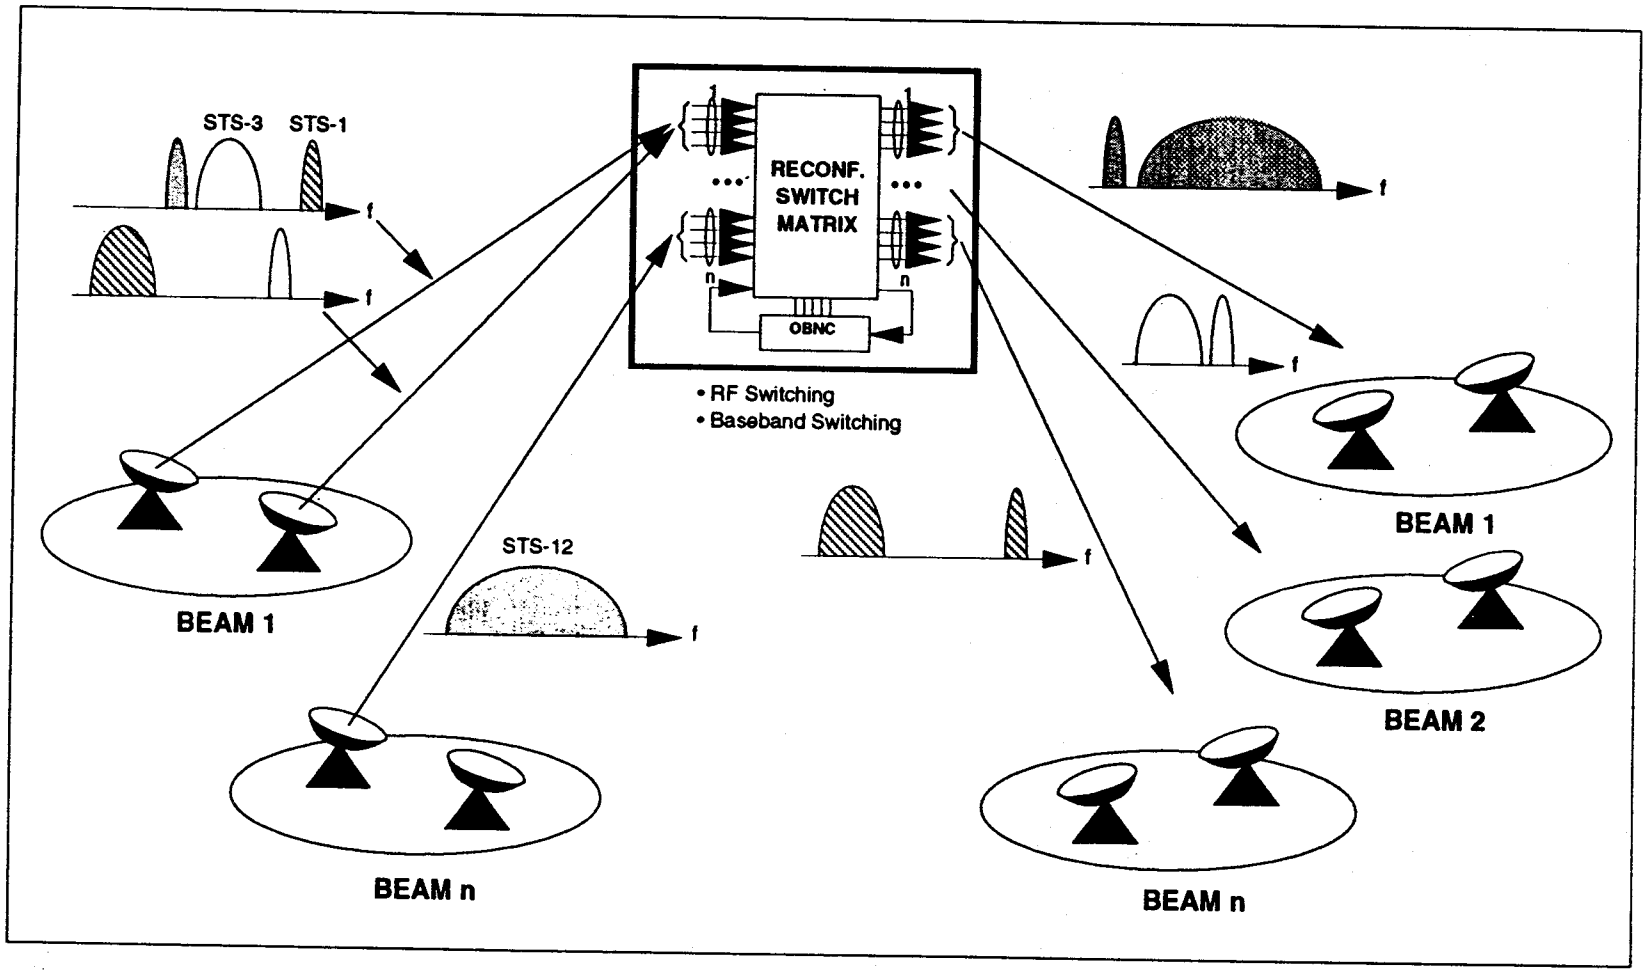

Figure 4-5. SS-FDMA Network

Frequency slot assignment and static switch reconfiguration are relatively simple, such that the OBNC should be able to implement the necessary traffic scheduling and assignment algorithms. To most efficiently utilize the space segment, it is desirable to have a capability of moving active carriers from their current positions to new ones within the allocated frequency band. This herding process can create a larger frequency slot for a higher bit rate carrier. Unfortunately, this feature is not easily implemented in the FDMA system without temporary loss of traffic. In this sense, traffic reconfiguration is static in the SS-FDMA system.

4.3.3 On-Board Cross-Connect System

In the previous two system concepts, earth stations perform proper multiplexing and demultiplexing functions to form multiple TDMA bursts or multiple FDMA carriers according to destination beams. In addition, these bursts or carriers are transparent and are delivered to the destination stations as they are transmitted. The SS-TDMA system requires all earth stations to operate at the burst rate and imposes scheduling constraints for hopping beam operation. The SS-FDMA system, on the other hand, requires multiple carrier operation to communicate with different beams.

The on-board cross-connect system not only provides additional flexibility but also optimizes earth station design to meet individual users' traffic needs. The cross-connect system concepts for TDMA and FDMA uplink are illustrated in Figure 4-6. A TDMA user station typically transmits one traffic burst in each frame, and the traffic burst contains B-ISDN channels to different destination stations. The TDMA burst is demodulated, demultiplexed into individual B-ISDN channels, and routed through a cross-connect switch to the proper output ports, where a continuous TDM signal is formed for downlink transmission. Similarly, the user earth station operating in FDMA transmits a single multiplexed carrier consisting of B-ISDN channels to different destinations. Each carrier is demodulated, demultiplexed, and routed to the output ports, where a single high-speed TDM signal is formed for each beam. In either case, a user earth station must deal with only a single burst or a single carrier, resulting in a more flexible circuit switched system and a lower cost earth station. Interconnection between the two types of systems is also possible by cross-strapping the switching payloads.

The TDMA traffic scheduling problem associated with beam hopping and on-board switching virtually disappears with the cross-connect system. In the cross-connect system, traffic scheduling for uplink and downlink beams can be independently performed, and any scheduling conflicts are resolved by an on-board buffer. (The on-board buffer basically acts as a time slot interchanger.) However, the use of hopping beams still requires burst mode downlink transmission.

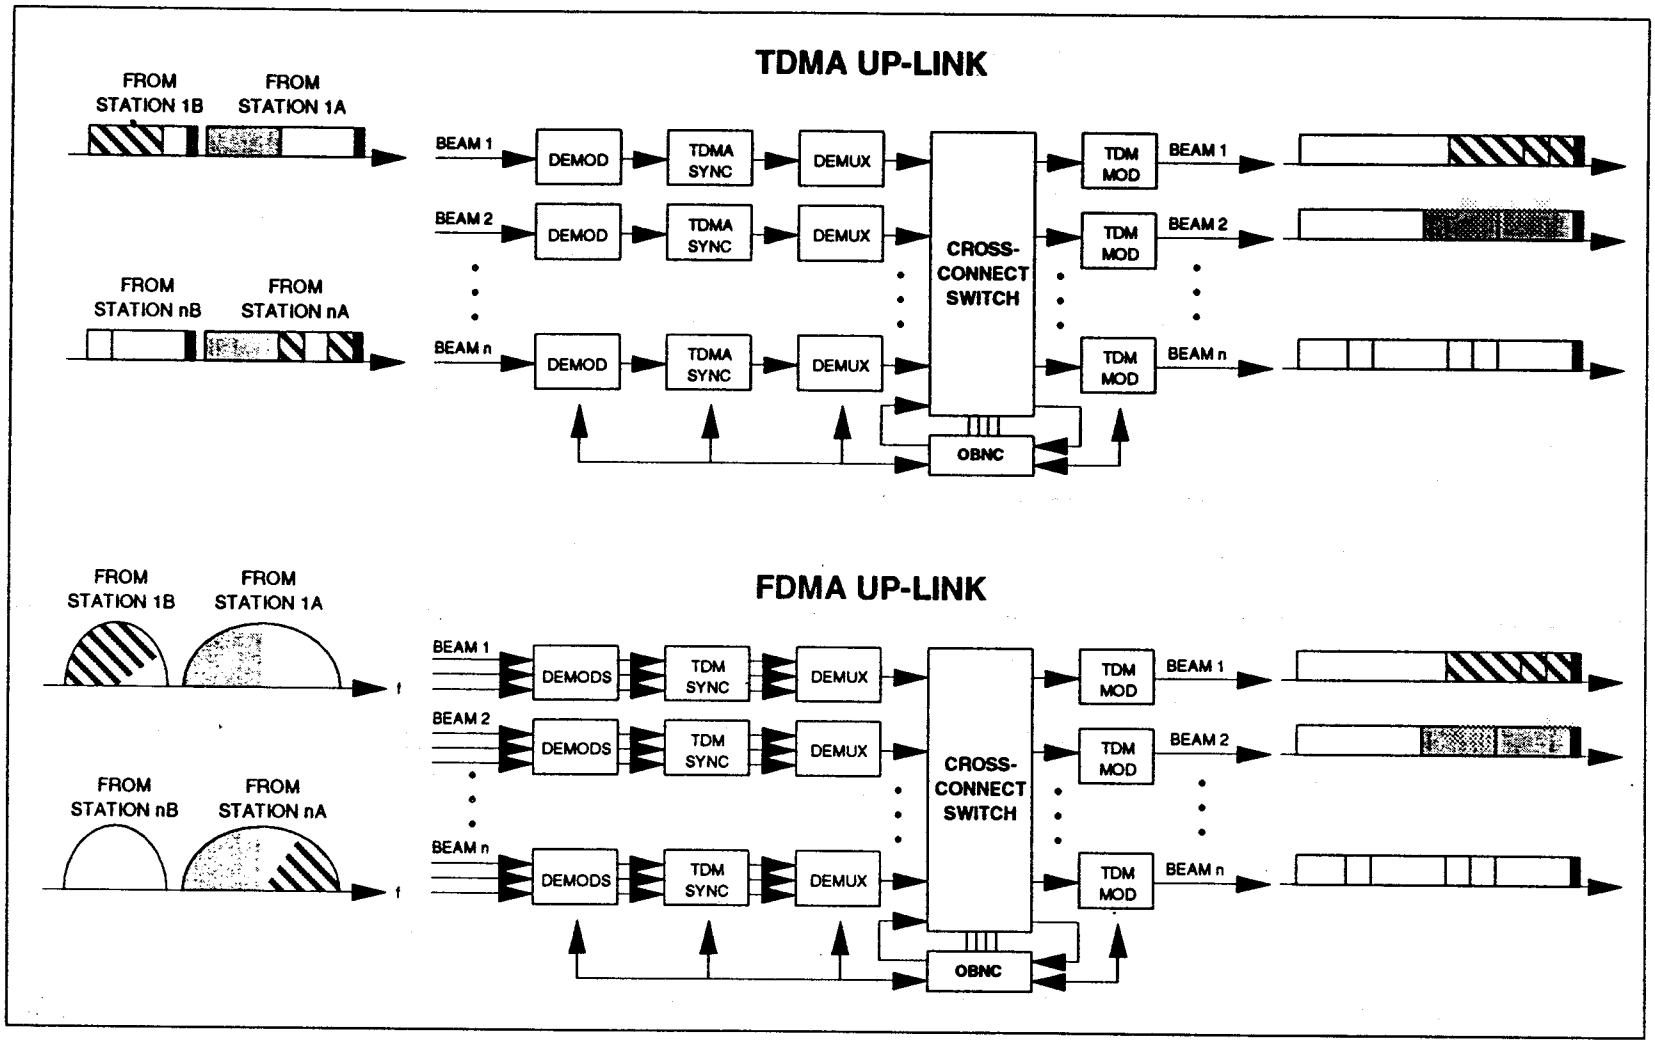

Figure 4-6. On-Board Cross-Connect System

Traffic reconfiguration by the OBNC is possible for the cross-connect system. Consider, for example, a Ka-band system consisting of 100 active spot beams, each having an equivalent of eight 51-Mbit/s channels. Furthermore, assume that the average connection duration for each channel is 20 minutes. Then, a channel request for either allocation or disallocation of capacity will arrive at the OBNC once every 750 ms on the average, or about 1.3 requests per second. The amount of requests can be easily handled by a microprocessor available on the market today.

4.3.4 User Earth Station Processing

Processing requirements at the user earth station include no processing, time division multiplexing (TDM), cross-connect of SONET frames, and ATM virtual path cross-connect. In the case of no processing, the terrestrial interface signal, after synchronization and conversion of the signal into a digital form, will be directly transmitted to the satellite. The other three cases will require either an external multiplexer/cross-connect switch or a baseband signal processor integrated with satellite access equipment. A further discussion is given in the following subsections. The terrestrial interface must conform to the B-ISDN (ATM and STM) or SONET requirements with the bit rates specified in Table 4-4. The terrestrial interfaces considered herein are primarily SONET with ATM or non-ATM payloads.

Table 4-4. B-ISDN and SONET Hierarchy and Bit Rates

STANDARD	INFORMATION STRUCTURE	BIT RATE (Mbit/s)
B-ISDN (ATM & SDH)	ATM/STM-1	155.52
	ATM/STM-4	622.02
	STM-16	2488.32
SONET	STS-1/OC-1	51.84
	STS-3/OC-3	155.52
	STS-9/OC-9	466.56
	STS-12/OC-12	622.08
	STS-18/OC-18	933.12
	STS-24/OC-24	1244.16
	STS-36/OC-36	1866.24
	STS-48/OC-48	2488.32

4.3.4.1 Time Division Multiplexing

In time division multiplexing, the earth station forms a TDMA burst or a TDM frame by multiplexing terrestrial channels having the same destination beam. If the earth station has traffic to multiple destinations, it may need to form multiple TDMA bursts

or multiple TDM carriers. (Note that this may not be necessary for an on-board cross-connect system.) The TDM concept and an earth station block diagram are shown in Figures 4-7 and 4-8, respectively.

Transmit terrestrial interface processing at the earth station includes SONET frame synchronization, timing adjustment, and buffering. The earth station uses a common clock to multiplex terrestrial channels and provides synchronization timing to the terrestrial interfaces. Timing correction is performed by adjusting the administrative unit (AU) pointer of the SONET frame to compensate for the frequency difference between the satellite clock and terrestrial clock. Data streams are fed into compression buffers to be multiplexed at a higher rate.

Receive processing is the reverse operation of the transmit side and includes demultiplexing and expansion of the received channels. AU pointer adjustment may be necessary if terrestrial clocks are used to send data to the network; otherwise, no timing adjustment is required. However, the latter case will deliver a clock with downlink Doppler to the terrestrial network or user equipment.

4.3.4.2 SONET Cross-Connect

In TDM processing, the terrestrial SONET frames received at the earth station must be presorted according to the destination beams or destination earth stations. However, this may not always be the case. It is often more economical to deliver a single high-rate SONET frame from a user premise to the earth station and to demultiplex it and regroup lower bit rate frames based on the destinations. The SONET cross-connect concept is illustrated in Figure 4-9. In the figure, Earth Station A interfaces with three terrestrial lines, i.e., two STS-3 and one STS-12 line. One STS-3 and STS-12 frames are demultiplexed into lower bit rate SONET frames at the earth station, which are remultiplexed to form satellite TDMA bursts or TDM bit streams. At the receiving earth station, all the satellite channels are demultiplexed and regrouped to construct terrestrial SONET frames. The cross-connect function at the earth station can be provided by external commercially available equipment and can be integrated with the earth station baseband equipment.

An earth station block diagram for cross-connect processing is shown in Figure 4-10. The earth station may be regarded as a network switching node and hence must perform frame overhead processing. The terrestrial frame overhead messages are terminated at the earth station and converted into special satellite orderwire messages to communicate with destination stations. Also, it might be necessary to adjust the timing of individual channels demultiplexed from a terrestrial SONET frame, since the frame timing of these channels is independently adjusted through individual frame offset indicators. Once the timing is adjusted, the rest of the earth station processing functions are basically identical to those of the TDM interface described in the previous subsection.

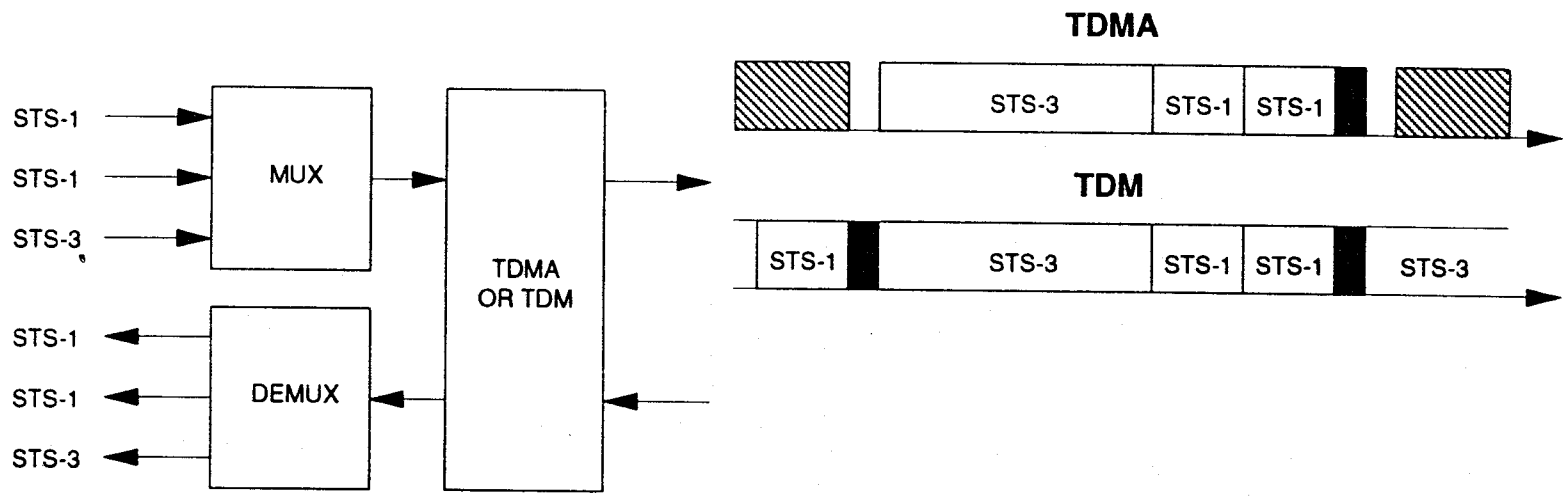

Figure 4-7. TDM Multiplexing

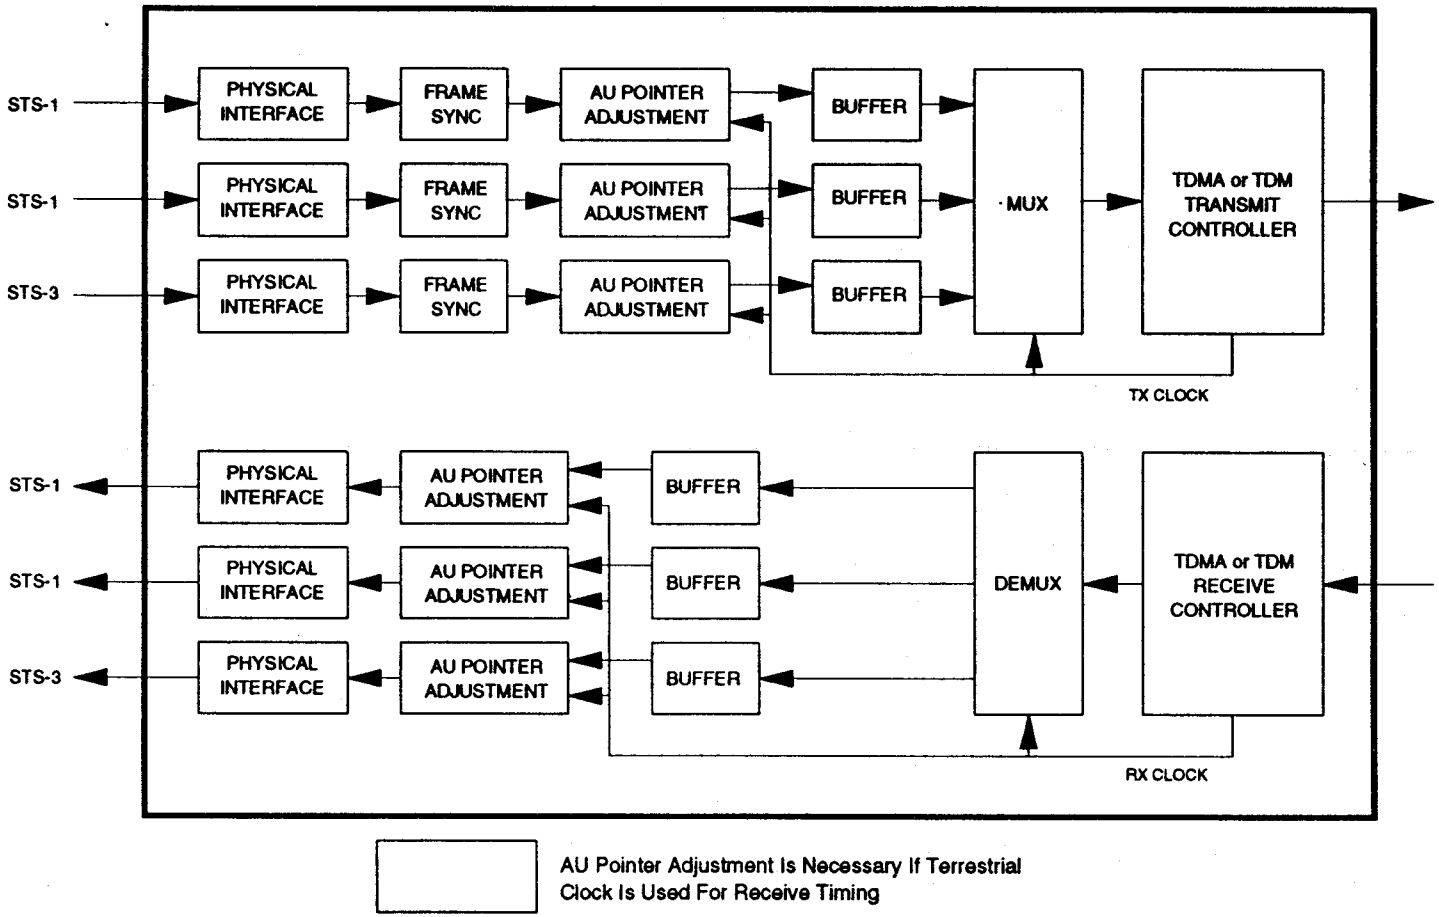

Figure 4-8. Terrestrial Interface Block Diagram for TDM Multiplexing

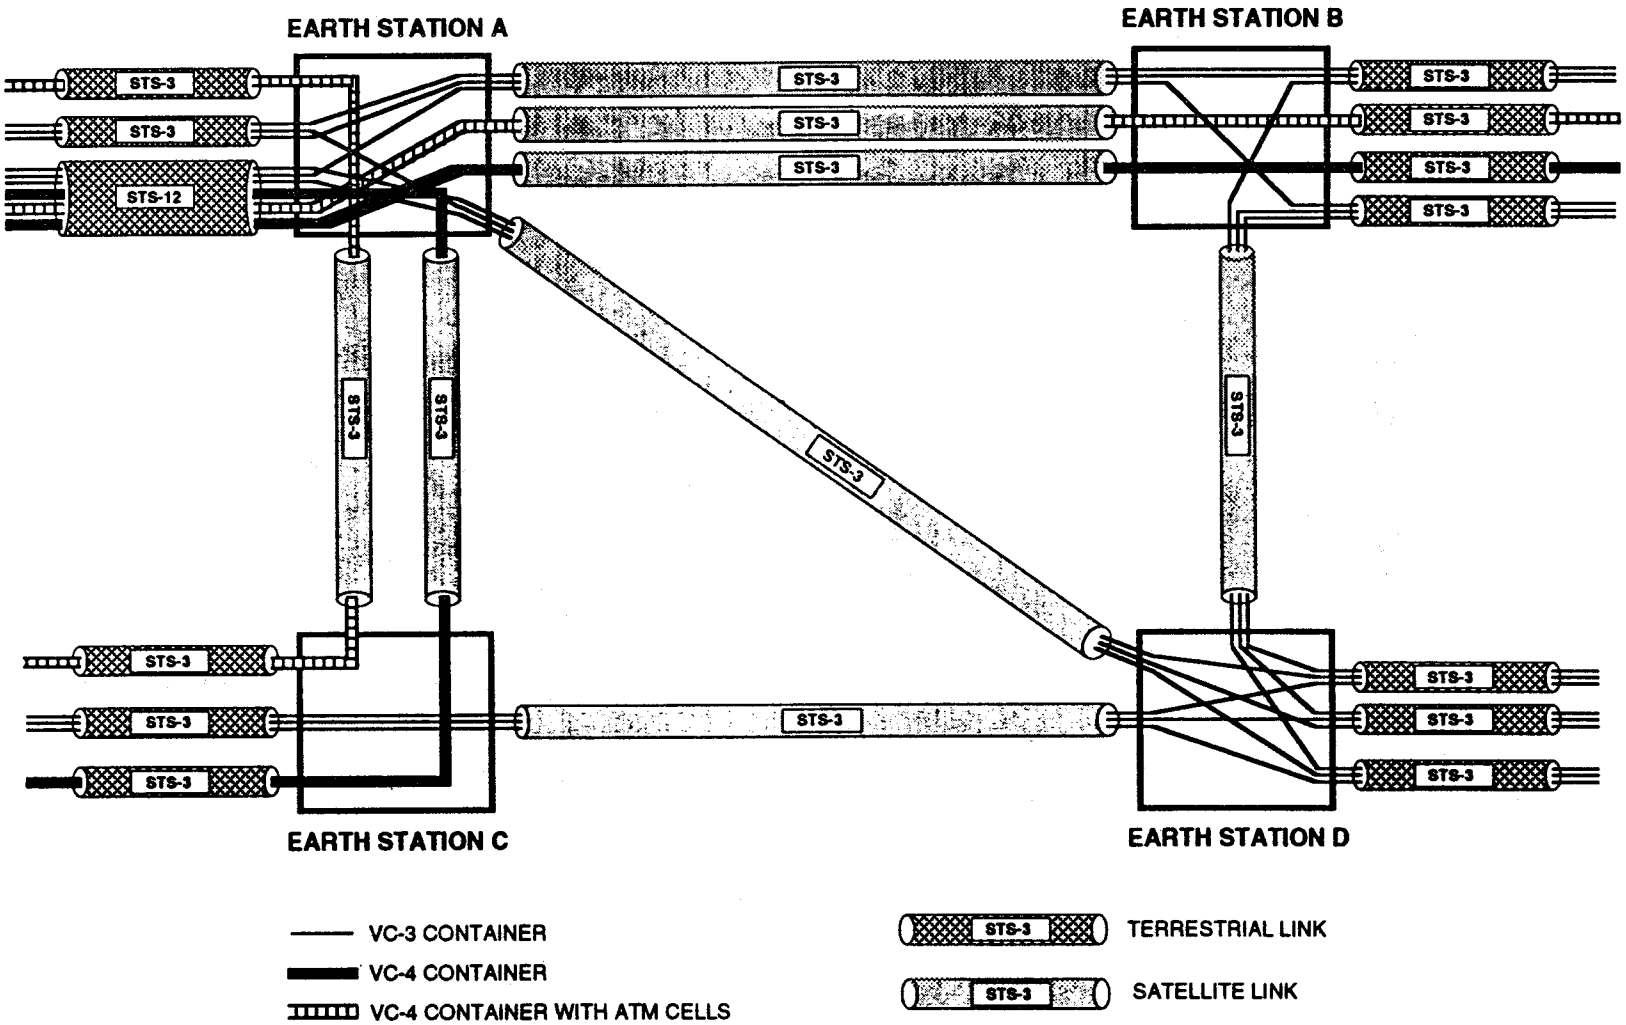

Figure 4-9. SDH Cross-Connect Concept

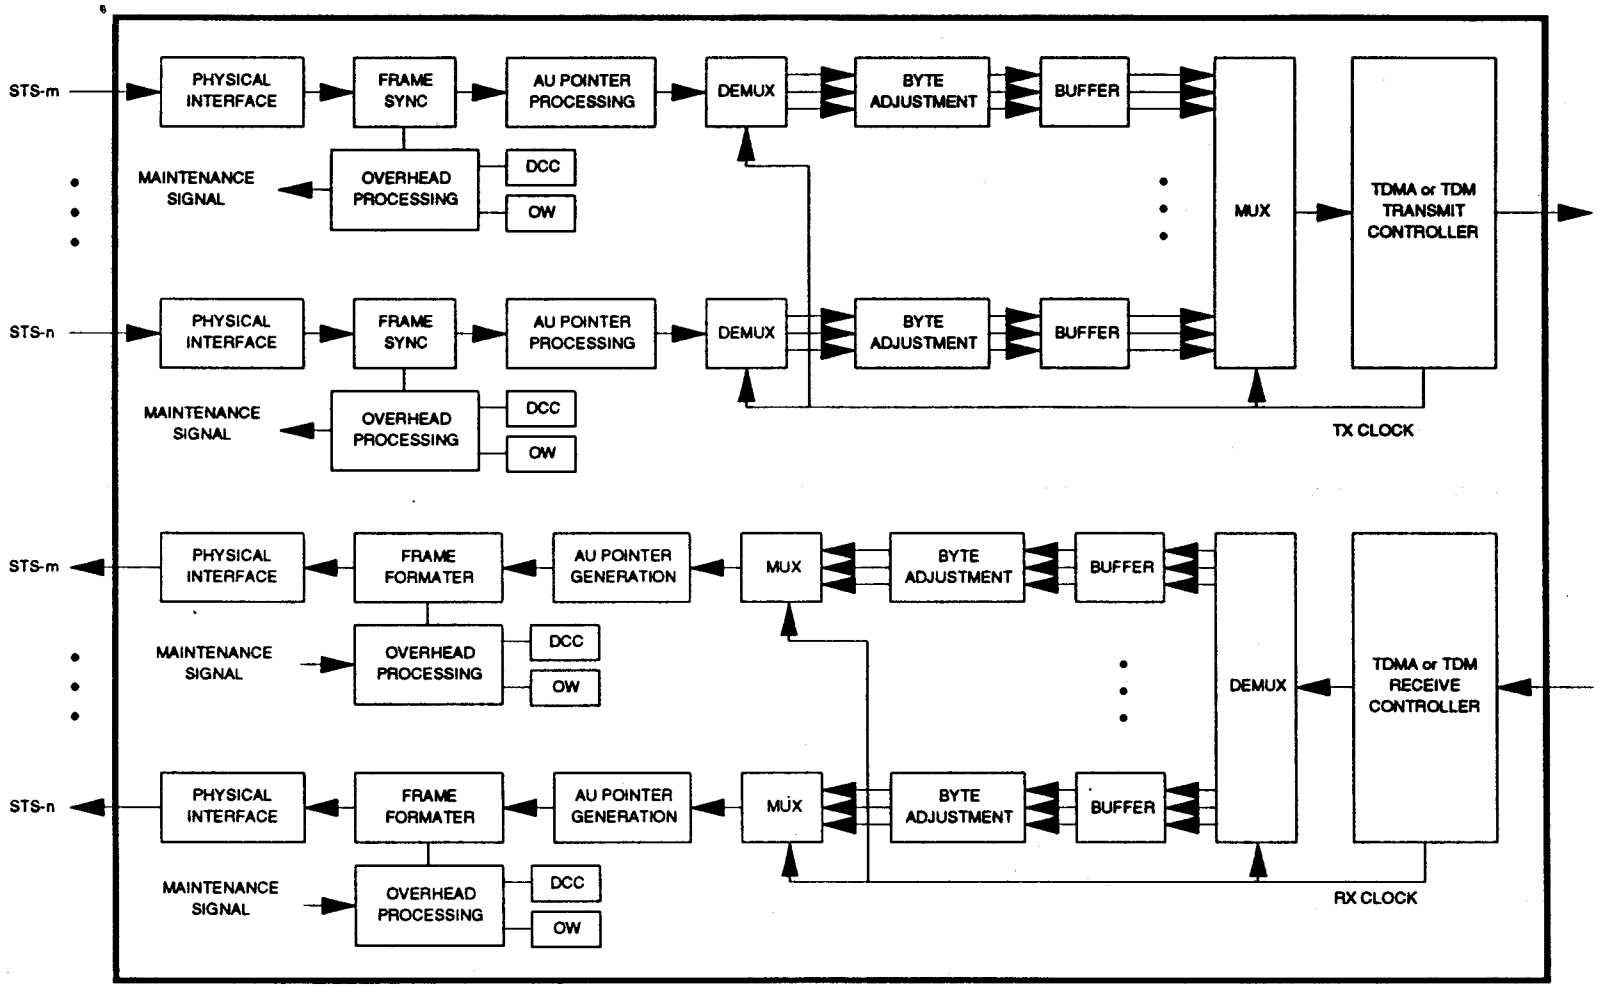

Figure 4-10. Terrestrial Interface Block Diagram for SONET Cross-Connect

4.3.4.3 ATM Cross-Connect

A cross-connect function can also be implemented for ATM cells at the earth station. An ATM cross-connect concept is illustrated in Figure 4-11. In the terrestrial network, ATM cells are transported using one of two formats. The first format uses the SONET frame with ATM cells in the payload. The second format is a pure ATM cell stream. In both cases, earth station processing for ATM cross-connect is almost identical, except for the terrestrial interface functions, such as physical interfaces and synchronization of SONET frames or ATM cells. As in the case of SONET cross-connect, the ATM cross-connect function can be performed using external equipment or can be integrated with earth station baseband equipment. The latter case will result in a more efficient space segment utilization, since satellite capacity can be allocated with a better granularity than the terrestrial B-ISDN bit rates.

The ATM cell contains a virtual channel identifier (VCI) and a virtual path identifier (VPI) in its header. The VCI designates a specific traffic channel which is associated with a traffic source. The VPI indicates a routing path of a particular ATM cell. The earth station sorts ATM cells according to the destination stations (or beams) based on their VPIs. In the example shown in Figure 4-11, one STS-3 frame contains an STS-1 frame with ATM cells, and the other STS-3 frame includes two STS-1 frames containing ATM cells. The earth station demultiplexes the STS-3 signals into STS-1 frames and processes ATM cells through a virtual path cross-connect switch. The output of the switch will be groups of sorted ATM cells with different destination beams. These ATM cell groups are transmitted to the satellite as subbursts in TDMA or as satellite channels in TDM.

One problem associated with ATM cross-connect is an unpredictable traffic flow to various destination beams. Without flow control, it is likely that the amount traffic to a certain beam will exceed the allocated space segment capacity. Thus, the earth station must perform flow/congestion control to minimize ATM cell loss.

4.4 Fast Packet Switched System

On-board fast packet switching provides flexibility and efficient bandwidth utilization for packet switched traffic. Its self-routing feature also benefits circuit switched traffic users if flexibility and frequent network reconfiguration are required. Application of fast packet switching is not limited to only B-ISDN traffic, but its generic structure can accommodate any type of packet or circuit-based traffic. Figure 4-12 depicts an on-board fast packet switching system with TDMA and FDMA uplink accesses.

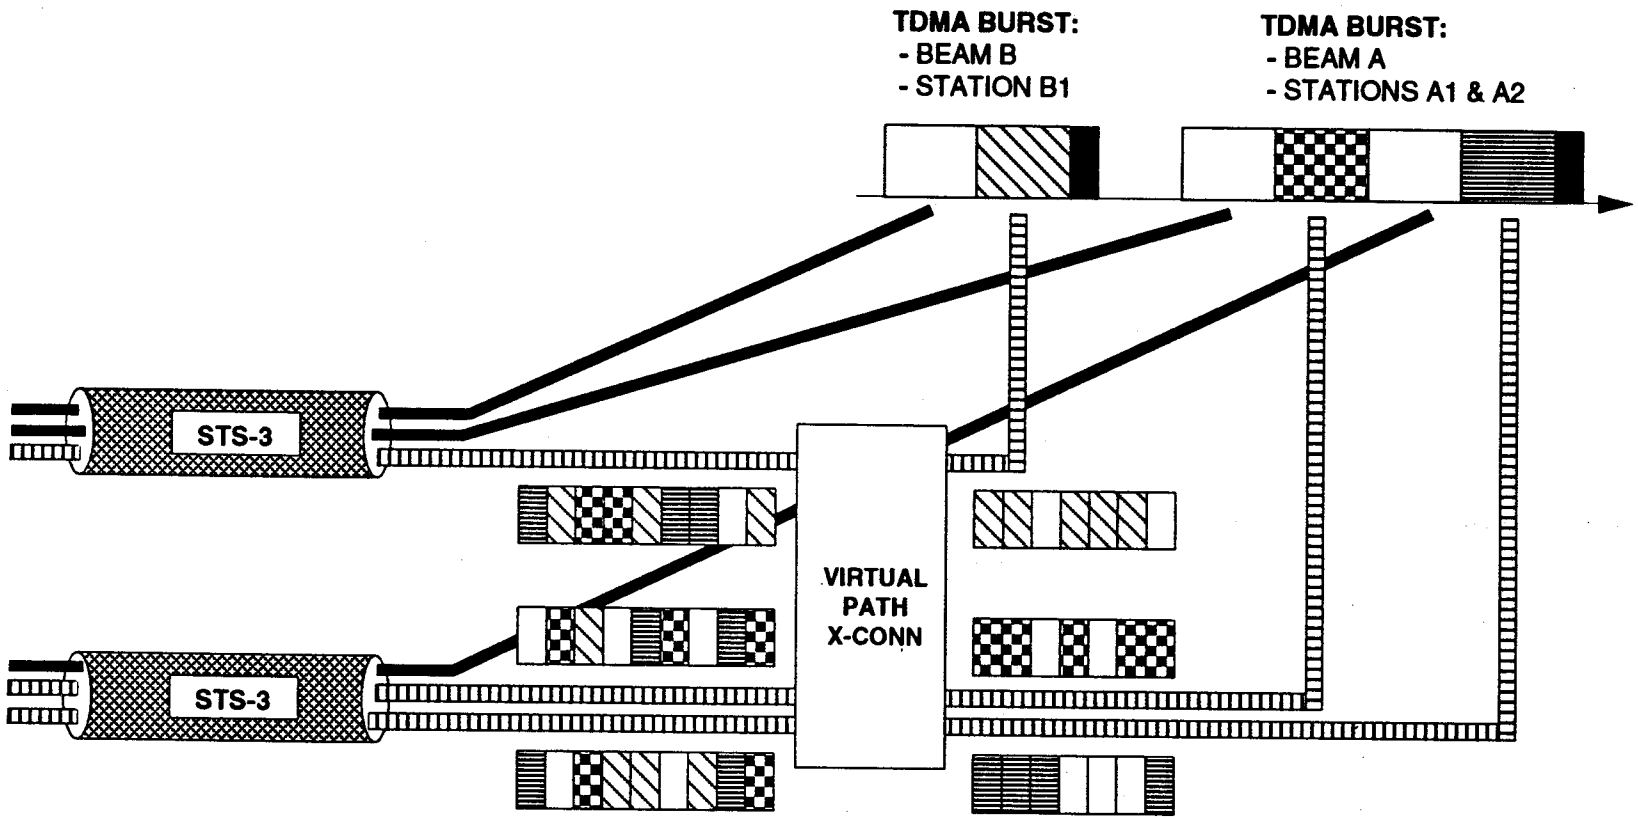

Figure 4-11. ATM Virtual Path Cross-Connect Concept

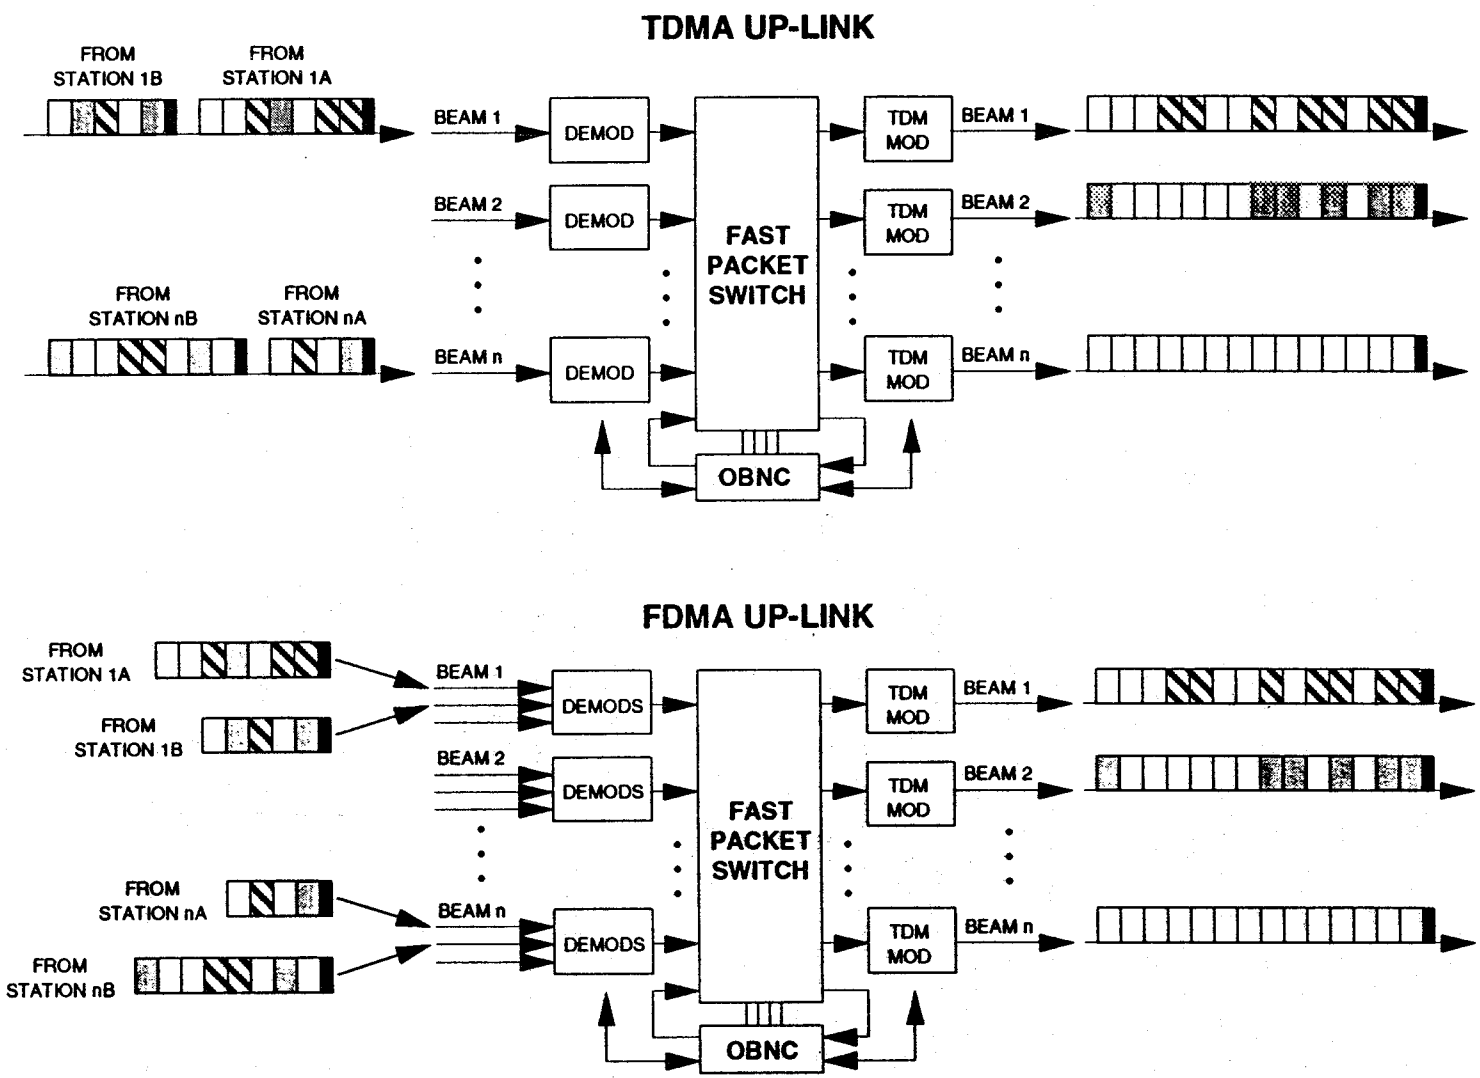

Figure 4-12. On-Board Fast Packet Switching

A potential fast packet switching architecture for B-ISDN includes:

A. a general purpose fast packet switch,

B. an ATM switch with permanent VPIs, and

C. an ATM switch with dynamically assigned VPIs

Architecture A uses a satellite virtual packet (SVP) to carry ATM cells as well as other types of traffic. The SVP is basically a transmission container, and the user can fill the container with any data that he wishes. The satellite fast packet switch simply routes SVPs based on the routing information included in the SVP headers. This architecture is most flexible and particularly suited for integrated operation of circuit and packet switched services. An earth station functional block diagram for Architecture A is shown in Figure 4-13.

In Architecture B, each earth station is assigned a set of distinct VPIs which are exclusively used for communications with prespecified earth stations. The VPI uniquely defines the connection of two specific earth stations. Some VPIs may be reserved for multicast connections on a permanent basis and others assigned on a demand basis. A 16-bit VPI gives over 65,000 different connections. A VCI is associated with a traffic channel, and a combination of a VPI and VCI uniquely identifies a specific traffic channel in the network. The on-board fast packet switch obtains routing information from VPIs by a table look-up method. Alternately, the VPI can be structured such that the first several bits constitute a routing tag. In this case, no on-board look-up table will be necessary. Although, Architecture B is customized to ATM cells, the same routing procedure is used for Architecture A.

Architecture C is conceptually the same as implementing an ATM switch on board the satellite. While earth station processing is minimal for ATM traffic, the on-board processor must perform similar functions as an ATM switch: assignment of VPIs (and VCIs), VPI translation from uplink to downlink, and a mapping of a VPI to obtain routing information. Architecture C is not recommended for implementation because of its complexity.

4.5 Sample Network Architectures

Three sample network architectures which are applicable to satellite B-ISDN services are investigated in more detail. Each sample architecture is geared towards a specific array of B-ISDN applications. Architecture 1 is intended for high speed transmission at B-ISDN rates of slowly varying circuit-switched traffic, as in trunking applications and HDTV distribution. Architecture 2 is more applicable to business networking and thin route traffic where traffic volume is usually low but traffic burstiness is high coupled with large variations in user traffic requirements. Architecture 3 focuses on NASA science networking applications, which include a wide variety of communications requirements ranging from high speed near-continuous transmission to low speed packet type computer networking.

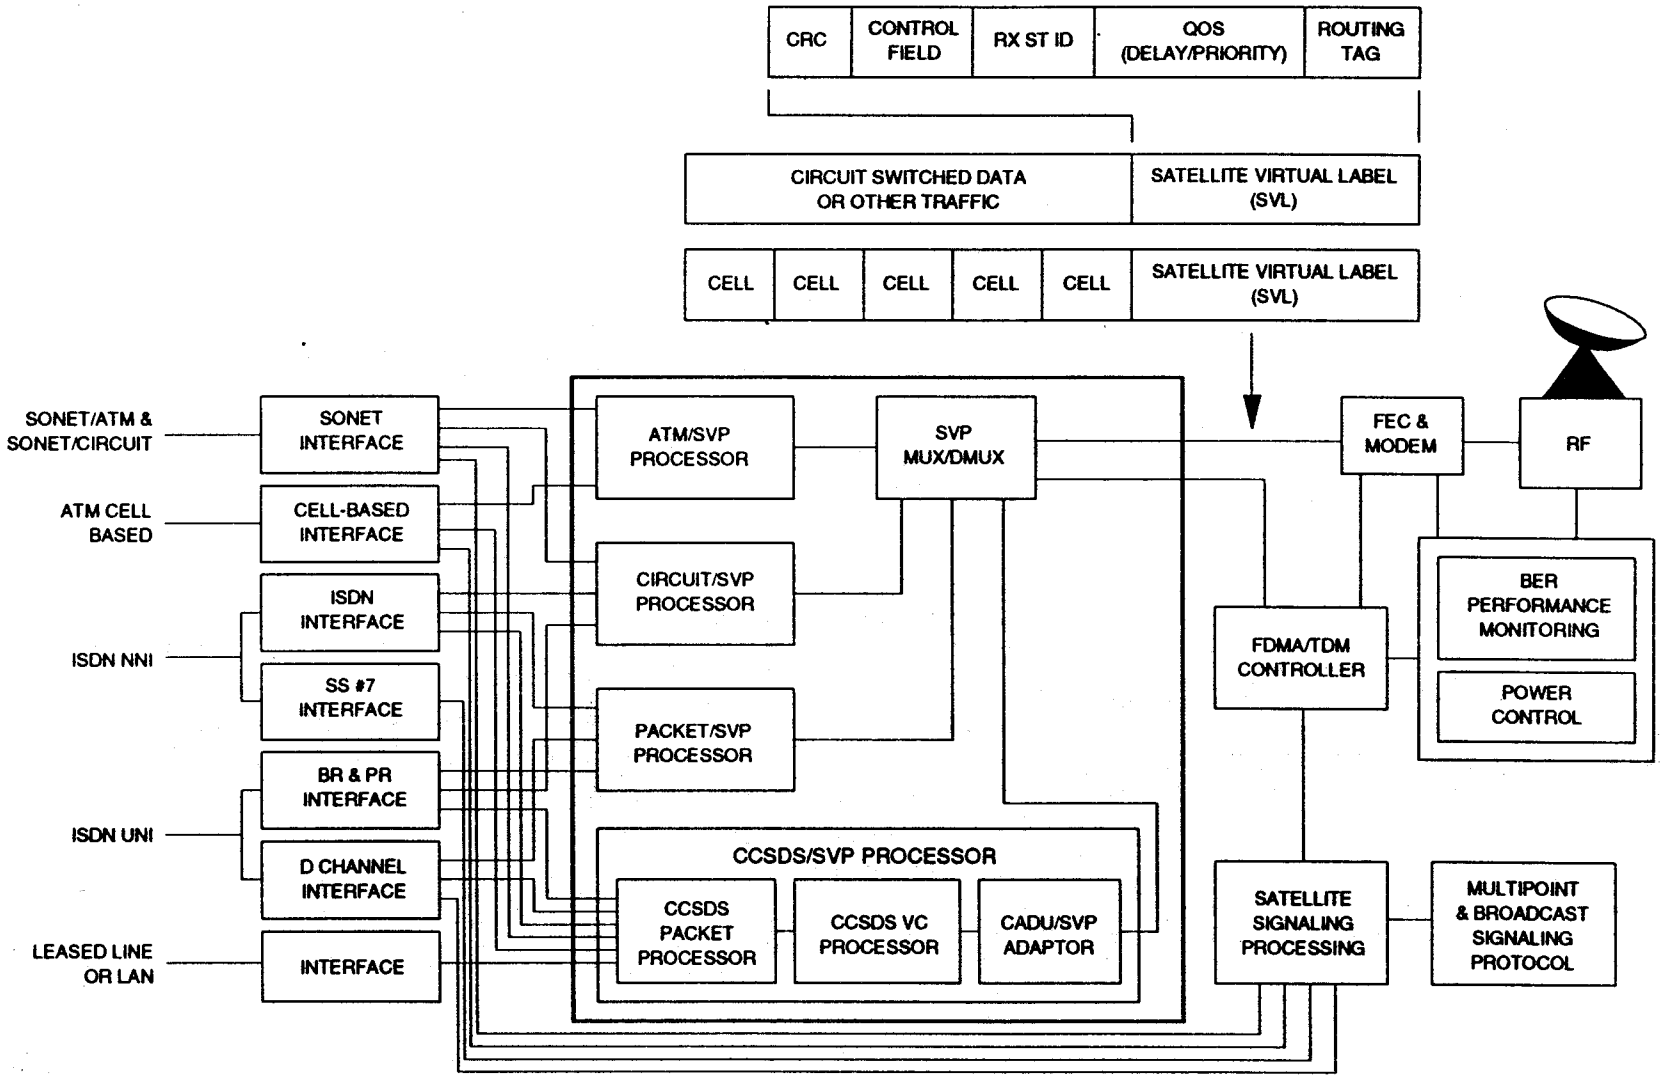

Figure 4-13. Earth Station Block Diagram for On-Board Fast Packet Switching

4.5.1 Architecture 1: High Speed Circuit-Switched Applications

This architecture is primarily intended to provide point to point and point to multipoint connections for high speed circuit switched applications. Applications include high speed trunking, HDTV program distribution, and cable restoration. Other applications may include large database transfers and medical image transfers.

The traffic requirements for such a system are for high bit rate transmission from bit rates starting at 51 Mbit/s to as high as 1.2 Gbit/s and possibly 2.4 Gbit/s. The traffic is mostly a circuit switched trunk type which does not require frequent reconfiguration. Satellite capacity is allocated for the establishment of connections on an as needed basis. In some instances, a very high speed interconnection is required between two points, as in the case of cable restoration. In others, a moderately high bit rate carrier is needed in broadcast mode as in the case of HDTV distribution to movie theaters by the respective production studios.

The requirements on the size of the earth station are rather relaxed in that for most of these applications we are not dealing with customer premise earth stations. For trunking, it is more likely to have the connection between two gateways which utilize large size earth stations, for example with 12-meter antennas. However, for broadcast applications, such as TV program distribution, small receive-only earth stations are desirable.

Satellite network architectures for high-speed circuit switched applications include SS-TDMA, SS-FDMA, and on-board cross-connect systems. The sample network architecture presented in the following uses SS-FDMA for uplink access in the Ka-band. Figure 4-14 shows the antenna coverage of the system. The satellite covers the continental United States (CONUS) with 16 steerable beams of 0.35°. Each steerable beam can be parked on any one of 200 spot areas and support a maximum throughput of 622 Mbit/s. A payload configuration is shown in Figure 4-15. On-board demodulators are programmable and can operate at any rate of an integer multiple of 51 Mbit/s but not exceeding 622 Mbit/s. Four on-board demodulators are allocated for each steerable beam. Programmability of these demodulators provides a fail safe operation without significantly degrading the built-in flexibility. As technology warrants in the future, the four demodulators can be replaced by a single multicarrier demodulator. Downlink transmission is continuous TDM at a bit rate of 622 Mbit/s to each steerable beam. The total system capacity is about 10 Gbit/s.

The interconnection between uplink and downlink beams is provided by a 64 x 64 baseband space switch. Although the connectivity of the SS-FDMA system is limited to 16 uplink and 16 downlink beams, any uplink carrier can be routed to any downlink beam within the constraints of beam capacity and a number of output ports (up to four). The baseband crossbar switch allows point-to-multipoint connections as well. The 64 x 64 switch matrix may be constructed from four LSI chips of 32 x 32 switch matrices, as shown in Figure 4-16, to form a strictly non-blocking switch.

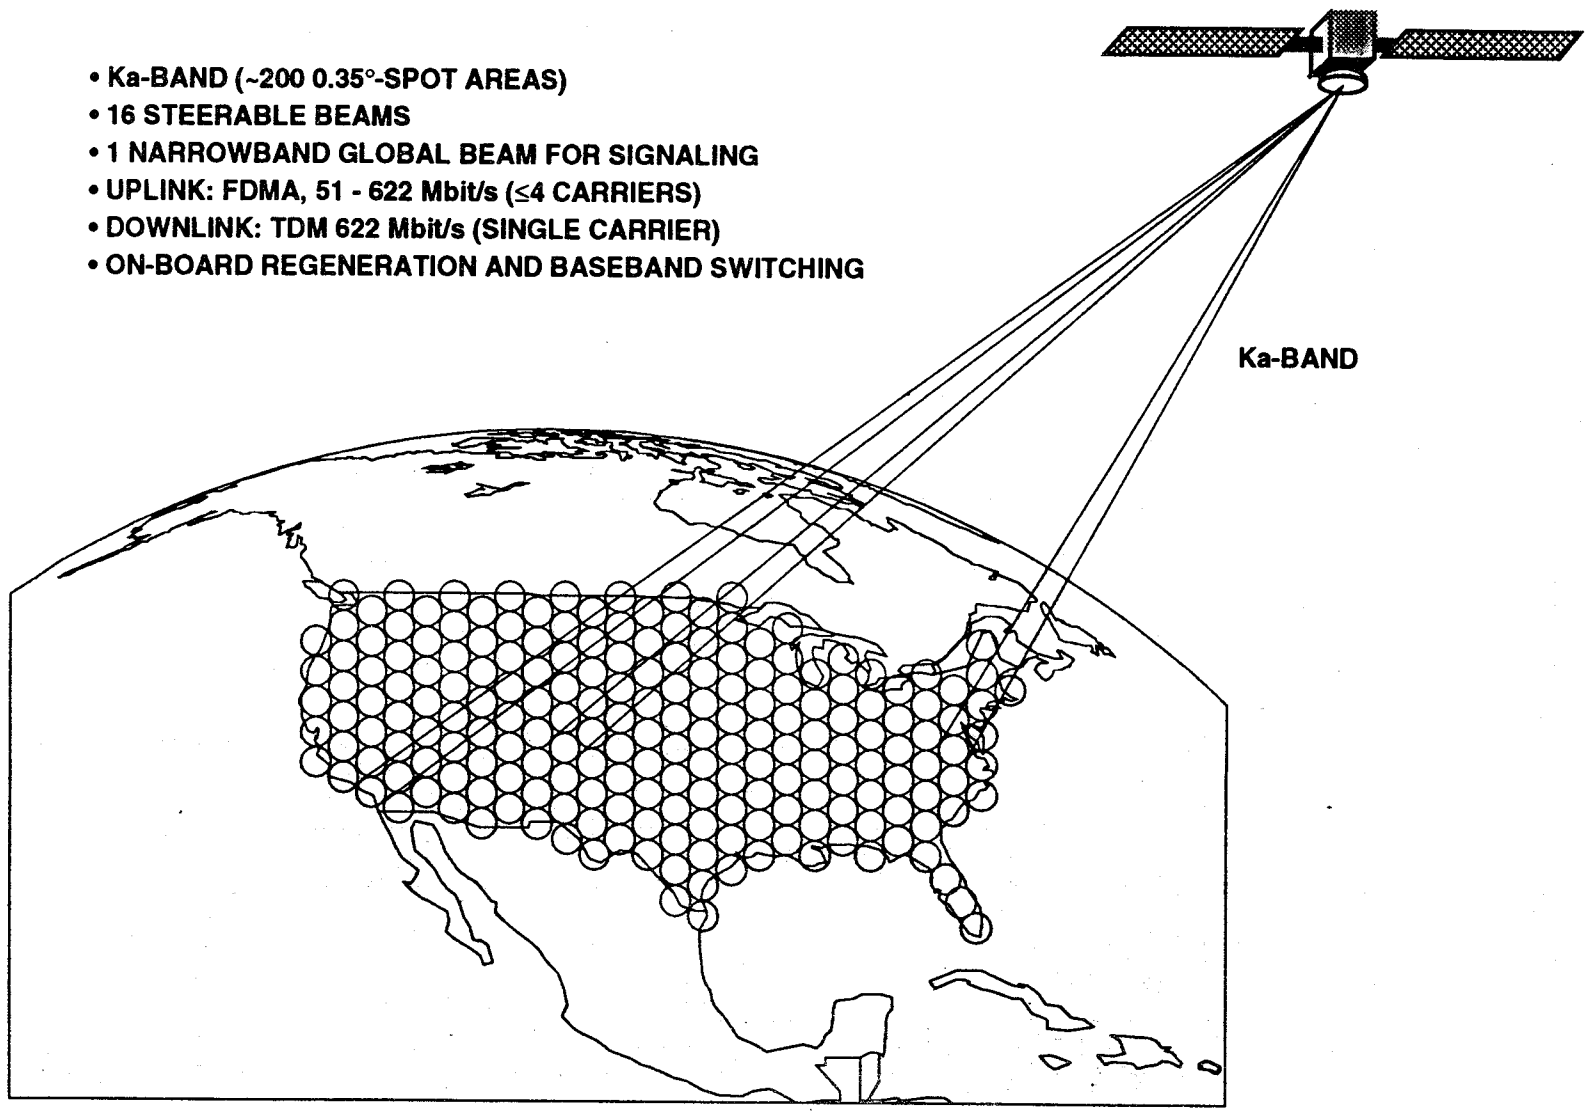

Figure 4-14. Ka-Band Antenna Coverage for High-Speed Circuit Switched System

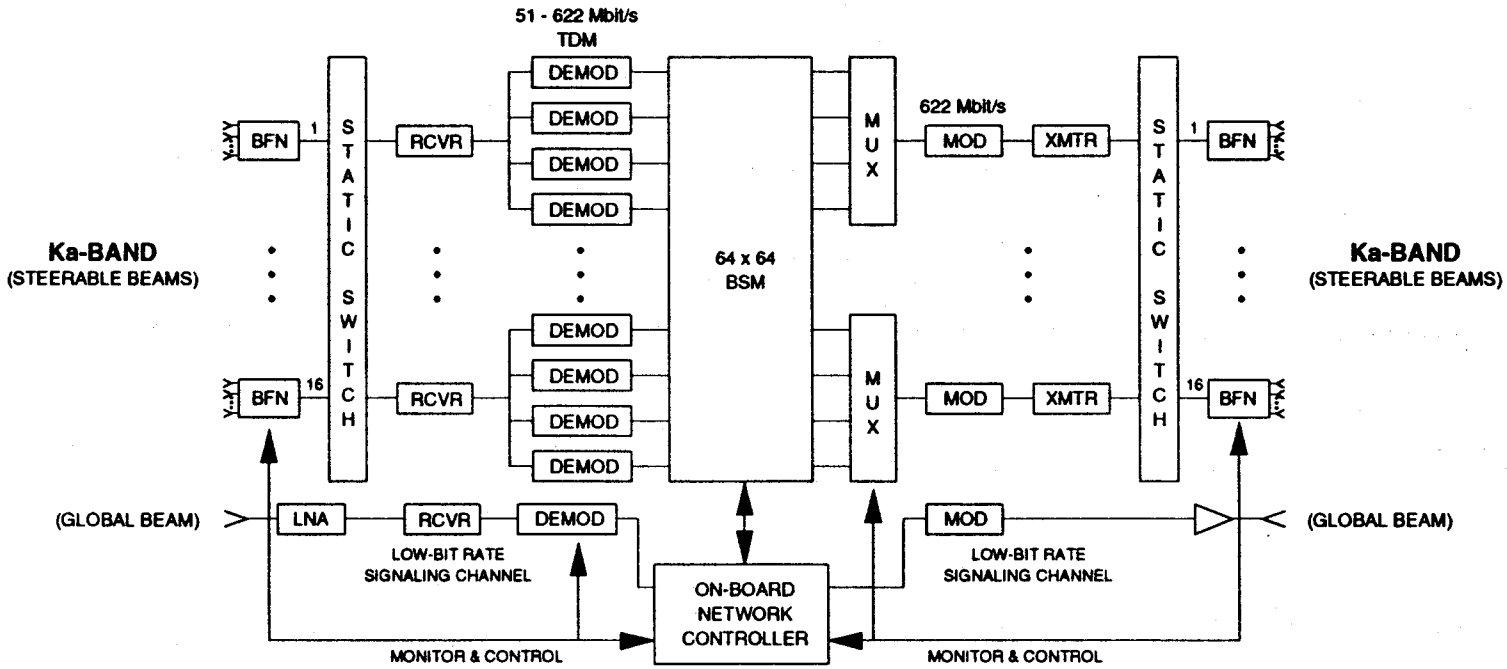

Figure 4-15. SS-FDMA Payload Configuration for Circuit Switched Applications

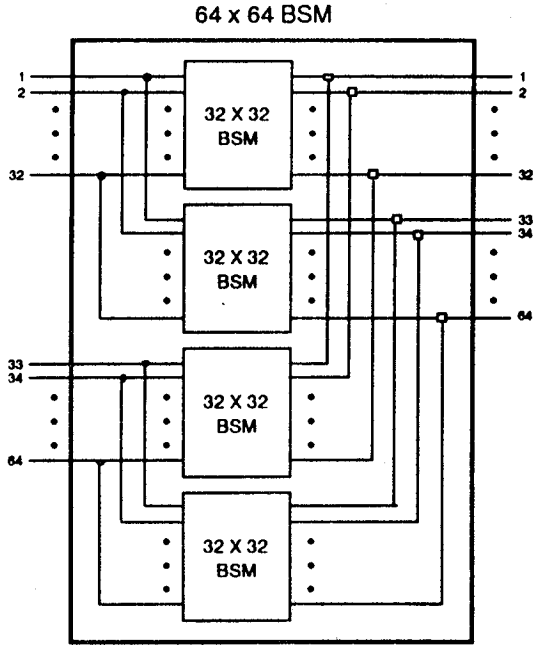

Figure 4-16. A 64 x 64 Baseband Switch Matrix Composed of 32 x 32 Switch Matrices

Since the system does not provide full connectivity, an additional global beam for a signaling channel is included in the design. This beam is used for receiving channel request and status messages from earth stations and transmitting station control messages from the satellite. Earth stations transmit signaling messages in random access packets and receive command/control messages from the OBNC in burst TDM. The OBNC reallocates its available resources (demodulators, BSM, modulator, and uplink/downlink carrier frequencies) to the requesting stations. If some beam does not have traffic, the OBNC will assign the resources to a needed spot area. The use of a low-speed global beam (e.g. 64 kbit/s) allows autonomous operation of the SS-FDMA system. Alternately, signaling messages can be communicated with a ground control station via the terrestrial network. This will eliminate the need for an additional global beam and simplify OBNC processing.

QPSK modulation along with a bandwidth-efficient Reed-Solomon (RS) code (255, 223) will require a transmission bandwidth of 35 MHz (at 51 Mbit/s) to 427 MHz (at 622 Mbit/s). The size of an earth station ranges from a 2.4-m antenna with a 3-watt HPA for 51 Mbit/s transmission to a 5-m antenna with a 100 watt HPA for 622 Mbit/s transmission, providing the respective availability of 99.6 and 99.9 percent at a bit error rate of 10^{-8}. The 2.4-m earth station with a 36-watt HPA can also support transmission at 622 Mbit/s, however, with a reduced availability. More bandwidth-efficiency and better transmission performance can be achieved by the use of octal PSK modulation combined with concatenated FEC coding of a high rate convolutional code and an RS code but with additional complexity at the earth station and the satellite.

The SS-FDMA system generally results in a simpler and lower cost earth station, in particular for single carrier transmission — small antenna and HPA size and simpler baseband equipment. The system, however, has a limitation in term of connectivity. A more flexible system for high-speed circuit switched applications uses SS-TDMA for access. A payload configuration for 622-Mbit/s SS-TDMA system is depicted in Figure 4-17. The system uses the same number of beams as SS-FDMA but includes a hopping capability. To effectively implement this system concept, the CONUS may be divided into 16 sectors, where each sector covers between 10 and 15 spot areas and shares a single 622-Mbit/s TDMA carrier. An earth station can transmit one 51-Mbit/s channel or can be dedicated to the entire 622-Mbit/s transmission. The on-board processor payload requires 16 active demodulators at 622 Mbit/s and the same number of modulators. The baseband switch size is 17 x 17 and is significanly lower than that of the SS-FDMA system. The overall payload complexity is, however, about the same or slightly simpler than the previous system. The OBNC communicates with user earth stations for channel requests/allocation through TDMA signaling bursts.

The earth station size required for this system is the same as the 622-Mbit/s SS-FDMA station, ranging from a 2.4-m antenna with a 34-watt HPA (99.6 percent availability) to a 5-m antenna with a 100 watt HPA (99.9 percent availability). Earth station baseband equipment must perform TDMA synchronization, burst compression/expansion, and multiple burst transmission (i.e., multiple traffic bursts and a synchronization/signaling burst). To implement hit-less network reconfiguration, the system must also be capable of performing coordinated time plan changes between earth stations and the OBP. A more complex and possibly higher cost earth station, however, provides flexible network interconnection and may be justifiable for a high-speed trunking service.

4.5.2 Architecture 2: Business/Thin-Route Applications

The second sample architecture deals with lower speed applications with bursty traffic as well as some circuit switched traffic. Network applications include private and business networks, thin-route communications, and emergency communications.

This type of network is characterized by extremely bursty traffic for data, a continuous low-speed circuit connection between two locations for stream data and voice, a mix of high and low speed data (e.g. VSAT hubs and VSAT user terminals), quick deployment, and a large number of users (private and business networks). Efficient utilization of satellite capacity and low cost user earth stations are critical in this environment to compete with terrestrial based services. Network user interfaces may not be limited to B-ISDN and include N-ISDN, LAN/MAN, and X.25/X.25.

Figure 4-18 illustrates the antenna beam coverage for such a system. The system uses narrow Ka-band spot beams for TDMA transmission at 155 Mbit/s and wider (1.5° beamwidth) Ku-band spot beams for low-speed multi-frequency TDMA (MF-TDMA) transmission. As in the previous case, the Ka-band payload employs 16 hopping beams, covering the CONUS, while the Ku-band payload consists of 12 fixed spot beams, each

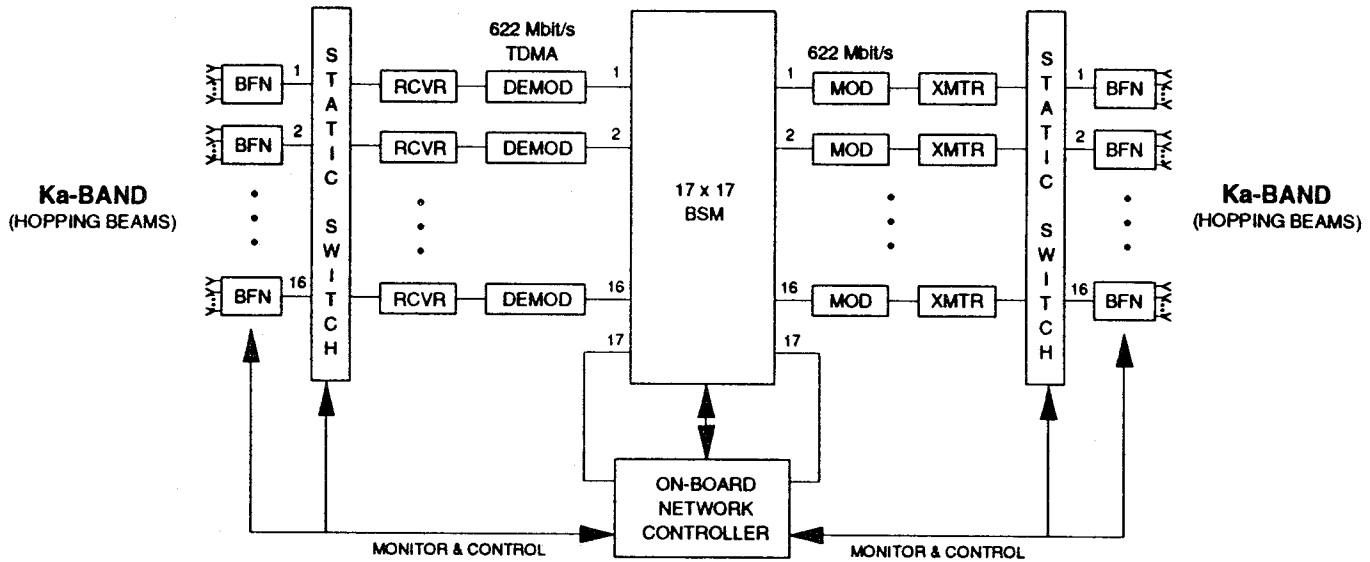

Figure 4-17. SS-TDMA Payload Configuration for High-Speed Circuit Switched Applications

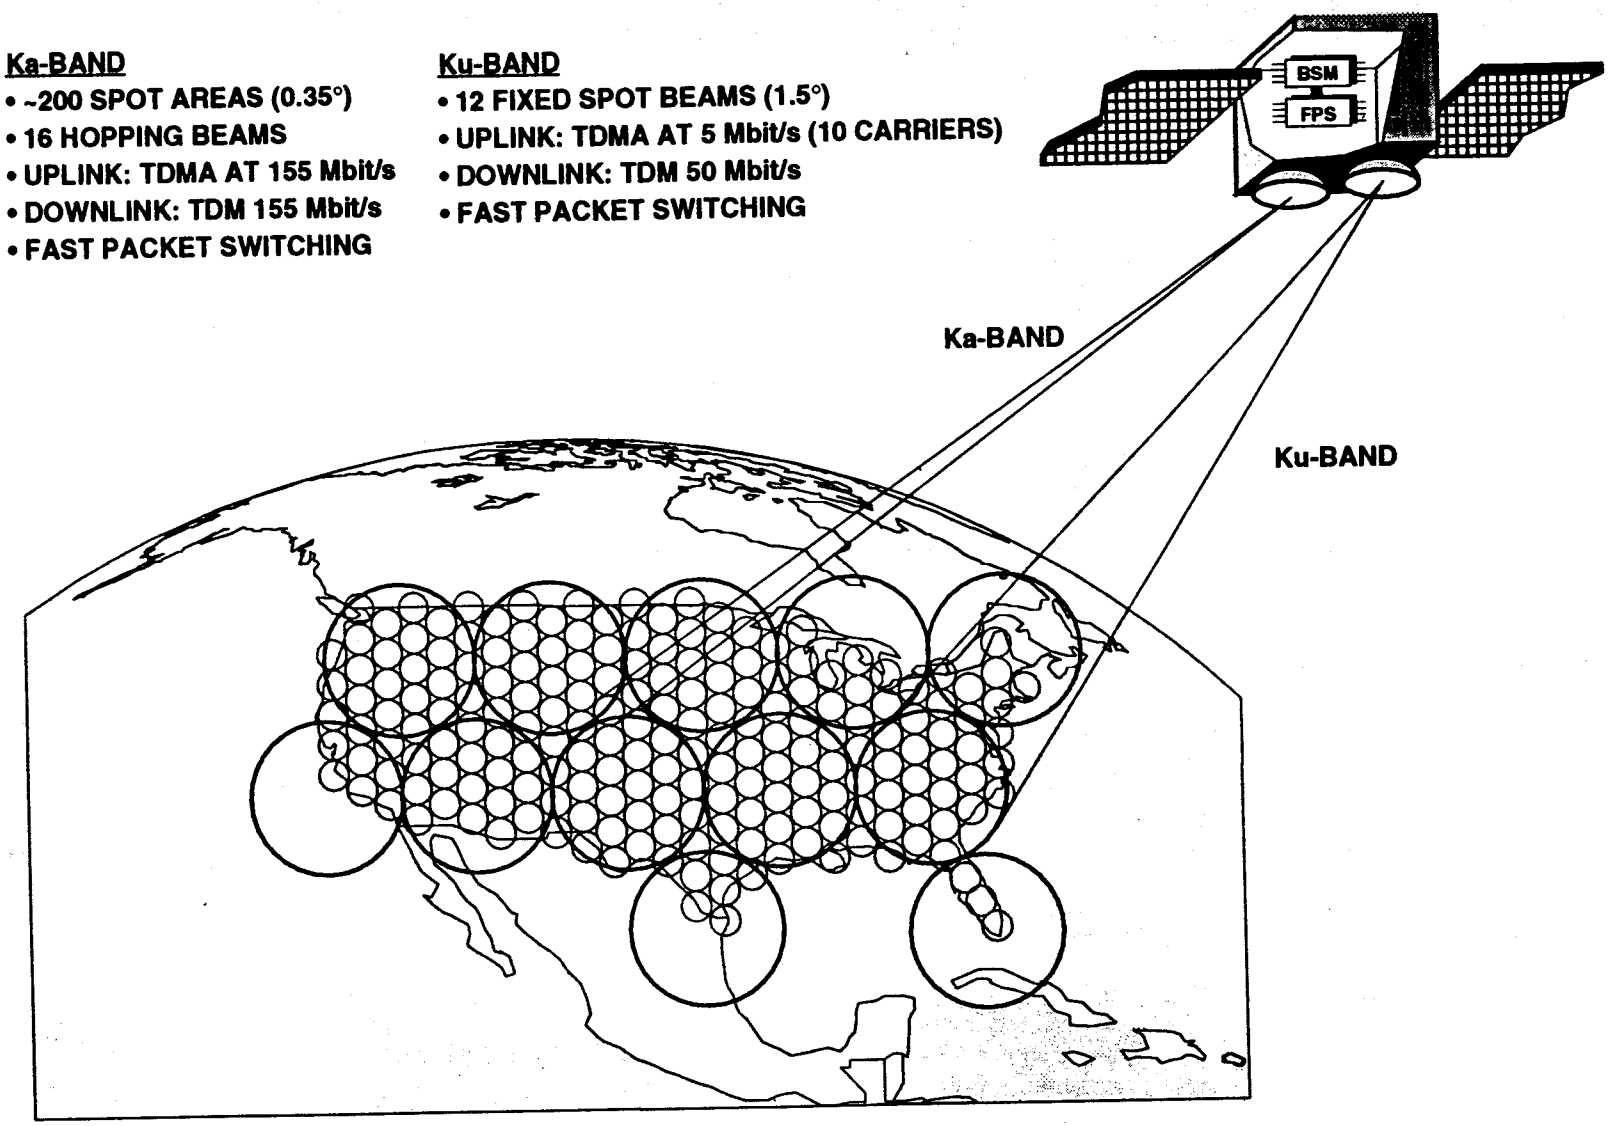

Figure 4-18. Antenna Beam Coverage for Business/Thin-Route System

having 10 uplink TDMA carriers at 5 Mbit/s. The downlink transmission at the Ku-band is a single 50 Mbit/s TDM carrier. The payload configuration of this system is shown in Figure 4-19. Flexible interconnection between Ka- and Ku-band beams is provided by a fast packet switch. For each Ku-band beam, the OBP utilizes a multi-carrier demodulator (MCD) for low power and low weight implementation. The system capacity is 2.48 Gbit/s in the Ka-band, 600 Mbit/s in the Ku-band, and a combined total of 3.08 Gbit/s.

Transmission is QPSK modulation with a rate-1/2 convolutional code and Viterbi decoding. A rate-1/2 convolutional code is not bandwidth efficient, but yields a better performance and graceful degradation at a higher link bit error rate. This is important for VSAT applications to achieve a high data throughput in a degraded link condition. A high availability Ka-band earth station requires a 2.4-m antenna with a 100 watt HPA to provide 99.9 percent link availability at a bit error rate of 10^{-8}. A 1.2-m earth station with a 42-watt HPA can also support 155 Mbit/s transmission, but its link availability will be reduced to 99.7 percent, primarily due to the downlink performance. A typical Ku-band earth station uses a 1.2-m antenna and a 10-watt SSPA to achieve 99.8 percent availability at a bit error rate of 10^{-8}.

The fast packet switching payload provides full connectivity between Ka- and Ku-band beams and is very flexibly and efficient for accommodating a mix of circuit and packet switched traffic. However, for a network with a relatively large amount of high-speed circuit switched traffic, a combination of SS-TDMA and fast packet switching payloads may result in a more efficient configuration. A cross-strapped configuration is shown in Figure 4-20. The Ka-band SS-TDMA system is used primarily for high-speed circuit switched applications, and 25-percent of its capacity (or a 100-percent of the Ku-band capacity) is cross-strapped with the Ku-band payload. While this architecture provides full connectivity between the two types of beams, traffic reconfiguration for Ka-to-Ka band will be less dynamic and also require more complex earth station control for time plan changes. The total system capacity is the same as that of the first configuration.

4.5.3 Architecture 3: NASA Science/Networking Applications

The traffic requirements for NASA science data networking and intra-agency communications include high-speed trunking and data distribution from fixed sites (mainly NASA centers), interactive data transactions for data archives, voice and video teleconferences, high resolution video/image transmission, and high-speed intersatellite link (ISL) transmission. The major NASA centers may require large earth stations for high-speed multiplexed data/voice/video/image transmission. On the other hand, experimenter, university, and scientist locations may not require high-speed data links, and affordable low-cost earth terminals are of the primary interest. These requirements can be translated into somewhat a mixed network configuration of the previous two architectures.

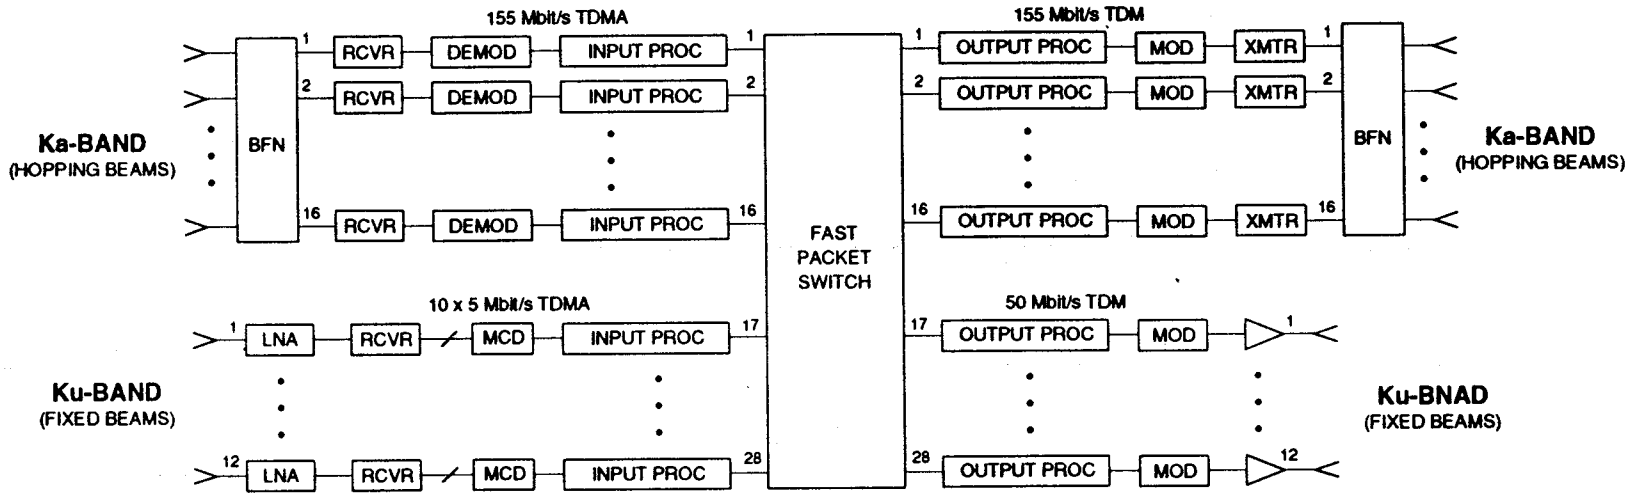

Figure 4-19. Fast Packet Switching for Business/Thin-Route Applications

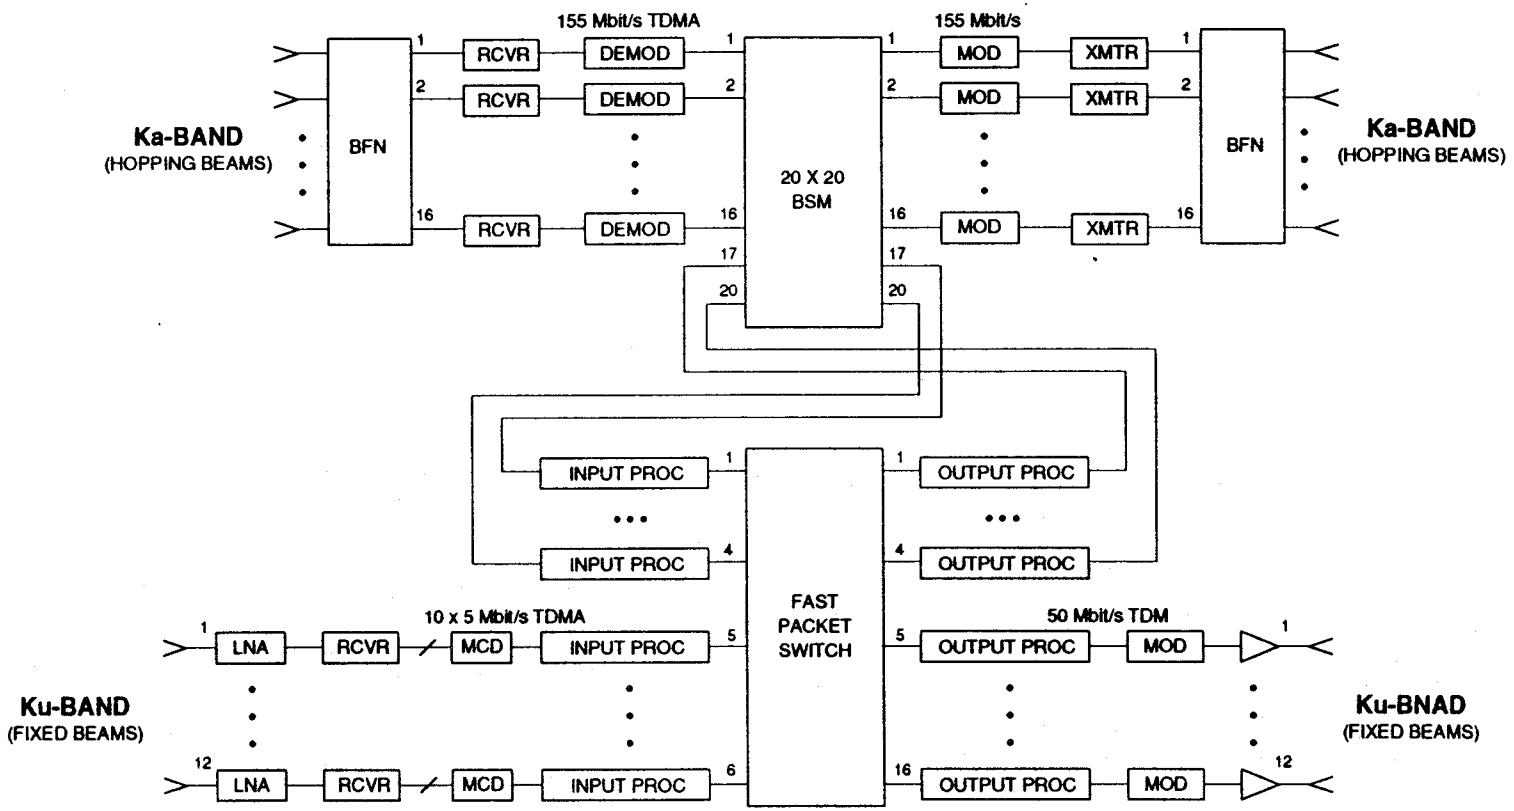

Figure 4-20. SS-TDMA and Fast Packet Switching for Business/Thin-Route Applications

A sample architecture for NASA applications is shown in Figure 4-21. The Ka-band system consists of six hopping beam sectors and 10 fixed beam locations for uplink coverage and six hopping beam sectors for downlink. Uplink transmission is either TDMA at 100 Mbit/s or burst TDM at 622 Mbit/s depending on traffic characteristics, and downlink transmission is burst TDM at 622 Mbit/s. The Ku-band system consists of six wider fixed spot beams, each supporting 10 TDMA carriers at 5 Mbit/s for uplink and three TDM carriers at 50 Mbit/s. In addition, a high-speed ISL is included in the system. The payload configuration is depicted in Figure 4-22.

The Ka-band payload accommodates 622 Mbit/s transmission from up to five fixed beam locations. Transmission is burst TDM in which an earth station forms bursts according to destination beams (bursts of up to 6 Ka-band beams and one Ku-band burst). These bursts are routed to either Ka-band beams in SS-TDM or Ku-band beams through an SS-TDM switch and a fast packet switch Earth stations in other Ka-band fixed beam locations and those in hopping beam areas access to the satellite in TDMA via 100-Mbit/s hopping beams. The Ku-band system is similar to Architecture 2, but provides three 50-Mbit/s downlink carriers per beam for additional capacity. The system capacity is 3.7 Gbit/s for Ka-band and 300 Mbit/s and 900 Mbit/s for Ku-band uplink and downlink, respectively. The throughput of a fast packet switch is 1.5 Gbit/s.

Earth station characteristics are similar to those described in Architectures 1 and 2: the Ka-band earth stations require an antenna size of 2.4 m to 5 m, and the Ku-band earth stations an antenna size of 1.2 m with a 10-watt SSPA.

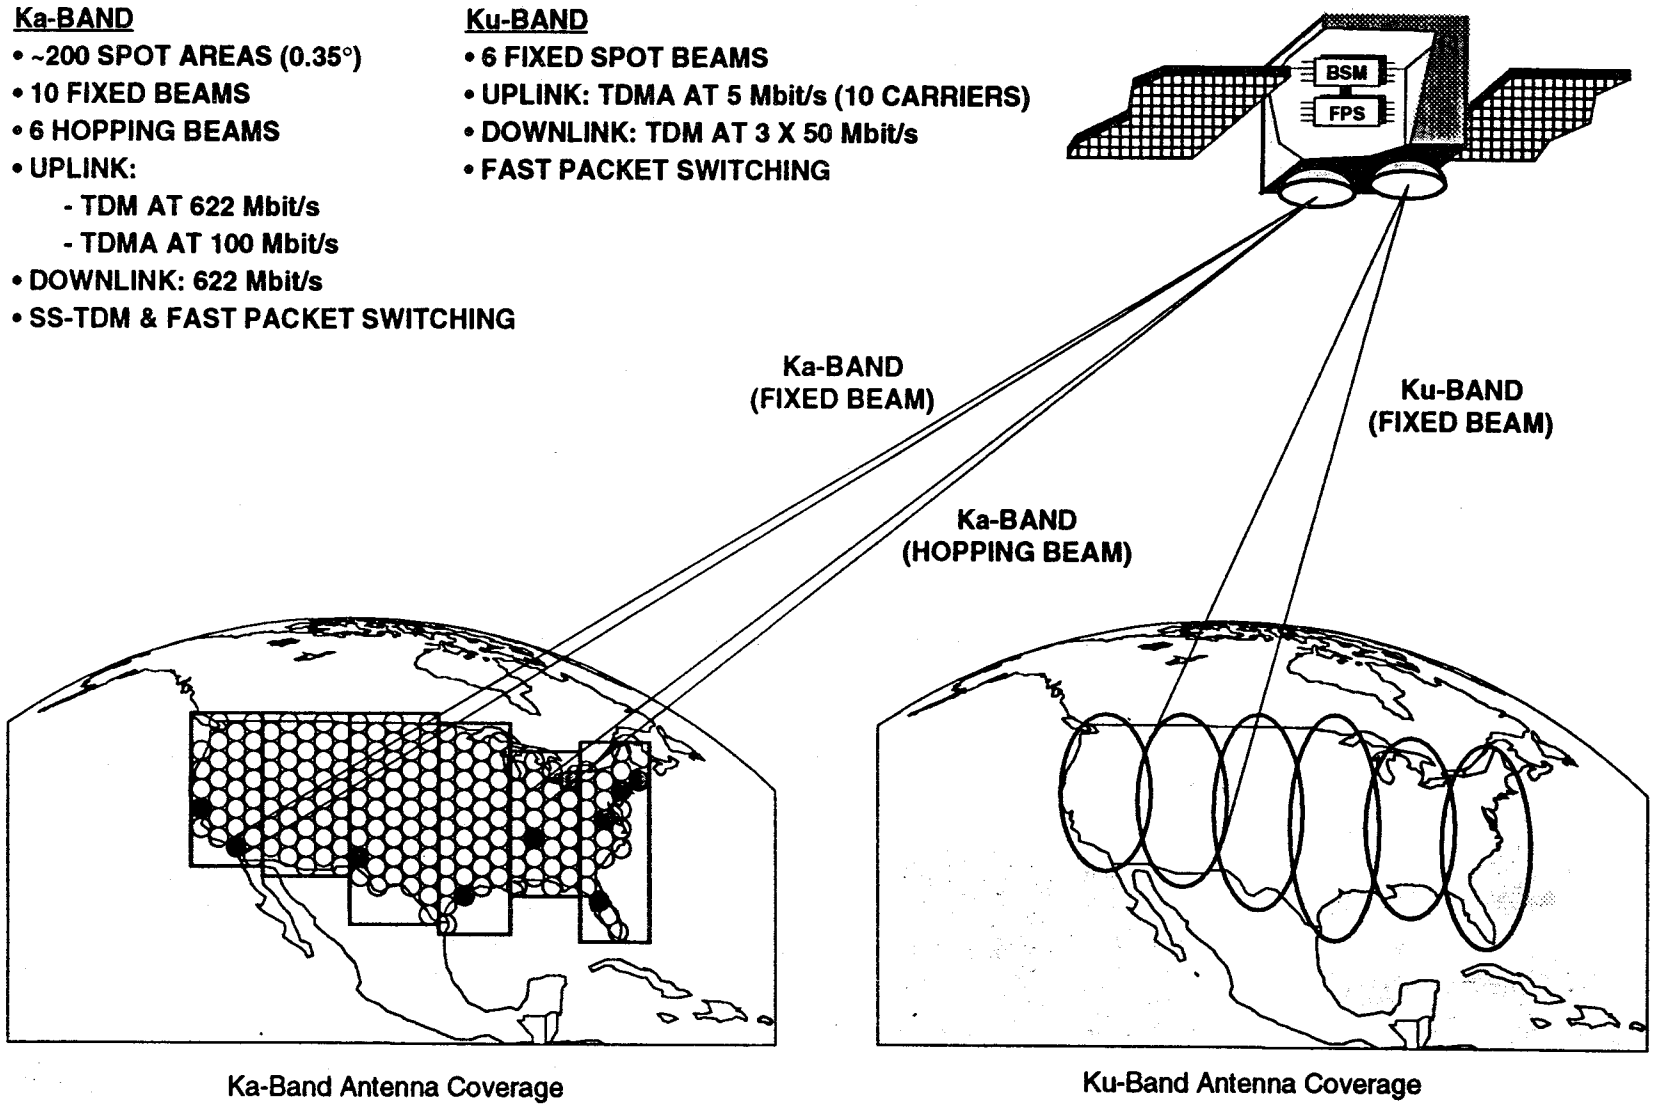

Figure 4-21. Antenna Beam Coverage for NASA Science/Networking System

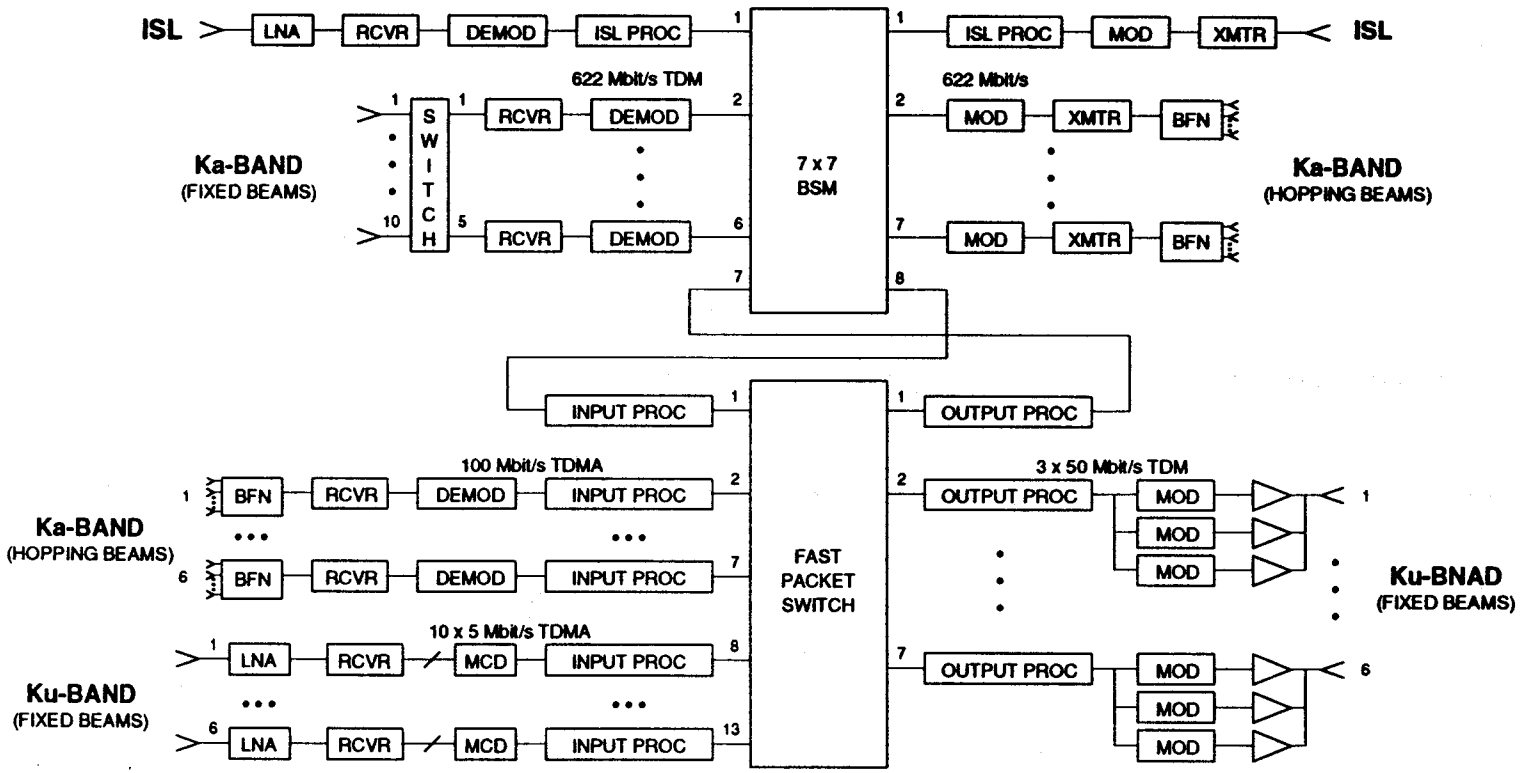

Figure 4-22. Payload Configuration for NASA Science/Networking Applications

4.6 References

[4-1] H. Gaggioni and D. Gall, "Digital Video Transmission and Coding for the Broadband ISDN," IEEE Transactions on Consumer Electronics, Vol. 34, No. 1, February 1988, pp. 16-35.

[4-2] K. Price and R. Jorasch, "Data Distribution Satellite System Architecture Concept," 1990 AIAA Communication Systems Conference, pp. 864-877.

[4-3] CCSDS 701.0-B-1, "CCSDS Advanced Orbiting Systems, Networks and Data Link," Consultative Committee for Space Data Systems, Blue Book, Issue 1, October 1989.

[4-4] R. Moat, "ACTS Baseband Processing," Conference Record, Globecom '86, Houston, Texas, December 1-4, 1986.

[4-5] G. B. Alaria, et al, "On-Board Processor for a TST/SS-TDMA Telecommunications System," ESA Journal, Vol. 9, 1985.

[4-6] On-Board Processing Satellite Network Architecture and Control Study, Final Report, NASA Contract NAS3-24886, COMSAT Laboratories, June 1987. Also, refer to: S. J. Campanella, B. A. Pontano, and H. Chalmers, "Future Switching Satellite," AIAA 12th International Communication Satellite Systems Conference, Virginia, pp. 264-273, March 13-17, 1988.

[4-7] Private communications with ESA personnel.

[4-8] T. Inukai, "An Efficient Ss/TDMA Time Slot Assignment Algorithm," IEEE Trans. Communications, Vol. COM-27, No. 10, October 1979, pp. 1449-1455.

Section 5

System Design Issues

In designing efficient and reliable satellite networks for B-ISDN services, one must address a number of design issues specifically associated with B-ISDN requirements and employ proper techniques for utilizing satellite and ground resources most effectively. This section addresses several system design issues along with alternative approaches. The system design considerations presented in the following are ATM cell concentration, ATM cell packetization alternatives, bit error rate and quality of service, ATM cell header protection, and impact of satellite delay. Some of these issues are not only associated with B-ISDN but also applicable to on-board fast packet switching for non-B-ISDN traffic and satellite data communications in general.

5.1 ATM Cell Concentration

Earth stations in the satellite B-ISDN network will be interfaced with different user network interfaces (UNIs) and network node interfaces (NNIs). To increase the transmission efficiency and share the satellite transmission link and resources among different users, the earth station performs a statistical multiplexing function. The immediate result of statistical multiplexing operation is output rate reduction; hence, improvement in transmission efficiency is expected. The buffer within the statistical multiplexer has the effect of smoothing out bursty traffic and reducing a peak-to-mean ratio. Intuitively, the higher the output transmission rate, the worse the transmission efficiency and the better the transmission delay. The larger the buffer size, the better the transmission efficiency, but the worse end-to-end delay. Therefore, to optimize delay and efficiency, a trade-off exists between the buffer size and the output transmission rate.

In this section, the impact of packet compression and decompression on satellite transmission will be analyzed. Specifically, the advantages of extracting idle cells and terminating the SDH overhead will be studied. Analytical queueing models will be derived for the statistical multiplexer in the next section to study the performance of the satellite B-ISDN.

5.1.1 ATM Interface Termination

5.1.1.1 ATM Idle Cells

In ATM, idle cells are inserted/extracted by the physical layer to adapt the rate of valid cells to the payload capacity of the transmission system. The ATM layered structure is shown in Figure 5-1 for reference. The ATM idle cell header structure is shown in Figure 5-2. The rate adaptation is termed cell rate decoupling. The function of idle cells is mainly synchronization. Cell synchronization (or cell delineation) is to recover the cell boundary based on the header error control (HEC) self-delineating algorithm. According to CCITT Recommendation I.321, the cell delineation method uses the correlation between the 4-byte header and the 1-byte HEC. Since traffic is bursty, without using idle cells, synchronization of the data stream cannot be maintained. If synchronization is not maintained, cell delineation must be performed frequently, which results in the loss of information and increase of end-to-end delay. In the case of protection switching (which will introduce a gap in the data stream), idle cells can be inserted and extracted from the data stream to maintain cell synchronization. This idle cell insertion/extraction concept has also been used in the Signaling System No. 7 (SS7) signaling link protocol where fill-in status units (FISUs) are sent to maintain signal unit alignment.

A block diagram of an ATM line card is shown in Figure 5-3 to illustrate the concept of cell rate decoupling. The ATM line card consists of one transmitter and one receiver. The cell rate decouplers for the transmitter and receiver have been enlarged in Figure 5-4.

At the earth station, it is not necessary to transmit the idle cells as long as the synchronization of the transmitted cells can be maintained on board the satellite. This synchronization issue is tightly coupled with the packet format chosen for satellite transmission. Idle cell detection and elimination is necessary at the earth station, since the earth station performs statistical multiplexing of several incoming lines. The hardware to extract the idle cells is very simple and can be implemented with a simple comparator (see Figure 5-3). As a result, a required satellite transmission capacity is less than the total capacity of all its terrestrial interfaces, and hence, the satellite transmission link can be utilized more efficiently.

5.1.1.2 Buffer Smoothing

Without providing buffers at the earth station, the satellite transmission capacity must be almost the same as the sum of the peak rates of the incoming lines to make the cell loss ratio (CLR) meet the QOS requirement. In this sense, statistical multiplexing becomes deterministic multiplexing. However, this non-buffer scheme introduces no queueing delay. If the earth station provides an infinite buffer size, then the satellite transmission capacity can be the same as the average rate of the incoming lines; however, the delay is also infinity. Evidently, a trade-off exists among four factors in

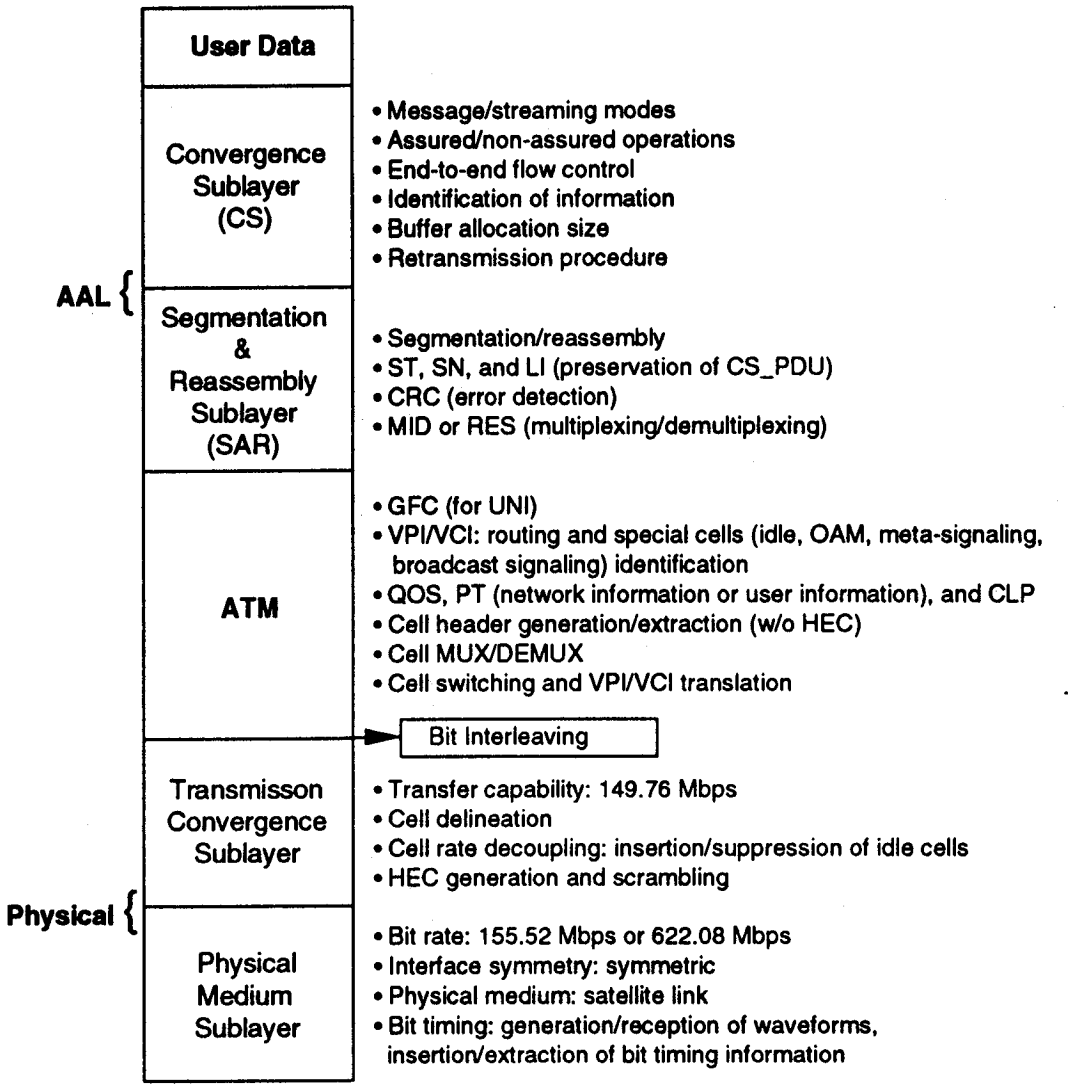

Figure 5-1. Function Requirements for Different Layers

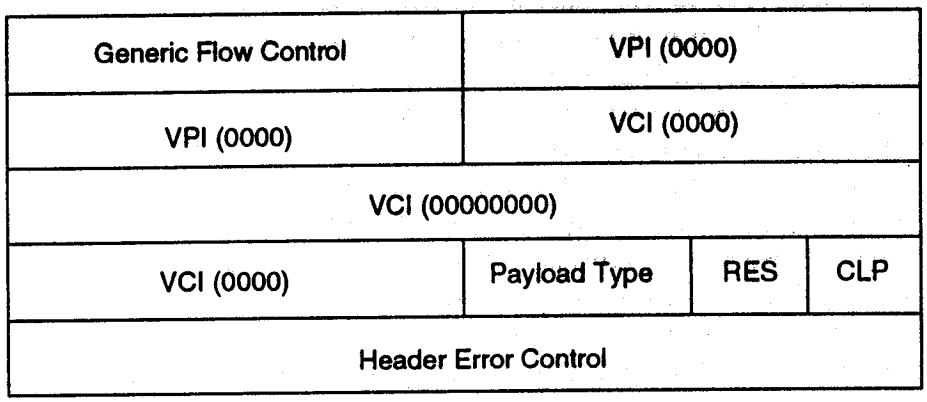

Legend:
VPI: Virtual Path Identifier
VCI: Virtual Channel Identifier
RES: Reserved Field
CLP: Cell Loss Priority

Figure 5-2. ATM Idle Cell Format for UNI and NNI

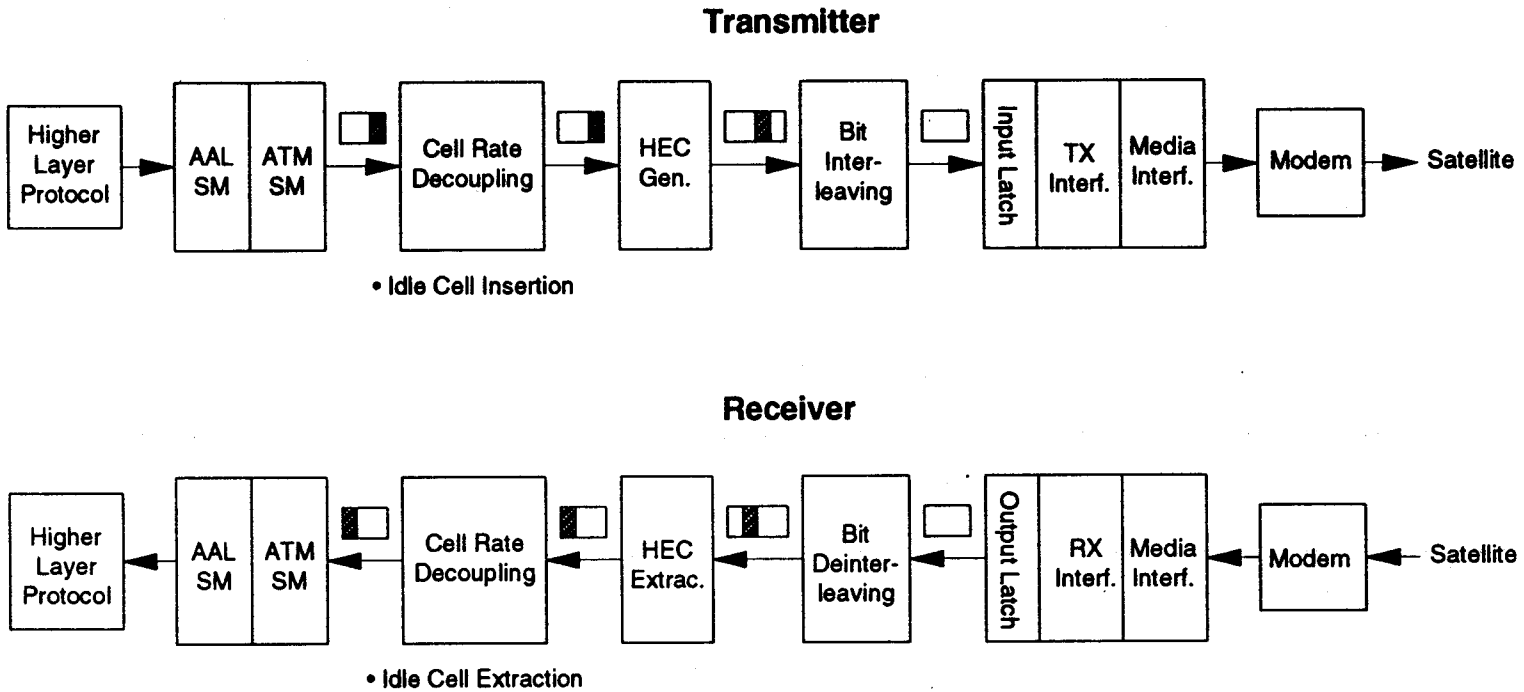

ATM SM: ATM State Machine
AAL SM: ATM Adaptation Layer State Machine
HEC: Header Error Control

Figure 5-3. ATM Line Card with Bit Interleaving

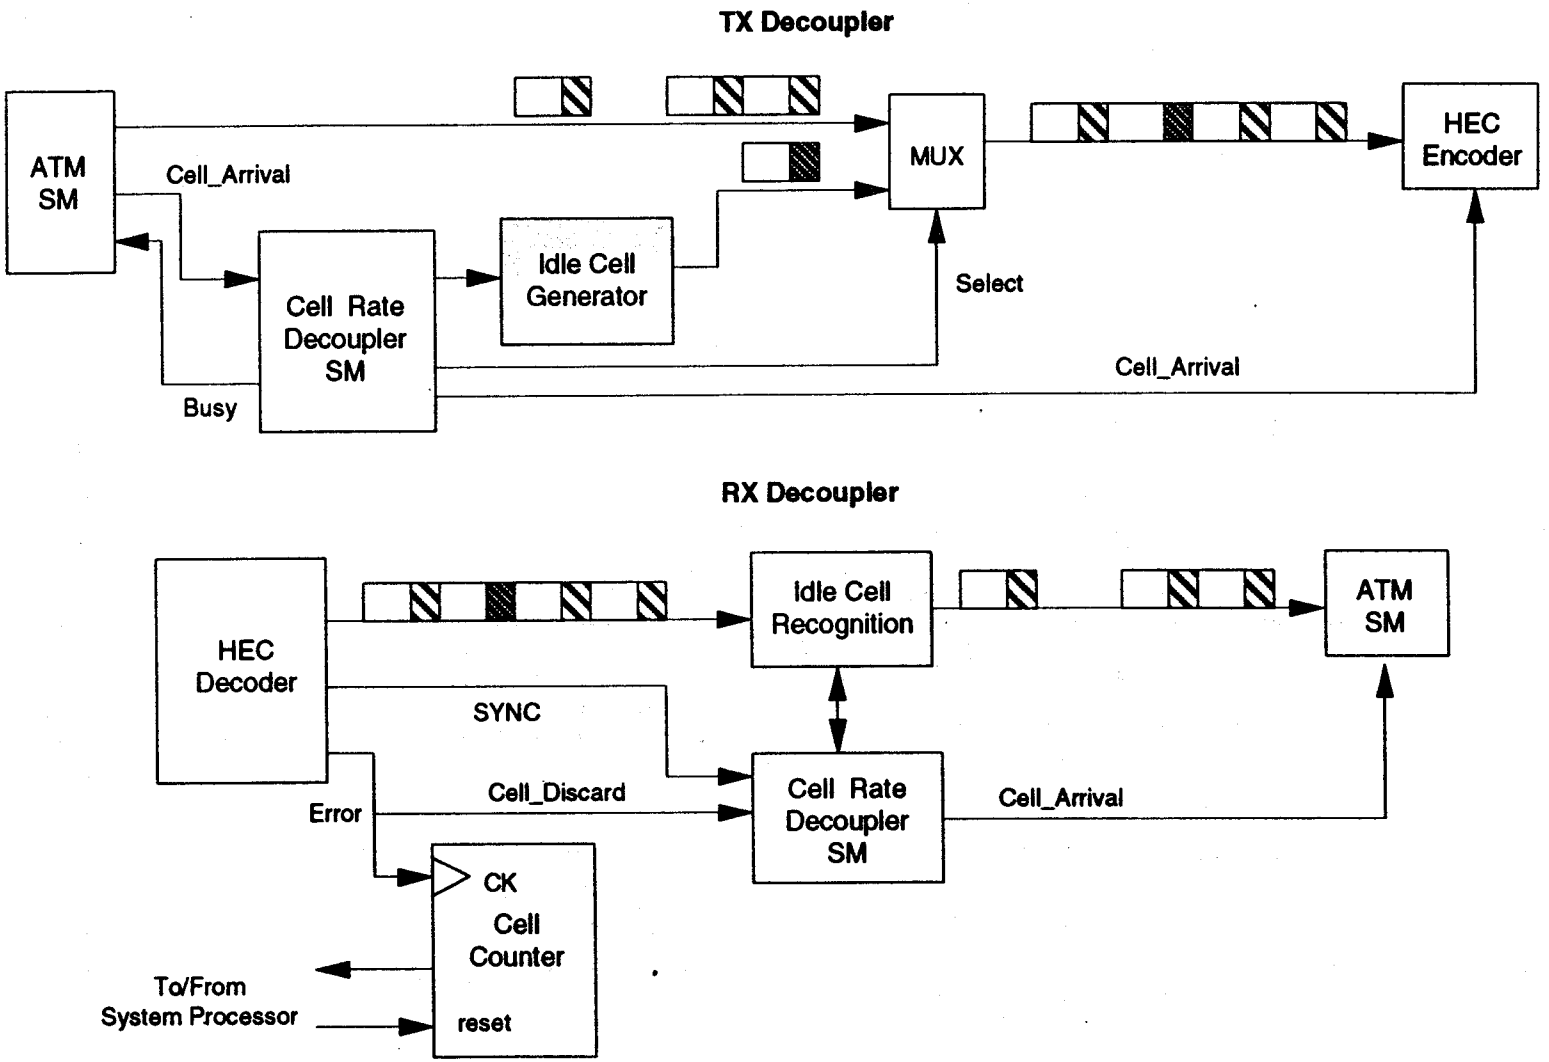

Figure 5-4. Cell Rate TX and RX Decouplers

the design of a statistical multiplexer: transmission capacity (or efficiency), buffer size, CLR, and delay. The buffer at the earth station will also help the satellite B-ISDN absorb the bursty traffic from different sources. Hence, the peak-to-mean ratio is decreased.

The size of buffers and their sharing policies among different priorities of cells must be carefully controlled to prevent such occurrences as buffer overflow, unfair resource allocation, network congestion, and other forms of performance degradation. Two buffer allocation control strategies are analyzed in the following. The first one is a complete sharing (CS) scheme. CS means that a common buffer at the output line is completely shared by all incoming lines. The second scheme is a complete partitioning (CP) scheme. In this scheme, buffers are partitioned into several disjoint areas based on the priority of the cells, and a cell may enter the multiplexer only if the area associated with the priority is not filled. A disadvantage of the CS scheme is that for highly asymmetric input rates, CS tends to heavily favor high input rates. This disadvantage may turn out to be an advantage for voice and image traffic, since the burstiness of B-ISDN traffic will affect the QOS for delay and loss sensitive traffic. The disadvantage of the CP policy is that buffers allocated to an inactive priority process are wasted, but sometimes this is a compromise to guarantee the QOS. Four queueing models will be analyzed for different buffer allocation strategies and queueing disciplines.

There are several approaches to analyzing a queueing system. The first one is exact analyses via numerical methods. For an ATM network, it is difficult to model the system without any simplified or modified assumptions. When a modification of the model is made, it is necessary to validate the model by using simulation techniques. The second approach employs a computer modeling to duplicate a physical system, which is essentially simulation. In this section, several queueing models for the statistical multiplexer will be analyzed. Queueing equations will be derived and numerical programs will be written to solve the equation. After these, extensive simulation will be performed to verify the results. Simulation will also be used to obtain the performance of the statistical multiplexer for several different traffic scenarios, which cannot be easily modeled using queueing equations.

Buffer smoothing has the effect of avoiding network congestion by reducing the burstiness of the input traffic. This effect will relieve congestion on board the satellite. In this context, it is expected that the traffic from earth stations to the satellite has been modulated and smoothed. Hence, the on-board buffer requirement to maintain the QOS will greatly reduced. The earth station queueing models will be extended to the on-board switch. In this sense, a star network (a satellite network) queueing model will be developed to design a B-ISDN satellite network to accommodate B-ISDN traffic and services.

5.1.2 SDH Overhead Termination

Two synchronous digital hierarchy (SDH) rates have been defined for ATM in CCITT G.707-G.709: 155.52 Mbit/s and 622.08 Mbit/s. The SDH frame structure is depicted in

Figure 5-5. The ATM cells will first be mapped into a C-4 container. Byte boundaries of the ATM cells will be aligned with the byte boundaries of the C-4 container. Then the path overhead (POH) is multiplexed with the C-4 container to form the VC-4 virtual container. Since the C-4 capacity (2340 bytes = 44 cells + 8 bytes) is not an integer multiple of the cell length, a cell may cross a C-4 boundary. The POH provides communication and maintenance service between the points of assembly of a container and disassembly. The fields in POH are depicted in Figure 5-5 and the functions of the path overhead are summarized in Table 5-1. The bits 3 through 8 of the multiframe indicator H4 are for the cell offset indicator. The H-4 offset will be set at the sending side to indicate the next occurrence of a cell boundary, i.e., the H-4 indicates the offset in bytes from itself to the first cell boundary in the C-4 payload. This H-4 offset can be used to supplement the initial cell delineation. The confirmation of the cell boundary is made using the HEC cell delineation algorithm. In this way, confirmation is done within a few cell periods instead of a few STM-1 frame periods.

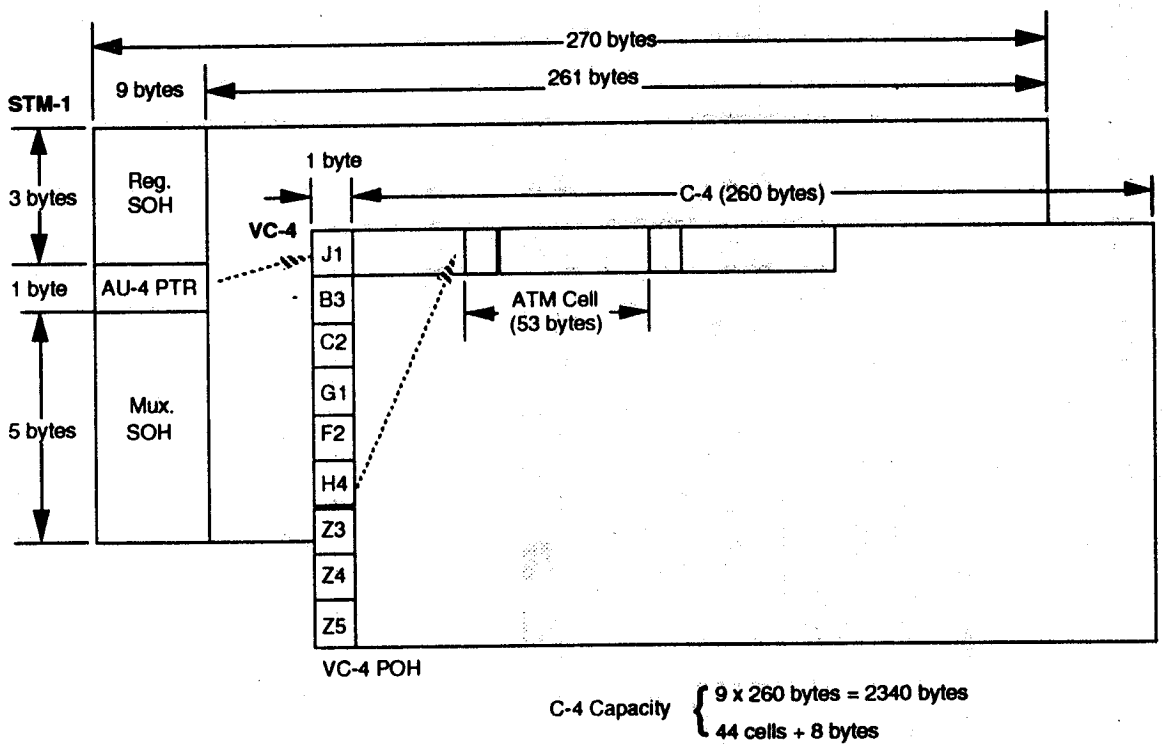

Figure 5-5. SDH Frame Structure

The VC-4 is combined with the administrative unit pointer to form the administrative unit (AU-4). The location of the first byte of the VC-4 POH is indicated by the AU-4 pointer. The AU-4 pointer is in a fixed position in the STM-1 frame, but the payload (VC-4) is not required to be synchronized with the STM-1 frame. Hence, frequency adjustment between the payload and the frame can be done using the AU-4 pointer. The SDH frame is formed by multiplexing the SOH and the VC-4. The fields in the SOH are depicted in Figure 5-6, and the functions of the SOH is summarized in Table 5-2.

Table 5-1. Path Overhead Functions

ELEMENT	ABBREVIATION	DESCRIPTION
VC-3/VC-4 Path Trace	J1	Indicates the beginning of the virtual container
Path BIP-8	B3	This byte contains the even parity of all bytes of the previous VC
Signal Label	C2	Indicates the composition of the VC3/VC4
Path Status	G1	Contains the path termination status and performance
Path User Channel	F2	Contains user communication between path elements
Multiframe Indicator	H4	Generalized multiframe indicator with payload specific application
Spare	Z3 -Z5	Reserved for future use

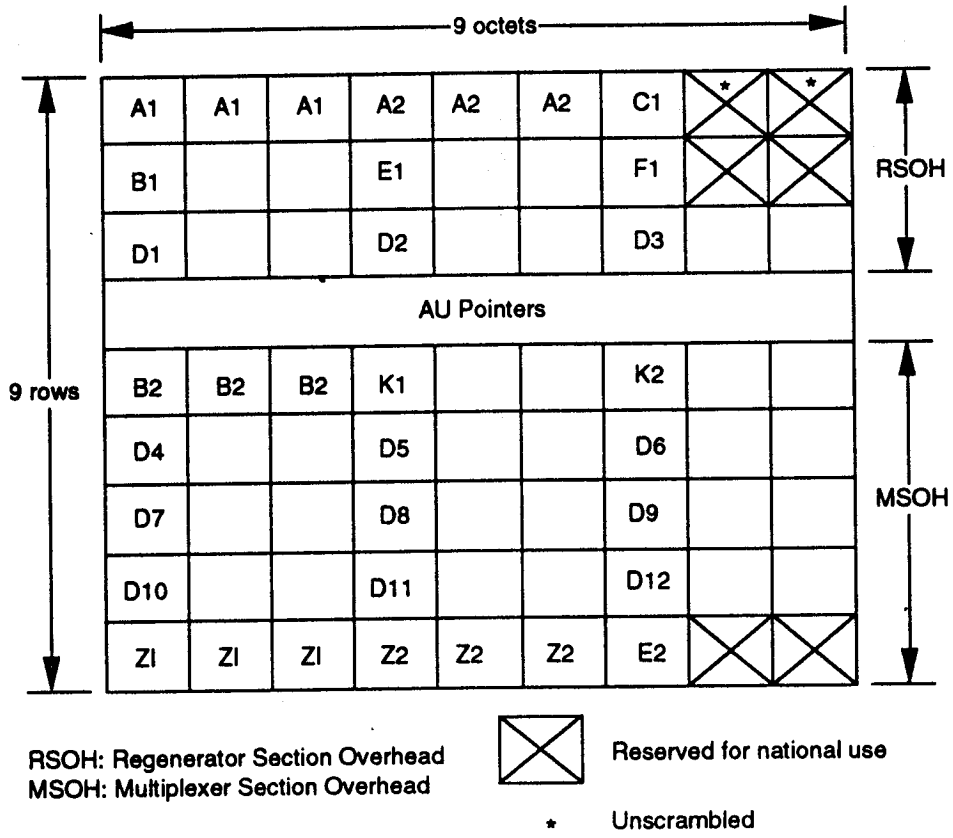

Figure 5-6. SDH Section Overhead

Table 5-2. Section Overhead Functions

ELEMENT	ABBREVIATION	DESCRIPTION
Framing	A1, A2	Provides framing
STM identifier	C1	Contains value of the multi-column, interleave depth coordinate
Data Com. Channel(DCC)	D1-D12	Provides a 192 and a 576 Kbit/s data communication channel
Orderwire	E1, E2	Provides voice orderwires
User Channel	F1	Reserved for user purposes
BIP-8	B1	Provides regenerator section error monitoring
BIP-Nx24	B2	Provides multiplex section error monitoring
APS Channel	K1, K2	Automatic protection switching channel
Spare	Z1, Z2	Spare

By following the procedures mentioned above, the SDH overhead can be terminated and the ATM cells can be extracted from the SDH frame. After extraction of ATM cells, elimination of idle cells becomes feasible. Hence the satellite link capacity can be used more efficiently.

The advantages of terminating the SDH and ATM at the earth stations are summarized as follows.

a. Better satellite resource usage: If the ATM interface is terminated at the earth station, idle cells can be extracted from the ATM stream, and the satellite link can be used more efficiently. Similarly, if the SDH interface is terminated at the earth station, extracting cells out of the frame and discarding empty cells is possible. A low-bit-rate frame such as the SONET STS-1 (51.84 Mbit/s) can then be used to transmit the remaining information frames.

b. Generation and processing of maintenance functions: After the earth station processes the maintenance function in the SDH overhead or the OAM cells, the earth station can send the maintenance signals to the local exchange to report failure detection.

c. Delay compensation for the data communication channels (DCCs) in the SDH overhead

d. Improvement of ATM cell loss rate performance by using bit interleaving: Since the ATM cell header provides a one-bit correction capability, provision of bit interleaving at the earth station can reduce the effects of bursty errors.

5.2 ATM Cell Packetization Issues and Alternatives

The basic difference between circuit switching and fast packet switching is that a circuit is routed based on the slot position and a packet is routed based on the label (tag) in the packet header. Since the routing information is contained in the packet header, the packet can be routed and switched within the network without the need of a centralized controller, as opposed to circuit switching. This fast packet switching technique is applied to the satellite network to accommodate B-ISDN traffic; hence, the on-board baseband processor uses the destination-directed packet routing technique, which is different from a circuit switching baseband processor such as used in ACTS. For the circuit switching baseband processor, the routing path within the processor has to be set up in advance, which involves interaction among the network control center, the transmit and receive earth stations, and the baseband processor. The destination-directed on-board packet routing eliminates the need of path setup of calls on the on-board switch. The advantages are the packet can be sent out immediately from the transmit earth station after the routing tag is appended in the header, and the on-board switching hardware is simplified and the processing requirement of the network control center is lessened.

For ATM cells, 24 bits virtual path identifier and virtual channel identifier (VPI/VCI) at user network interface (UNI) and 28 bits VPI/VCI at the network node interface (NNI) are available for routing information. Issues related to the VPI/VCI such as the use of virtual path to simplify the processing requirements and the call setup delay, and the elimination of VPI/VCI translation at the on-board switch are discussed first. Then the satellite virtual packet concept is introduced for unified routing, control, and management in satellite B-ISDN.

5.2.1 VPI/VCI Processing

At the earth stations, for routing, signaling, and synchronization purposes, either VPI processing or VPI/VCI processing must be performed depending on the application. Before discussing the difference between VPI processing and VPI/VCI processing, the concept of virtual paths (VPs) is introduced. VPs are a bundle of virtual channels that have a common destination. A VP allows many virtual channels to be processed, transported, and managed as a group. In the satellite B-ISDN network, a VP provides a direct logical link between any two earth stations. The decision of using VPs or VPs/VCs lies in the performance requirements of the satellite network and the complexity difference between VPI processing and VPI/VCI processing. Basically, the difference is the following.

a. Addressing space flexibility: The total number of bits in the VPI/VCI is 24 bits at the UNI and 28 bits at the NNI. The VPI only has 8 bits at the UNI and 12 bits at the NNI. Therefore, the number of connections which can be set up using the VPI/VCI is larger than that using the VPI only.

b. Bandwidth allocation efficiency: If only VP processing is supported, then each time a demand is made for more virtual path identifiers, at least 2^8 (or 2^{12}) virtual channel connections will be allocated. This is not efficient if the increase (or decrease) of number of connections is not large. But this VP concept is well suited for video and telephone trunking applications.

c. Node processing complexity: By definition, VP processing only processes VPI; hence, the hardware complexity of node processing is low and equipments to implement VP processing are simpler.

d. Call setup delay: For VPs, a predefined route is provided for each VP in the satellite network. Therefore, at the transit node, the routing table for the VPI does not need to be rewritten at call setup time using signaling. Whenever one cell coming into the transit node, the incoming VPI will be compared with a routing table and an outgoing link can be found. However, at the end node, the VPI still has to be rewritten at every cell setup. The call setup delay will be minimum if the earth station and the on-board switch act as transit nodes.

e. Satellite network management: In a VP, all the virtual channels with same origination and destination will be grouped into a single bundle for routing, control, and management purposes. Therefore, it is simpler to manage a VP instead of many VCs.

VPI can have either a local or a global significance within the satellite network based on the VPI uniqueness of the NNIs interfaced with the station. If each NNI interfaced with the earth station has a unique VPI within the satellite network, then VPIs have a global significance; in other words, the VPIs are shared by all the NNIs interfaced with the satellite network. If VPIs can be reused at different NNIs interfaced with different earth stations, VPIs only have a local significance. The advantages and disadvantages of VPI with a local significance and a global significance are discussed below.

a. Topology flexibility: VPI with a local significance provides more flexibility in adapting the network topology for network expansion (such as node addition and link addition) and node failure. In contrast, VPI with a global significance has less flexibility. For example, once a node is failed, it is hard to reuse the VPI space reserved for that node.

b. Node processing and delay: Since VPIs with a global significance have a unique VPI for each NNI, the transit node only performs routing while VPI retranslation is not necessary. The result is that node processing cost is small and the call setup delay is minimal.

c. Addressing space: The addressing space for VPIs with a global significance is shared by the NNIs within the satellite network. This may not be feasible for a large network since the number of NNIs is large.

Two scenarios are provided for the VPI with a global significance and the VPI with a local significance in Figures 5-7 and 5-8, respectively. In the first scenario where the VPI has a global significance, the earth station is only responsible for cell routing and multiplexing and the VPI has an end-to-end significance within the satellite network. Since all the VPIs have to be distinct inside the satellite network, the number of NNIs, which the satellite network is interfaced with, is limited by the VPI addressing space. In the second scenario, since VPI only has a local significance, VPI retranslation is performed at the earth stations and at the space segment. Although the VPI addressing space can be reused within the satellite network, the extra delay and more complex logic incurred by the VPI retranslation at each node are the main disadvantages.

5.2.2 Satellite Virtual Packet Format Alternatives

From the discussion in Section 2, the VPI with a local significance is more advantageous than the VPI with a global significance in terms of flexibility and the addressing space. However, the VPI with a local significance requires the VPI to be translated at every VP terminator (such as a switch or a multiplexer). Since the space segment contains a fast packet switch, VPI retranslation on-board will result in a larger delay, higher processing cost, and more memory requirements. To utilize the strength of the VPI with a local significance and to avoid VPI retranslation on-board, the satellite virtual packet (SVP) concept is introduced. SVPs are created by appending a header (with routing tag) to one cell or a group of cells which are destined to the same down link beam for unified routing, control, and management purposes. The header of the SVP is termed as the satellite virtual label (SVL). The necessary components in the SVL will be introduced in the next subsection. Three possible formats for grouping the cells are shown in Figure 5-9. First, cells are grouped into fixed size packets at the earth station. Since the packet length is fixed, the packet boundary can be determined as long as the packets are in synchronization. Packet synchronization is the first step for on-board processing/routing. Two cases can be identified based on that whether the packet boundary is aligned with the burst (or frame) boundary or not (see Figure 5-10). If the packet boundary is in alignment with the burst (or frame) boundary, the packet can be extracted from the packet stream as long as the burst unique word (or the frame synchronization marker) has been detected. If the packet is not aligned with the burst (or frame) boundary, after the burst unique word detection, the on-board processor has to perform packet synchronization, which involves a temporary storage of the partial SVP. In general, the uplink TDMA burst synchronization or TDM frame synchronization are maintained using a fixed unique word or a synchronization marker. To have a low probability of false synchronization, the marker should not be too short. Hence, this marker may cause a high overhead if the TDMA burst or TDM frame is short. If error control is applied to the synchronization marker, the overhead will be higher. Secondly, cells are grouped into variable size blocks. In this case, cell length indication is necessary to be included in the SVL or a begin flag and an end flag are necessary to be appended at the beginning and the end of the packet (see Figure 5-11). Thirdly, SVP only consists of one cell. In this case, the SVL seems unnecessary; however, the routing tag within the SVL can enlarge the VPI/VCI addressing space and,

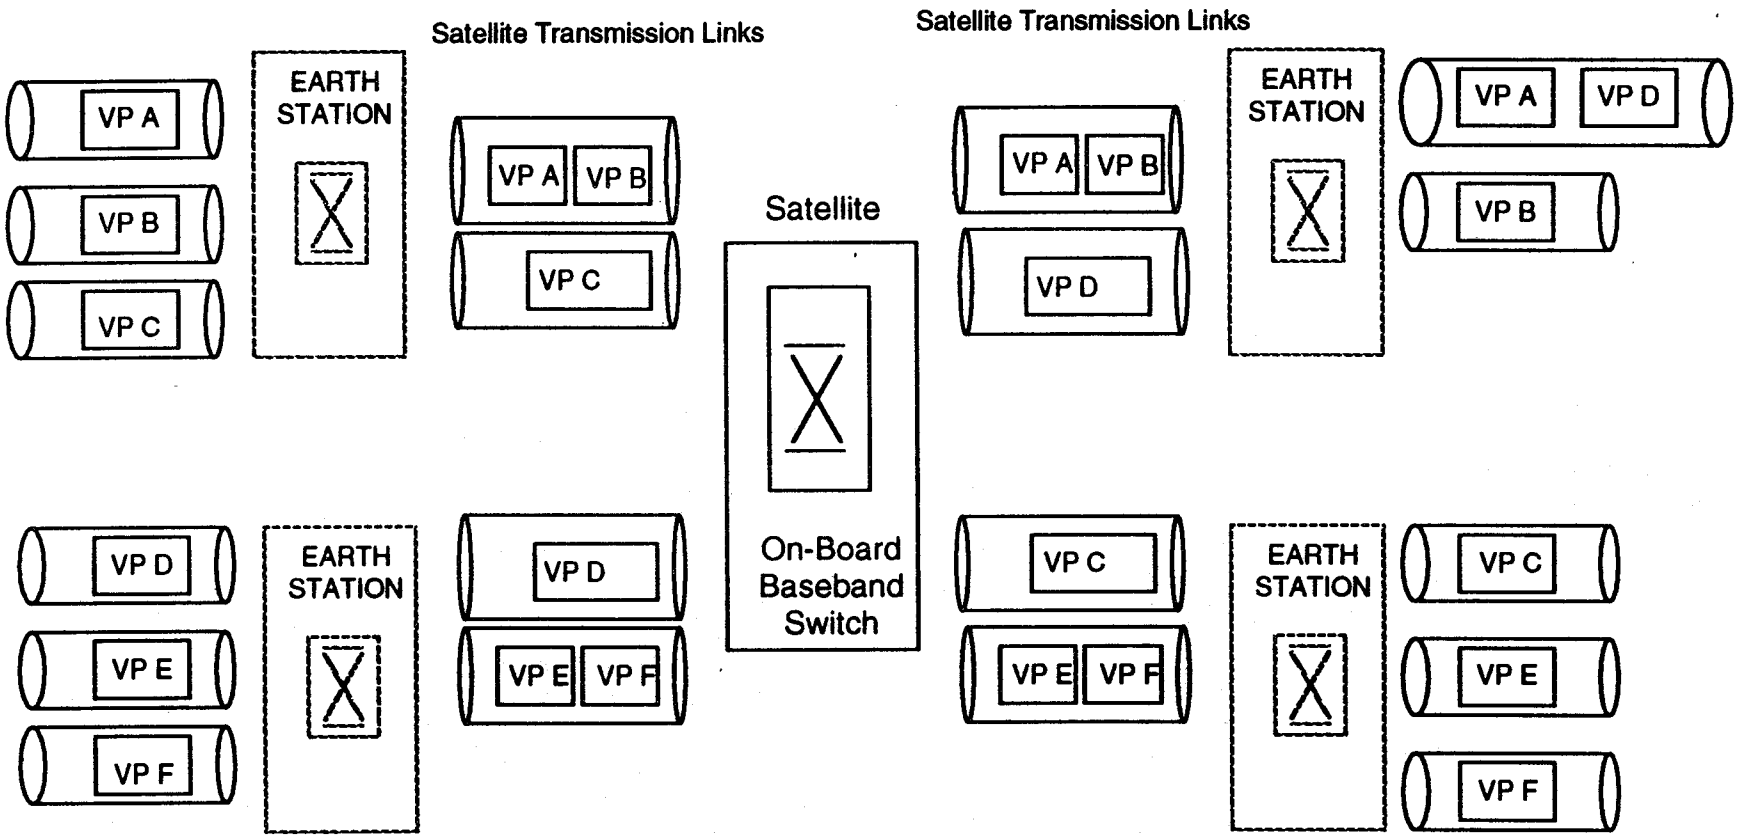

Figure 5-7. VPI with Global Significance

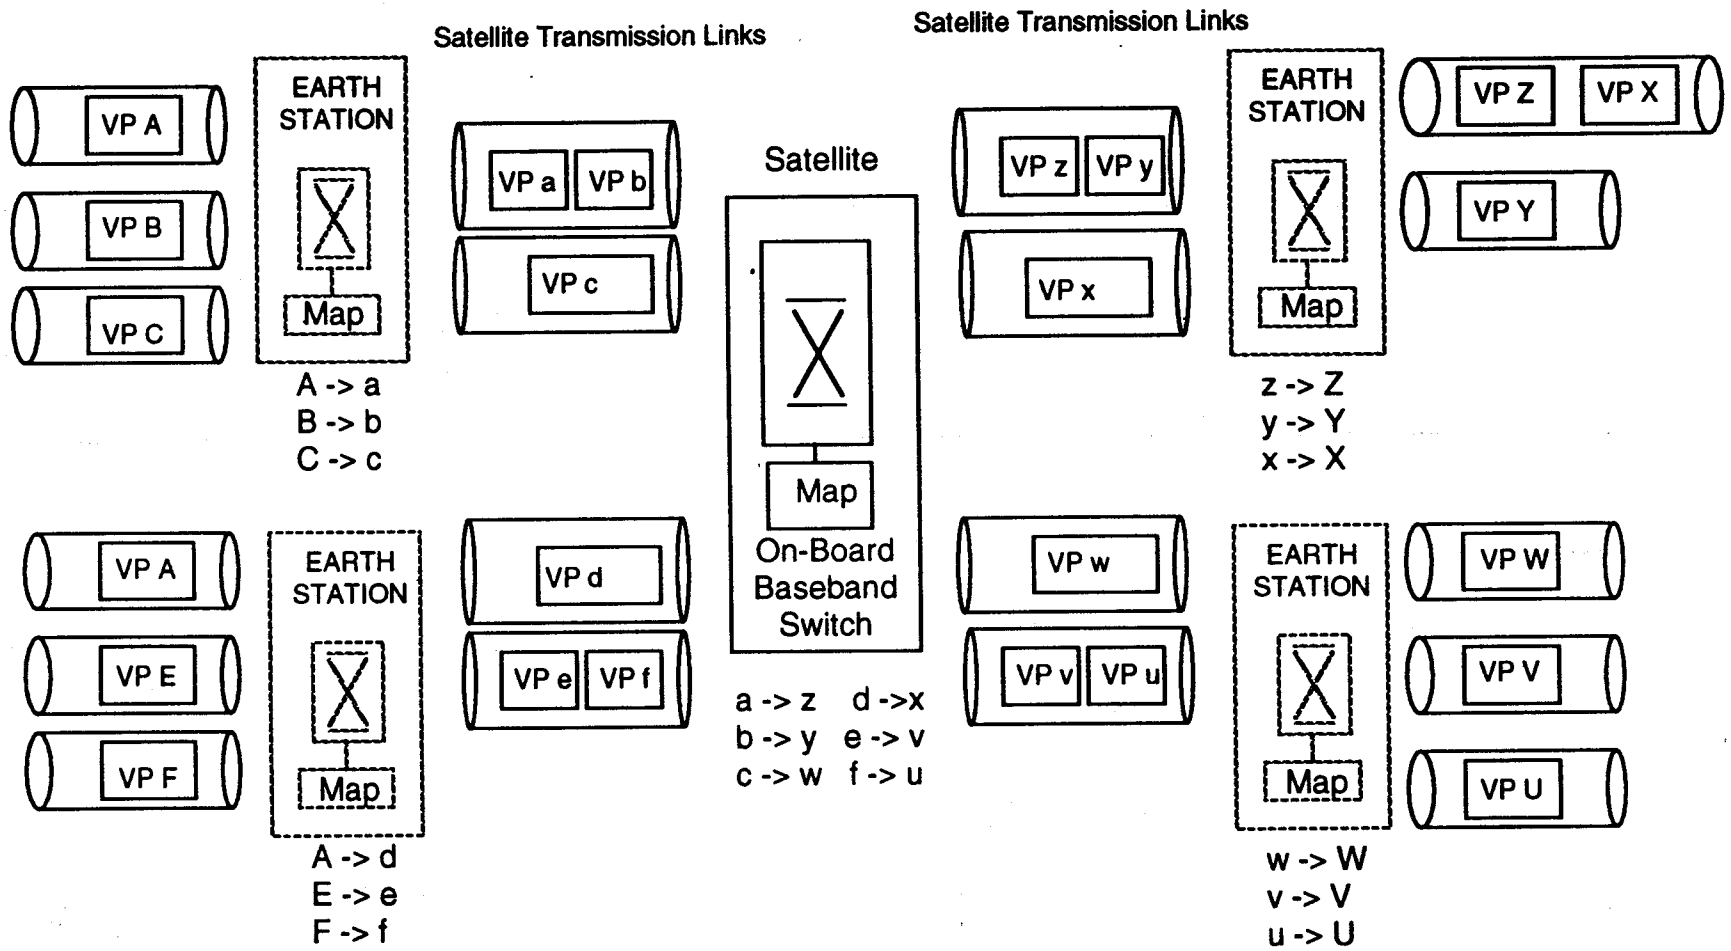

Figure 5-8. VPI with Local Significance

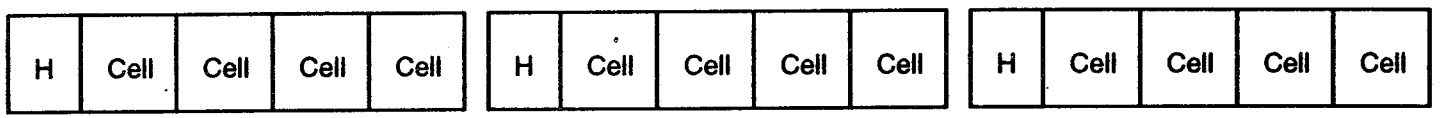

Packet with a Fixed Size

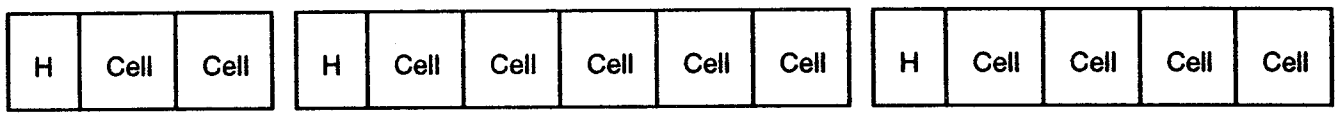

Packet with a Variable Size

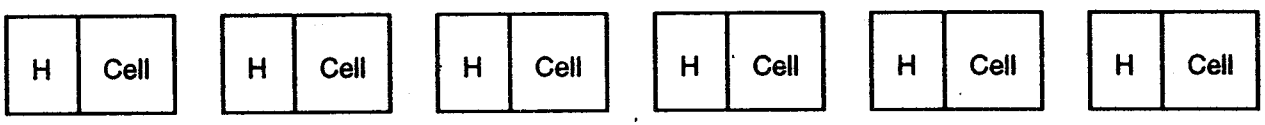

Packet with Individual Cell

H: Packet Header

Figure 5-9. Satellite Virtual Packet Format Alternatives

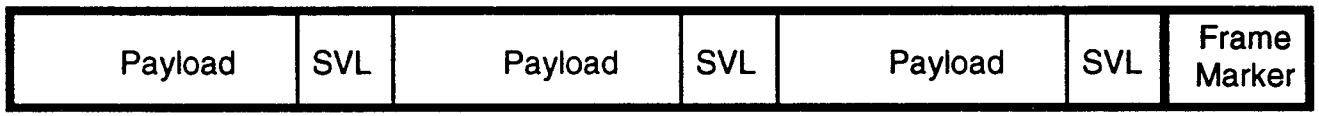

Packet Boundary is Aligned with the Frame Boundary

Packet Boundary is not Aligned with the Frame Boundary

Figure 5-10. Two Possible Frame Formats

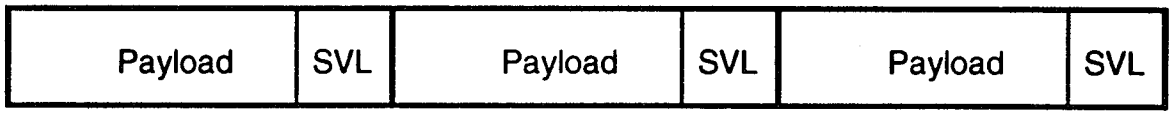

Fixed Packet Size Format

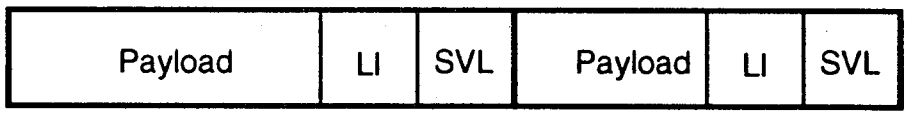

Variable Packet Size Format

Legend:

BEGIN: begin flag
END: end flag
LI: length indication

Figure 5-11. Differences Between Fixed Size Packet and Variable Size Packet

therefore, avoid the on-board retranslation of the VPI/VCI. This approach has the simplest implementation but it also has the largest overhead. Another disadvantage of the this one-cell packet is that it is hard to perform bit interleaving using only one cell because the interleaving depth is not large enough to combat the burst error. This issue is discussed in the "Protection of ATM Cell Headers from Burst Errors" section.

One advantage of using the SVP containing more than one cell is that the error code parameters of the SVP can be chosen freely since the error code does not have to be limited by the 53-byte cell length. The only difference between the first two approaches is that one has a fixed packet size while the other one has a variable packet size. The variable size approach is more flexible in allocating the capacity based on the traffic pattern and intensity; hence, the efficiency of the transmission link is higher. Also the packetization delay (the time required to fill a SVP with cells) of formatting SVPs at the earth station can be optimized. However, the SVL format and the packet structure are more complicated, which results in higher delay and more complex logic in processing the SVP and SVL. There are two options for the on-board processor to process the variable size SVP. The first one is the processor segments the variable size SVP into fixed size small packets so that the switch and its associated memory and clock signals can be operated in slotted mode. One variation of the first approach is that the SVP packetization process only supports several sizes, for example sizes of 2, 4, 6, and 8 cells. The on-board switch operates in slotted-mode where the slot size is 2 cells. The header of a SVP carries an indication of the size of the SVP. The on-board switch examines the indication and determines the number of transmission through the switch required for each packet. This issue is discused in the Appendix D "Traffic Simulation of Fast Packet Switched Networks". The second one is that the switch and its associated memory and clock signals are operated in mini-slotted mode to accommodate the variable size packet. The fixed size approach, although, has lower transmission link efficiency, it is advantageous in the simplicity of on-board processor and packet switch design and feasibility of slotted mode operation. Another problem with the fixed size approach is that the SVP packetization delay at the earth station may be intolerable if the SVP size is large and the incoming traffic intensity is low. This situation can be remedied by using a timer for each SVP at the earth station. If the timer expires and the SVP has not collected enough cells, the SVP will be padded with empty cells and sent out. To conform with the fixed size nature of ATM cells, a fixed size SVP is chosen in this study.

The size of the SVP is the next issue which has to be determined. Since the size of the SVL will be fixed, the more the cells are put into the SVP, the higher the transmission capacity (bandwidth) utilization will be. A larger SVP also increases the packet interarrival time, which lightens the speed requirement for on-board SVL processing. However, a larger packet will have a longer packetization delay at the earth station, and it is inefficient if not enough number of cells can be filled into the packet. A larger packet size will result in larger buffer requirement, longer end-to-end delay, and worse delay jitter. A larger packet will also increase the number of bits of the SVP payload in error. Since the buffer requirement is a main concern for the space segment and the delay determines the quality of circuit emulation service, it is envisioned that one SVP

should only consist of several cells. See Appendix D for more results about the SVP delay through the satellite network.

The QOS control will apply to the SVP so that all the cells within the packet will have the same QOS. For example, the cell loss priority (CLP) of all the cells in the same SVP should be the same. If the QOS is not considered, the most demanding cell will decide the QOS for the whole packet. By performing QOS control, the network resource can be divided among different classes of services and the optimal traffic control is performed for each class to satisfy the efficient network resource usage and the demanded quality for each service.

5.2.2.1 Satellite Virtual Label

The necessary fields of the SVL are proposed to contain the following (see Figure 5-12): the routing tag, the sending earth station address, the receiving earth station address, the QOS, the control field, and the CRC.

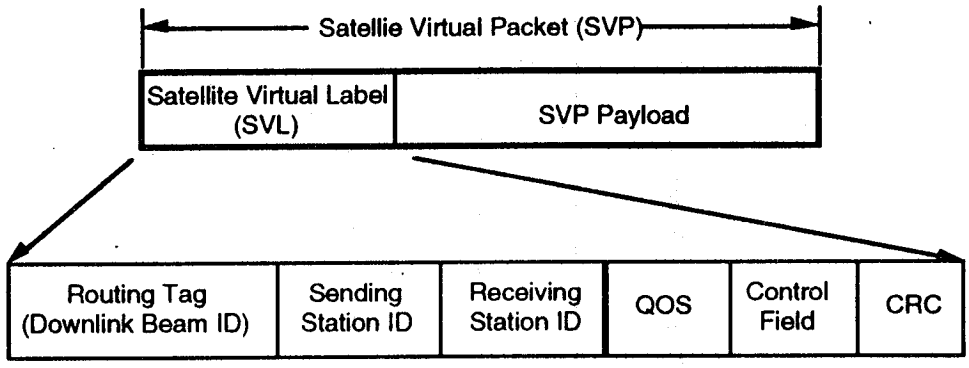

Figure 5-12. Satellite Virtual Label Format

The routing tag is to identify the down link beam and is also used for routing through the on-board switch. The structures of the on-board fast packet switch are studied in the "Fast Packet Switch Architectures" Section, and several self-routing fast packet switches are chosen for detailed design . This self-routing packet switch will route each packet based on the binary bit representation of the routing tag in the point-to-point connection case. In the terrestrial network, the appending of the routing tag is done at the input port of the switch. For the satellite network, to reduce the on-board switch complexity, appending of the routing tag may be performed at the earth station instead at the on-board switch. The size and format of the routing tag depends on the fast packet switching architectures and connectivity (point-to-point or point-to-multipoint). For point-to-point applications, the routing tag should consist of at least (Log_2 N) bits, where N is the size of a switch or the number of the output ports. For point-to-multipoint applications, the multicast routing tag may be appended at the earth station depending on the size of the routing tag. One possible routing tag format uses a series of 1's and 0's. The positions of all the 1's means all the destined output ports. Since this routing tag size is the same as the size of switch, the transmission link efficiency may be

low if the switch size is large. Hence the other alternative is to carry the original VPI, and the multicast routing tag is obtained by performing an on-board translation at the input port (or at the output port if multicast connectivity is achieved using a bus to connect all the output ports). It is emphasized that if the routing table is fixed, then the complexity of the on-board switch can be minimized. However, to be able to rewrite the routing table provides more flexibility for VPI reuse. The point-to-point and point-to-multipoint routing tags are shown in Figure 5-13.

An inherent advantage of appending the routing tag at the earth station is that fault-tolerance of the on-board switch can be easily achieved by simply changing the routing tag of the packet at the earth station. As a result, this packet can be routed through a fault-free path within the switch. All the issues related to the routing tag are examined extensively in "Fast Packet Switch Architectures" and Fast Packet Switch Designs" Sections.

The receiving earth station ID is to identify the receiving earth station within the same down link beam. Point-to-multipoint connections can be achieved easily if the receiving stations are within the same down link beam since no duplication of packets is necessary. The sending earth station ID is to identify the sending earth station so that each sending earth station can use the full addressing space of the VPI/VCI. If the sending earth station ID is not included, then two sending earth stations cannot send packets with the same VPI/VCI to the same receiving earth station; hence, the receiving earth station has to partition its VPI/VCI addressing space to all of the correspondent sending stations. The sizes of the receiving and sending station IDs are proportional to the number of earth stations in the network.

As mentioned above, the QOS field is for QOS control. All the cells within the SVP should have the same QOS. For example, all the cells within the SVP will have the same cell loss priority (CLP). Based on this QOS field, the on-board switch can perform priority control to guarantee the QOS for certain services, and also can drop the low priority cells if congestion occurs. It is suggested that 2 bits are necessary for QOS control; one is for the cell loss performance and the other is for the delay performance. The control field is to identify the SVP is an information packet or a control packet. If the SVP is a control packet, it will be routed to the on-board autonomous network controller. Also the on-board autonomous network controller will generate control SVPs, and these packets will be broadcasted to all the earth stations. One bit is enough for the control field. The last field is the CRC field which has the capability to correct and detect errors in the SVL. The selection of the CRC code is very important since an inappropriate CRC code will degrade the CLR performance. It should be mentioned that the CRC field can also be used for synchronizing the SVPs as the cell delineation algorithm performed in the ATM cell synchronization procedure.

5.2.2.2 Timing Consideration

To meet the delay requirement for the circuit emulation services and variable bit rate video services (AAL class 1 and 2 services), it may be necessary that the SVP has to be

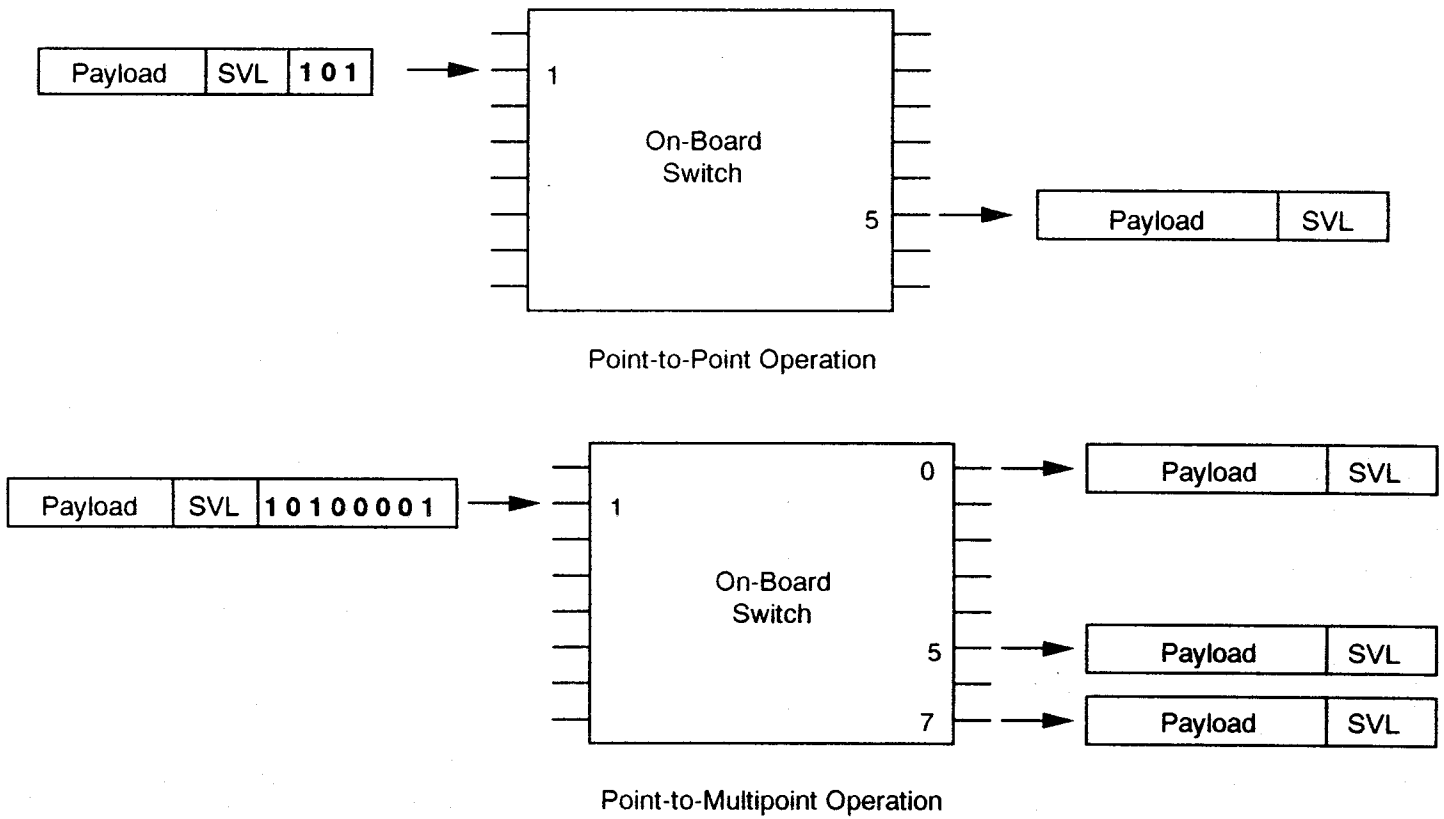

Figure 5-13. Point-to-Point and Point-to-Multipoint Routing Tags

sent out even though it is only partially filled. The empty portion of the SVP can be filled with idle cells. By inserting the idle cells, the length indicator is not necessary to be included in the SVL since idle cells can be extracted easily at the receiving earth station.

5.2.3 SDH Packetization

Synchronous digital hierarchy (SDH) and synchronous optical network (SONET) will be used to support B-ISDN traffic in Europe and Northern America, respectively. SDH (or SONET) supports both ATM connections and circuit switched connections. Several options are available to carry the SDH traffic through the satellite links. The first one is to transport the whole frame. In this case the bit rate requirement of satellite transmission will be the same as the SDH bit rate requirement (either 155.52 Mbit/s or 622.08 Mbit/s). The second option is to transmit the SDH payload (such as VC-4 container) and SDH overhead separately. In this scenario, the SDH payload can not be decomposed into cells at the receiving earth station until the OHs are received. Basically the requirements of this approach are the same as the first approach except the bit rate requirement is reduced. The third approach is to packetize the SDH payload and SDH OH into SVPs. For unified control and easy maintenance, it is proposed that the SDH will be converted into the SVP format, which is the focus of this section. Another advantage of packetizing the SDH signal is that point-to-multipoint SDH connections can be easily achieved. For the SDH information payload, a multicast fast packet switch on-board the satellite can send the SVPs to multiple destinations; hence, duplication of packets can be done at the space segment. For the SDH overhead, each path in the multipoint configuration has to be monitored and the data communication channel has to operated in multipoint mode (such as the ARQ protocol), the duplication of the SVPs for the OHs should also be done in the space segment. The multipoint tag in the SVP may be appended at the earth station to minimize the space segment complexity. The operation of point-to-multipoint fast packet switching is discussed in the "Fast Packet Switch Architectures" Section.

5.2.3.1 SDH Payload Packetization

The SDH signal can be divided into SDH information payload (container) and overheads. The overheads consist of section overhead, path overhead and administrative unit (AU) pointer.

5.2.3.1.1 SDH/ATM Cells

If the SDH payload contains cells, then the location of the first byte of the virtual container (VC-4) path overhead is indicated by the AU pointer. The VC-4 consists of a container (C-4) and the path overhead. The cells can be extracted from the C-4 container by processing the H4 offset within the path overhead. After this, the cells will be grouped into SVPs as discussed above. The interface requirement at the earth station is shown in Figure 5-14.

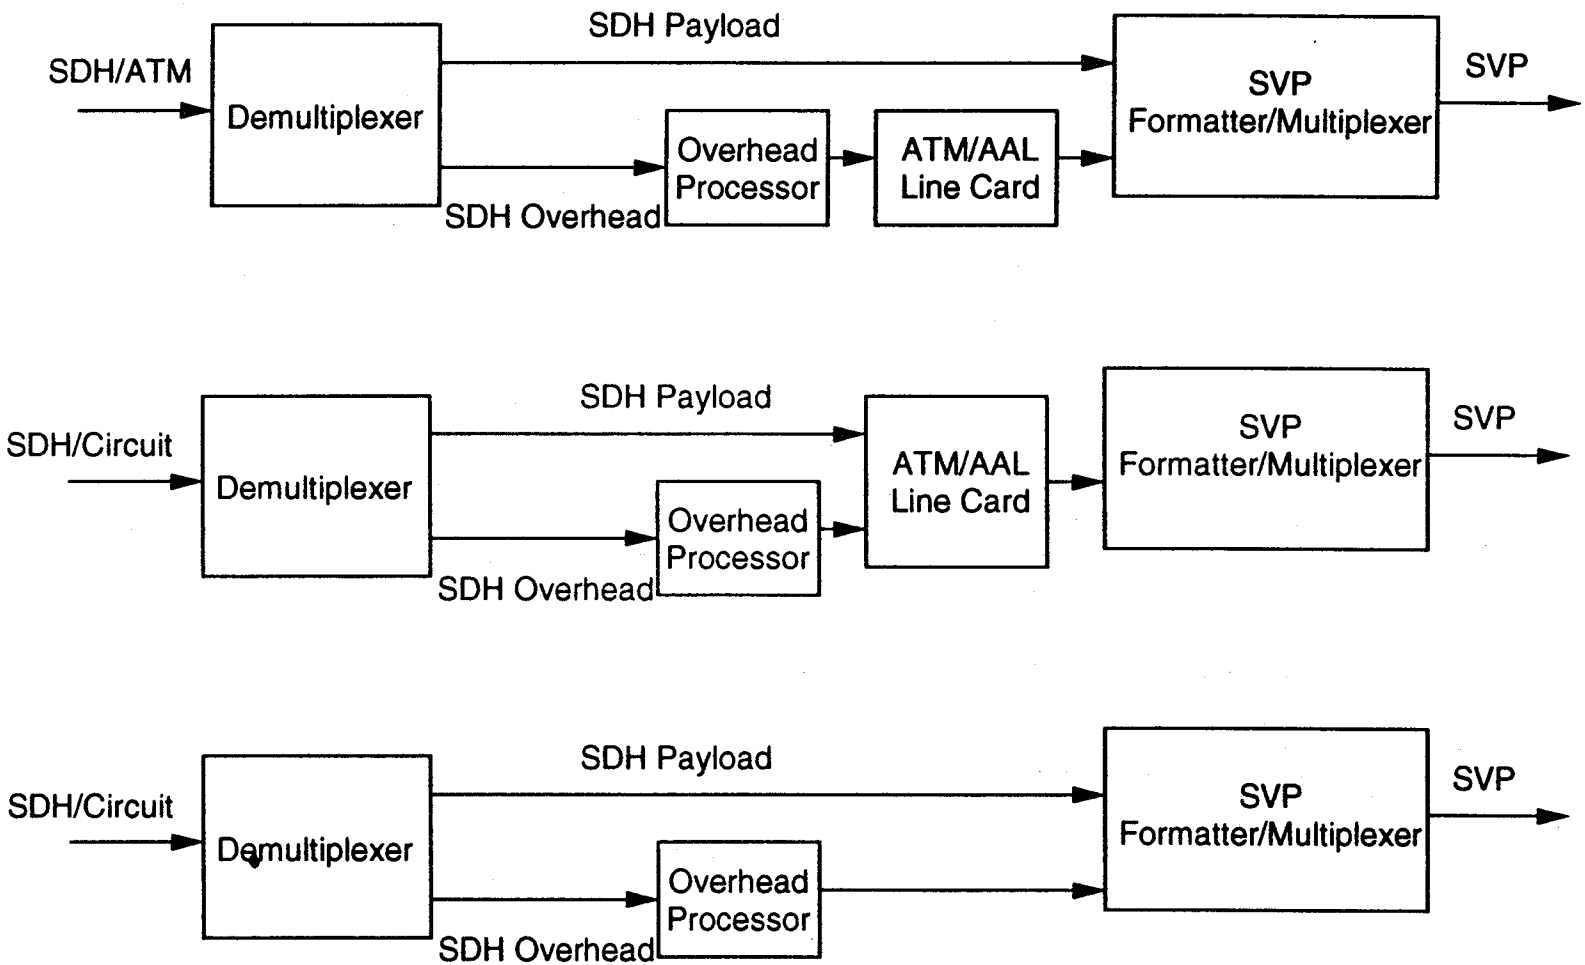

Figure 5-14. SDH Packetization Interfaces

5.2.3.1.2 SDH/Circuit

If the SDH payload contains circuits, then one approach is at the sending station, all these circuits will be adapted to cell formats first by employing an AAL/ATM line card, which is currently under development at COMSAT Labs. After this, these cells will be grouped into SVPs. At the receiving stations, the cells will be extracted from the SVPs first, and then the cells will be assembled back to slot format and be put into the SDH payload. It is envisioned that the earth station will not support many different hierarchical multiplexing schemes so that the hardware requirement of the earth station can be reduced. The interface requirement at the earth station is shown in Figure 5-14. Another approach is to put circuit slots directly into the SVP payload, and designate one bit in the SVL to identify that this packet consists of circuit slots. Since the SVP payload is transparent to the space segment, whether the SVP contains cells or slots has no impact to the space segment. These SVPs can be treated as circuit data, which exhibit periodic and deterministic natures and will simplify the capacity allocation algorithm and congestion control schemes.

5.2.3.2 SDH Overhead Packetization

The SDH contains standardized overhead bits for operation, maintenance, communication, and performance monitoring functions. The advantage of processing the SDH overhead has been discussed at "ATM Cell Concentration" Section. These overheads consist of section overhead (SOH), path overhead (POH), and AU-4 pointer. Since these overheads are in the circuit slot format, one scenario is to employ the AAL/ATM line card to adapt the slots into the cell format. These cells will be packed into the SVP format. The second scenario is to allow the SVP to accommodate the circuit slot format. Using the SVP format, the important overhead functions such as framing, AU-4 pointer and H4 offset can be protected easily using the error control functions. Since not all the overhead functions are important, it might be possible that the only a portion of the OHs are transmitted and the untransmitted OH portion will be regenerated at the receiving stations. At minimum, the earth stations should support the SOH, which means the section overhead should be transmitted. If the earth station acts as a path end point, the path OH has an end-to-end significance; hence, all the POH will be transmitted through the satellite network. The data communication channel (DCC) will be delay compensated at the sending station first, and then be sent to the satellite using the SVP format. These communication channel SVPs will be routed to the appropriate destination without any processing on-board. For multi-destination SDH streams, the destination information is contained in a separate signaling channel (such as SS7). This information can be used to format the point-to-multipoint routing tag either at the earth station or on the satellite. As mentioned before, the duplication of SVPs (which contain SDH overhead) for multiple destinations should be performed at the on-board switch. The interface requirement at the earth station is shown in Figure 5-14.

5.3 Bit Error Rate and Quality of Service

This section addresses the error performance of the ATM cells through the satellite links. The necessary corrective measures once the error performances of some services do not meet the QOS are discussed in the "Transmission System Design" and "Protection of ATM Cell Headers from Burst Errors" sections.

5.3.1 Satellite BER Characteristics

The QOS of ATM connections through the satellite network depends on two factors: the transmission errors (the bit error rate and the statistics of error bursts on the link), and the amount of congestion experienced at the earth station and space segment. If burst error protection is applied on the ATM cell header or the SVP header, the dominate source will be in congestion. To maintain the QOS of some services, the network should be able to selectively discard the low priority cells. The effects of QOS due to congestion at the earth station's multiplexer and the space segment's switch are discussed in Appendix C and D. Please refer to them for more detail. In order to maintain the QOS such as the bit error rate and cell loss ratio (CLR) through the satellite network, all these factors have to be tackled at the same time so that no one dominates the error performance. This section focuses on the effects of the satellite transmission errors to the QOS of ATM connections.

The error patterns of satellite links exhibit random errors and burst errors. The statistics of burst errors in the satellite links depends on the channel characteristics, the modulation scheme, the FEC coding scheme, and the scrambling scheme. Compared with the random bit error rate (BER) of fiber links, which is on the order of 10^{-12}, the satellite links can provide about the same error performance in the clear sky condition; however, the performance of the satellite links is degraded if precipitation happens. Different error control schemes are proposed for the ATM cell transmission via satellite links to improve the error performance in the "Transmission System Design" and "Protection of ATM Cell Headers from Burst Errors" sections.

5.3.2 Error Performance Requirements for Different Protocols

In this subsection, the BER and QOS of different protocols are surveyed. This information helps determining the satellite system availability and BER.

5.3.2.1 Signal System Number 7 (SS7)

As per CCITT Recommendation Q.701, Blue Book, the possibilities of message transfer failures have been specified, and in Q.706 the design objectives for the message transfer part for different parameters have been established. These requirements are summarized as follows.

a. Loss of messages in message transfer protocol (MTP): lower than 10^{-7}.

b. Message out-of-sequence in MTP: The probability is expected to be lower than 10^{-10}; this value also includes duplication of messages.

c. Delivery of false information or undetected errors in MTP: Undetected errors may result in the delivery of false information; the probability is expected to be lower than 10^{-10}.

d. Unavailability of a signaling route set: The unavailability of a signaling route set should not exceed a total of 10 minutes per year.

e. Long-term bit error rate of 64 kbit/s links: lower than 10^{-6}.

f. Medium-term bit error rate of 64 kbit/s links: lower than 10^{-4}.

SS7 is one of the AAL Class 4 services (connectionless data transfer); hence, these requirements give an indication of the error performance requirements of AAL class 4 services.

5.3.2.2 IEEE 802.6 WAN

The objective of the undetected errors should not be greater than 5×10^{-14} per byte of delivered data. The proposed connectionless service sets a delivered error rate of less than 5 in 10^{13} delivered messages.

5.3.2.3 Error Performance of ISDN

The error performance objective for a 64 kbit/s circuit-switched connection has been stated in CCITT Recommendation G.821. It is summarized as follows.

Degraded minutes: fewer than 10% of one-minute intervals to have a bit error rate worse than 10^{-6}. Severely errored seconds: fewer than 0.2% of one-second intervals to have a bit error rate worse than 10^{-3}. Errored seconds: fewer than 8% of one-second intervals to have any errors (equivalent to 92% error-free seconds).

NTT has conducted studies to determine the error performance (the long term BER) requirement for high quality ISDN services [5-1]. The results are summarized in Table 5-3.

The study also shows that to ensure video transmission integrity for B-ISDN, a discarded cell ratio of 10^{-9} to 10^{-10} is required.

Table 5-3. Error Performance Requirements Based on NTT Study [5-1]

SERVICE	BER
Voice (64 kbit/s)	3×10^{-5}
Data and Fax (16 kbit/s, 64 kbit/s)	3×10^{-6}
Facsimile	3×10^{-6}
Data (384 kbit/s - 6,144 kbit/s)	1×10^{-7}
4 MHz Video (Linear PCM)	5×10^{-7}
4 MHz Video (DPCM, 32 Mbit/s)	1×10^{-7}
4 MHz Video (Inter-frame Coding)	1×10^{-9} - 1×10^{-7}

5.3.2.4 SMDS

Switched multi-magabit data service (SMDS) is a high performance, connectionless, and packet switched data service. The QOS of SMDS can be characterized by the availability of service, accuracy, delay and throughput. The accuracy objective is 10^{-10} for errored service data units (SDUs) and 5 X 10^{-8} for undelivered SDUs, where SDU is the basic data unit for the connectionless service.

5.3.2.5 ATM Performance

The performance of an ATM connection depends on the BER of the transmission media and the probability distribution of the bit errors. The fiber system exhibits random errors characteristics; however, there are error bursts if protection switching occurs in the network such as a failure of a repeater. The satellite network exhibits burst errors mainly due to the adoption of forward error correction (FEC) coding for the channels. The ATM cell header has a 1-bit error correction capability which is effective in the fiber system but this capability is hindered by the burst errors in the satellite channels. Bit interleaving is one of the corrective measures, which increases the burst error correcting capability and improve the ATM transmission error performance through the satellite channels. The quality of some services, proposed in the CCITT meeting [5-2], is provided in Table 5-4. The CLR and CIR stand for cell loss ratio and cell insertion ratio, respectively.

For video and stereo sound transmission, since no retransmission is possible, the CLR requirement is high. For data transmission, since ARQ can be used for error control, the CLR is lower. The BER decides the quality of the sound and the video; therefore, telephony, video and stereo sound all have the same BER. Date transmission has a stricter BER requirement so that all received information can be guaranteed almost error free.

.c9.Table 5-4. Quality of Service Requirements Proposed at CCITT Meeting [5-2]

SERVICE	BER	CLR	CIR
Telephony	10^{-6}	10^{-3}	10^{-3}
Data transmission (2 Mbit/s - 10 Mbit/s)	10^{-8}	10^{-6}	10^{-6}
Video telephony/video conference	10^{-6}	10^{-8}	10^{-8}
Broadcast video	10^{-6}	10^{-8}	10^{-8}
HiFi Stereo	10^{-6}	10^{-7}	10^{-7}
User signaling/remote process control	10^{-5}	10^{-3}	10^{-3}

The value of CLR proposed in [3] for continuous bit rate services is between 10^{-6} and 10^{-9}. The design criteria of CLR for a fast packet switch proposed by many researchers is 10^{-9}.

Based on the above information, it is assumed that the satellite system has a 99.9% availability of BER 10^{-8}, and at the clear sky without any sun outage and fading, the BER can achieve 10^{-11}. These issues are discussed in the "Transmission System Design" section. It should be mentioned that this is the general requirement for the channels. For different services, different error performances may be required. If this general error channel requirements cannot satisfy the specific service requirement, corrective measures such as FEC or ARQ have to be provided. The FEC code alternatives are introduced in the "Protection of ATM Cell Headers from Burst Errors" while the ARQ procedures are discussed in the "Impact of the Satellite Delay" section. Note the error performance monitoring is provided at the network control center (NCC).

5.4 ATM Cell Header Protection

In general, two major sources of errors exhibit in the satellite links: the random errors and the burst errors. The random errors can be tackled using the forward error correction (FEC) code and the burst errors can be alleviated using the FEC code along with the bit interleaving scheme. For different FEC codes, different bit interleaving schemes should be incorporated to increase the burst error correction capability and to achieve a better performance.

The satellite virtual packet (SVP) concept has been proposed in the "ATM Cell Packetization Issues and Alternatives" section to simplify network routing and control. Since the SVP is not limited by the 53-byte cell size, different FEC schemes can be selected to improve the error performance of SVP transmission via satellite links. The issues of how to improve the error performance of the SVPs through satellite links are studied in this section. As mentioned above, the error patterns of satellite links exhibit random errors and burst errors. The statistics of burst errors in the satellite links depends on the channel characteristics, the modulation scheme, the FEC coding scheme,

and the scrambling scheme. To combat the burst errors, special operation has to be performed on ATM cells or SVPs. In this section, the methods to protect the ATM cell headers, SVP payload, and SVP header (SVL) against burst errors are also discussed.

Several other areas worth studying are the methods to improve the AAL protocol data unit (PDU) error performance and the synchronous transport module (STM) virtual container loss through the satellite links. Enhancement of the AAL protocol data unit error performance is accomplished by bit interleaving the PDU if a correction capability is provided within the AAL layer. Nevertheless, this AAL PDU error performance is not an issue if error protection is provided for the SVP payload (since the AAL PDU is part of the SVP payload). Hence, how to improve the AAL PDU error performance is not discussed here, and the focus is on how to improve the error performance of the SVPs. If the assured operation in the AAL class three and four services is necessary, then efficient automatic repeat request (ARQ) protocol should be used to guarantee that the AAL PDU is delivered with the exact same content. The efficient ARQ protocol for the AAL PDU is discussed in the "Impact of Satellite Delay" section.

5.4.1 ATM Cell Delineation, Error Detection and Error Correction

Within the 5-byte ATM cell header, there is a 1-byte header error control (HEC). As per CCITT Recommendation I.321, the consequence of errors in the ATM cell header is shown in Figure 5-15. If an error is undetected in the cell header, this cell is called the valid cell with an errored header. There are two consequences of this kind of cells. First if errors occur at the VPI/VCI field, and the errored VPI/VCI address has been assigned by the network, then this cell will be misdelivered to a wrong destination. Secondly, if the address has not been assigned, then this cell will be discarded by the network. A cell is sent to the correct destination if there is no error in the header or the error can be corrected, and this cell does not suffer from network congestion.

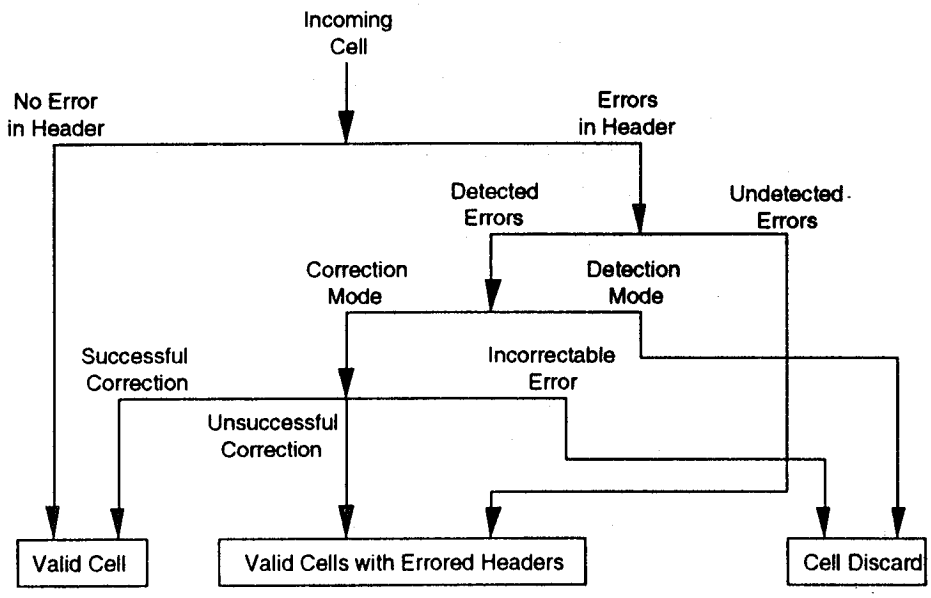

Figure 5-15. Consequences of Errors in ATM Cell Headers

The receiver has two modes of HEC operation (correction and detection modes) as shown in Figure 5-16. The default mode is the correction mode. In the correction mode, both the error correction and detection capabilities are invoked. In the detection mode, only the detection capability is invoked. The main reason two modes are adopted in the receiver is that when the receiver operates in the detection mode, the probability of detecting multiple bit errors is increased compared with the correction mode.

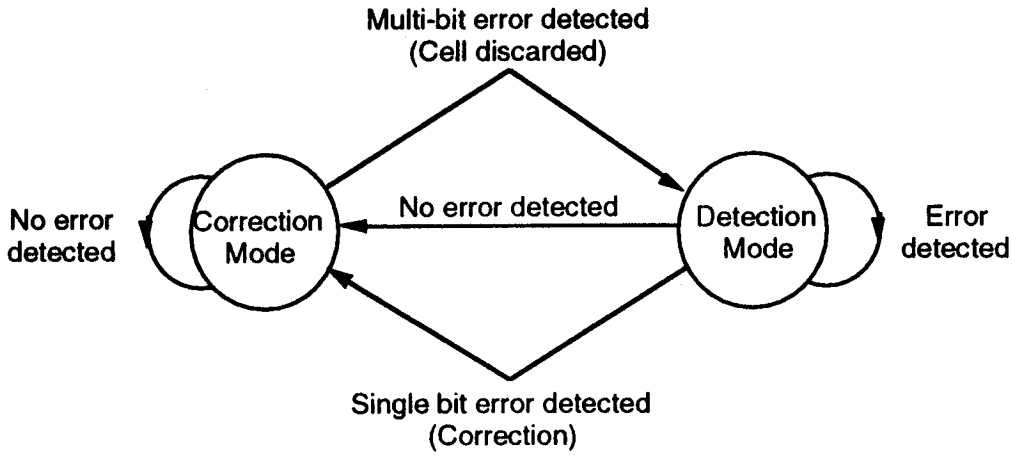

Figure 5-16. HEC Receiver Modes of Operation

When a single bit error is detected in the correction mode, the cell header is corrected and the operation mode is switched to the detection mode. When multiple bit errors are detected in the correction mode, the operation mode is also switched to the detection mode. When the receiver is in the detection mode and no error is detected in the incoming cell, the receiver is switched back to the correction mode. If there are still errors in the incoming cell, the receiver stays in the detection mode. The advantage of switching to the detection mode is that if the single bit error is the start of a burst error, then the probability of miscorrection is reduced if the decoder only performs error detection. However, there is an associated penalty with this switching mechanism. In the detection mode, the incoming cell is not corrected even though there is only one bit error within the cell. The result is cells with a 1-bit error are dropped in the detection mode. This penalty is more severe in the satellite environment since the cells are bit interleaved before transmission. Hence 1-bit error patterns in the cell header are expected to be dominant. In order to compensate this penalty, the bit interleaving scheme should assure that the adjacent cells are separated as far as possible so that the burst errors do not affect two adjacent cells at a time. Hence, the operation mode at the receiver can always transit back to the correction mode from the detection mode quickly.

The generator polynomial of the HEC is a (40,32) shortened cyclic code. In other words, the message consists of 32 bits and the code word consists of 40 bits, and there are 8 control (parity) bits. The $g(x) = 1+x+x^2+x^8 = (1+x)(1+x^2+x^3+x^4+x^5+x^6+x^7)$, where $(1+x^2+x^3+x^4+x^5+x^6+x^7)$ is an irreducible and primitive polynomial. To analyze the error correction and protection capability of the shortened cyclic code, the original code

has to be known. The original code is the $(2^7-1, 120)$ cyclic code. The error protection function provided by this code can recover all the single bit errors and a certain percentage of multiple bit errors. The error detection capability of this code can be known by analyzing the error detection capabilities of (1+x) and $(1+x^2+x^3+x^4+x^5+x^6+x^7)$, separately. First, a polynomial is divisible by (1+x) if this polynomial contains an even number of terms, This property implicitly implies the error detection capability of (1+x). If the error polynomial consists of an odd number of terms, then this error polynomial cannot be divisible by (1+x). Hence with (1+x) as a factor in the HEC generator, it can detect all errors with odd weight. The $(1+x^2+x^3+x^4+x^5+x^6+x^7)$ is a primitive polynomial. The property of a primitive polynomial with degree m is that the smallest positive integer n for which this polynomial divides x^n+1 is $n = 2^m - 1$. Applying this to the error detection capability, this primitive polynomial cannot divide any error polynomial x^n+1 for $1 \leq n < 2^m - 1$. Hence, this primitive polynomial can detect any weight 2 error in a code word of 2^7-1 bits or less.

The recommended cell delineation method is using the correlation between the 4-byte header and the 1-byte HEC. The cell delineation state diagram is shown in Figure 5-17. The parameters m and n have to be chosen to make the cell delineation process as robust as possible. Robustness against false misalignment due to bit errors depends on m. Robustness against false delineation in the resynchronization process depends on n.

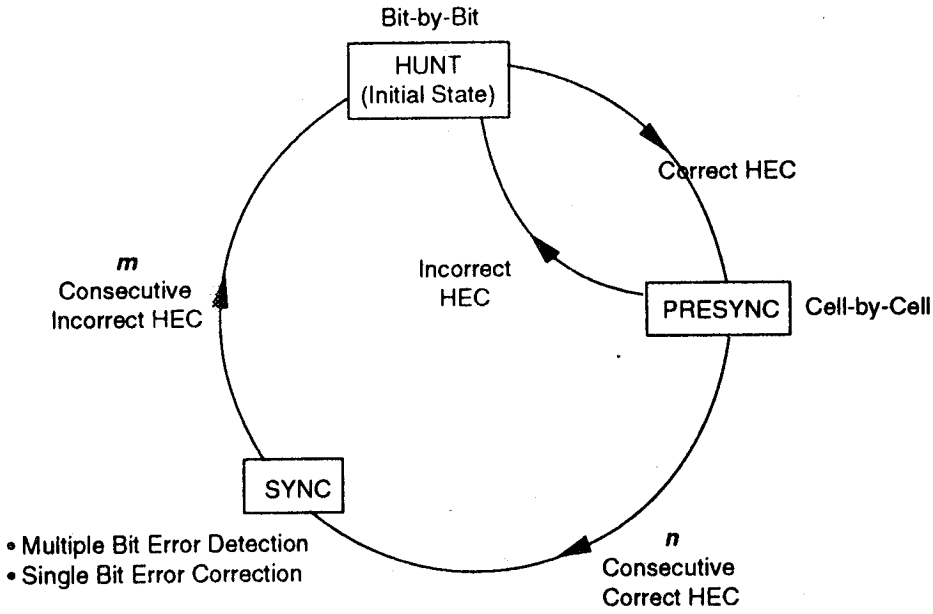

Figure 5-17. Cell Delineation State Diagram

5.4.1.1 Error Detection Capability in Correction and Detection Modes

The ATM HEC detection capabilities in correction and detection modes are compared below. First assume the error operation mode is in error detection mode. In general, for

an (n,k) cyclic code, the length of the code word, v(x), is n, the length of the message, u(x), is k, and the length of the generator polynomial, g(x), is n-k+1. The decoding procedure is to compute the syndrome, which is obtained from dividing the received code word, r(x), by the generator polynomial, g(x). If the syndrome is not equal to zero, then the decoder declares that the received code word is in error, and the result is that the cell is discarded. However, when the syndrome is equal to zero, the decoder can only guarantee that the received word is one of the 2^k code words whose syndromes are zero. Hence these 2^k code words are categorized as the undetectable error patterns, which means if the error changes the original code word into one of these 2^k code words, the error is not detectable. Assume the probability that the appearance of all possible messages is equal. Then the probability of an undetected error is 2^k - 1/ 2^n. For the ATM cell, this probability is 2^{32} - 1/ 2^{40}, which is approximately equal to 0.0039.

Secondly assume the decoder is in the correction mode. In this case, the (n,k) cyclic code is capable of correcting 1-bit errors. Now the received code word is not in error if the syndrome of the received code word is equal to zero or if the syndrome has the error pattern which corresponds to an 1-bit error. When the syndrome is equal to zero, the received code word might be one of the 2^k code words. And each code word with length n, there are n possible 1-bit error patterns. Hence the total number of received code words, which result in the zero syndrome (after correction), is 2^k (n+1). Assume the probability that the appearance of all possible messages is equal. Then the probability of an undetected error is $(n+1)(2^k - 1) / 2^n$. For the ATM cell, this probability is $41(2^{32} - 1) / 2^{40}$, which is approximately equal to 0.16.

The results above show that when the decoder is used in the correction mode, the error detection capability is reduced. This may result in a large probability of undetected errors in ATM cells if burst errors are encountered. However, for the satellite application since bit interleaving is used as a corrective measure for the burst errors, the advantage of using the detection mode and the correction mode alternatively is not significant.

5.4.1.2 Protection of ATM Cell Headers from Burst Error

5.4.1.2.1 Bit Interleaving

A satellite system, which exhibits burst errors in the transmission links, can still utilize the 1-bit error correction capability of ATM HEC by spreading the bits of the error burst among many cell headers so that only a 1-bit error is experienced in the ATM cell header. This procedure of spreading the bits in the error burst among different cell headers is called bit interleaving. The AAL might also provide the error correction capability, which is still under study at the CCITT meeting. If the AAL does provide the error correction capability, then bit interleaving the cell payload (48 bytes) improves the performance of the adaptation layer. If the AAL only has error detection capability, bit interleaving over the cell payload may increase the number of cells which contains errors.

How to perform bit interleaving is a design issue. Basically, this is decided by the complexity of the bit interleaving logic, the BER performance improvement, and the cell delineation algorithm (the decoding algorithm). In general, bit interleaving can be classified into two categories: intra-cell and inter-cell bit interleaving. In either case, the interleaving size or the interleaving depth (the number of cells being interleaved) should be large enough to successfully correct the burst errors. The objective is that the cell header should have at most 1-bit error after bit interleaving. To perform bit interleaving, temporary storage of cells is necessary, which results in the cell queueing delay. The interleaving depth (number of cells) is constrained by the cell transfer delay requirement. To reduce the interleaving delay, high information transmission rate should be used; hence, it is advantageous that bit interleaving is performed after the cell streams are multiplexed. In this context, it is more advantageous of performing bit interleaving on the satellite virtual packets instead on the cells since SVPs are the results of multiplexing several different input cell streams.

5.4.1.2.2 Pseudorandom Interleaving

Although bit interleaving alleviates the burst error problem, it might create new problems if a periodic noise source is encountered in the transmission links. The interleaving depth and the transmission speed have to chosen carefully so that no new error bursts occurs. The other way of avoiding creating new error bursts due to periodic noises is using a pseudorandom interleaving scheme. One way of implementing the pseudorandom interleaving scheme is using scrambling.

5.4.1.2.3 Inter-Cell Bit Interleaving

The inter-cell bit interleaving means at the transmitter the cells are stored in one direction while the bits are read out in the vertical direction. At the receiver, a reverse operation is performed. In Figure 5-18, the transmission order of the bits of multiple ATM cell is shown for without and with bit interleaving at the transmitter. If the interleaving depth exceeds the length of the error burst, ATM cell headers only have at most 1-bit errors.

Because of the detection-correction alternating mechanism used in the cell HEC delineation algorithm, it is best that in two adjacent cells, there is only one cell containing errors. In this way, all the cells containing an 1-bit error are corrected. Otherwise if both adjacent cell contain an 1-bit error, one cell is dropped due to the detection mode has no correction capability. This can be achieved by arranging the sequence of reading the cells from the memory at the transmitter so that the positions of any two bits of the two adjacent cells are separated as far as possible. If intra-cell bit interleaving, which will be discussed latter, has been performed, then instead of reading our bits from cells, reading one byte from each cell is feasible.

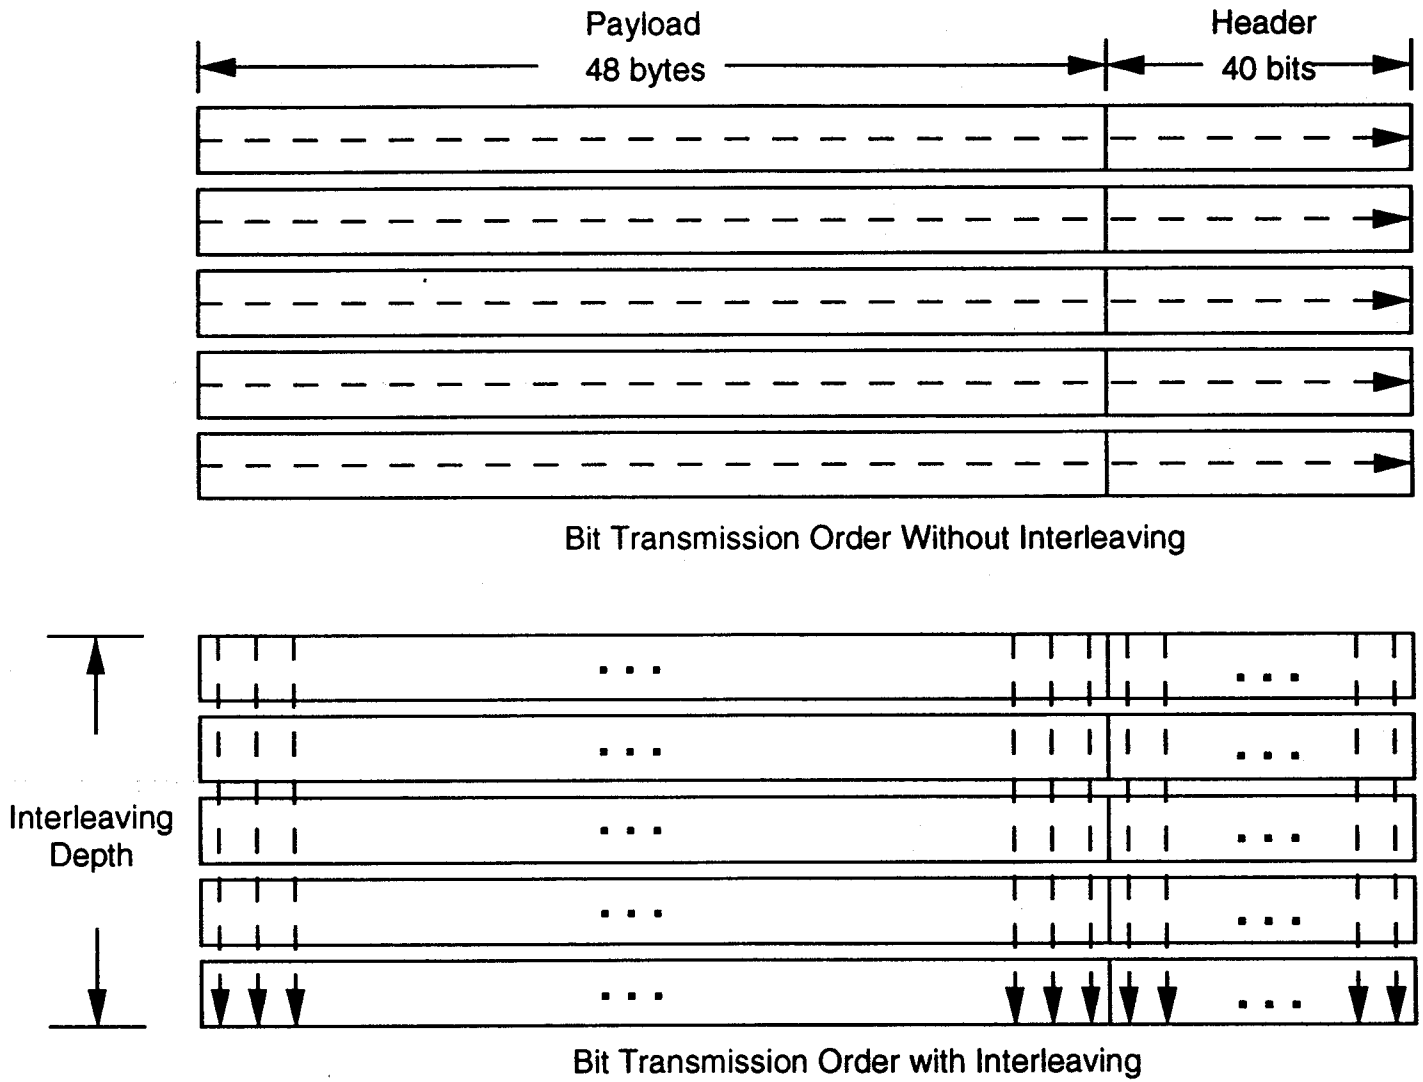

Figure 5-18. Inter-Cell Bit Interleaving

5.4.1.2.4 Intra-Cell Bit Interleaving

Intra-cell bit interleaving means that the header bits of one cell are spread among the same cell payload (see Figure 5-19). Since the number of bits in the header is 40, each bit in the header may be multiplexed with 7 bits in the payload for 40 times. After intra-cell bit interleaving, the format of the first 40-byte data consists of 1-bit header and 7-bit payload, and the format of the last 13-byte data consists of pure 8-bit payload. Without performing intra-cell bit interleaving, the headers of different cells are multiplexed together after inter-cell bit interleaving. If a long error burst occurs, it is possible that this error burst corrupts more than one bit in the header of the same cell. Since the ATM cell has only a 1-bit error correction capability, this error burst results in loss of many cells. By contrast if the header bits are spread around the cell payload, then after inter-cell bit interleaving the distance of any two bits in the same header is large; hence, each time an error burst can only destroy one bit in the header. To achieve the same effect of intra-cell bit interleaving is to bit interleave a very large number of cells using inter-cell bit interleaving, where the number (or the interleaving depth) is larger than the size of the error burst. However, as mentioned before, the inter-cell interleaving depth is constrained by the delay requirement. The advantage of intra-cell bit interleaving is that the number of cells needed for the inter-cell nit interleaving can be reduced with the cost of higher hardware complexity. Hence, inter-cell bit interleaving may be a necessary compromise between the cell loss ratio and the cell transfer delay.

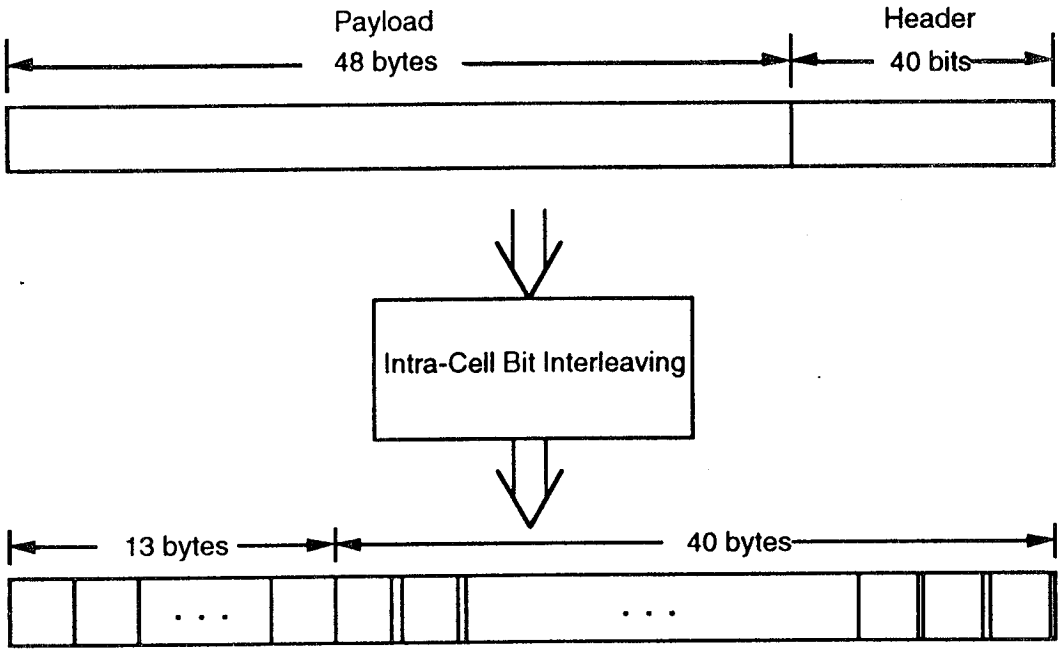

Figure 5-19. Intra-Cell Bit Interleaving

5.4.1.2.5 Protection of SVPs from Burst Errors

As mentioned in the "ATM Cells Packetization Issues and Alternatives", the size of a SVP is a design issue based on many considerations. According to the simulation results provided in Appendix D, the SVP packetization delay (the time required to fill the SVP with cells) dominates the queueing delay at the multiplexer of the earth station. Hence, the suggestion is that the SVP size should be ≤ 4 to minimize the packetization delay or some other corrective measures should be provided (such as a timer for each SVP) if a SVP with a larger size is used. One of the advantages of performing SVP packetization is that the bit interleaving can be naturally achieved among the cells within the SVP and the SVL (see Figure 5-20). Hence the packetization delay of a SVP occurring at the earth station already includes the write cycle of the interleaving delay at the interleaving buffer.

There are two variations of interleaving the SVP depending on whether the SVL is bit interleaved or not. The first one is not to bit interleave the SVL with the cells within the SVP. However, the SVL is protected using a burst error correcting code. In this scheme, the on-board processor complexity is reduced since the SVP payload is independent from the SVL. To route and process the SVP, the processor only has to decode the SVL. No bit deinterleaving buffer is required in this case. The second one is to bit interleave the SVL with the cells. In order to decode the SVL, bit deinterleaving of the SVP is required on-board. Although the impact of burst error to the SVL is reduced due to bit interleaving, this scheme results more buffer requirement for the on-board processor and more cell end-to-end delay.

If the SVP size is small, the bit interleaving scheme among the cells within a SVP to protect the SVP from burst error is not effective since the interleaving depth is much smaller than the average burst error. There are two corrective measures. The first one is to use more SVPs for bit interleaving. However, this approach is exactly the same as using a SVP with a larger size. Hence, the same problem occurs, i.e., the interleaving delay (which is dominated by the time for the interleaving buffer to fill with SVPs) is large. If this is the approach to be employed, it is better off to use a SVP with a large size since the overhead (the SVP header) associated with one SVP is smaller than the overhead using several SVPs. The second approach is to use a burst error correcting code such as the Reed-Solomon code to protect the ATM cells within a SVP in conjunction with the bit interleaving scheme. However, the coding overhead associated with the scheme is high which results in low link efficiency, and the on-board processing complexity is increased.

5.4.1.3 Cell Loss Ratio

Several factors can lead to loss of cells: random errors and burst errors in the transmission links, and congestion in a statistical multiplexer or a packet switching node. As mentioned before, FEC coding is one way of improving the performance of CLR. However, since coding increases the intensity of traffic and the bandwidth

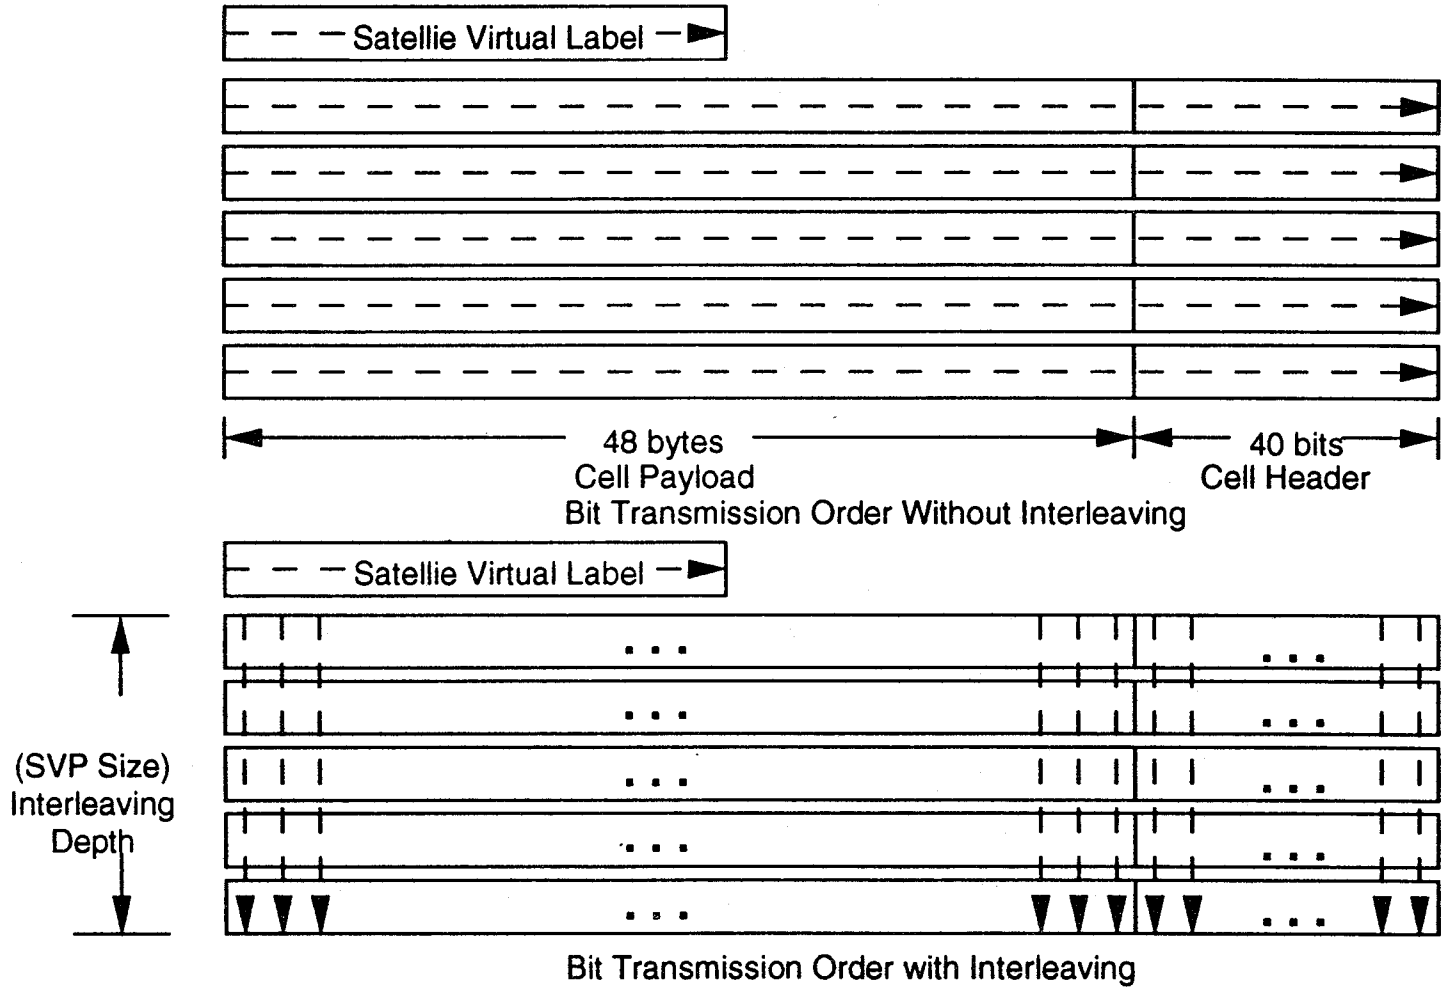

Figure 5-20. SVP Bit Interleaving

requirement, congestion might be more server and the CLR might be higher. The objective of coding is that the cell loss ratio after decoding should be smaller than the cell loss ratio without any coding.

The cell loss ratio (CLR) due to congestion within a statistical multiplexer has been analyzed in Appendix C. In this subsection, the CLR due to random errors and burst errors of the transmission links is studied. The CLR improvement of using inter-cell bit interleaving to combat burst errors is also analyzed.

5.4.1.3.1 CLR due to Random Errors

Assume the BER of the satellite channels is P_e. If the receiver is operated in the correction mode, the cell is lost if the number of bits in error in the cell header contains is more than one. The probability that there is more than 1 bit in error in the cell header is equal to

$$1 - \sum_{k=0}^{1} \binom{40}{k} (P_e)^k (1 - P_e)^{40-k} \text{ , where } \binom{40}{k} = \frac{40!}{k!\,(40-k)!}$$

Remember that the probability of undetected errors in the header, given that the cell header contains errors and the receiver operates in the correction mode, is $[2^{32}(41) - (41)] / 2^{40}$. Although these cells whose errors were not detected are not lost, they are dropped by the higher layer protocol. However the loss of these cells cannot be measured at the physical layer of the network. Hence the CLR measured at the physical layer of the network should be modified as:

$$CLR = [1 - \sum_{k=0}^{1} \binom{40}{k} (P_e)^k (1 - P_e)^{40-k}] \; (1 - [2^{32}(41) - (41)] / 2^{40})$$

In the error detection mode, the cell is lost unless the cell header has no error. Using the same principle as above, the CLR can also be derived.

$$CLR = [1 - (1 - P_e)^{40}] \; (1 - [2^{32} - 1] / 2^{40})$$

The effect of the random errors to the CLR at the receiver is shown in Figure 5-21 for different operational modes.

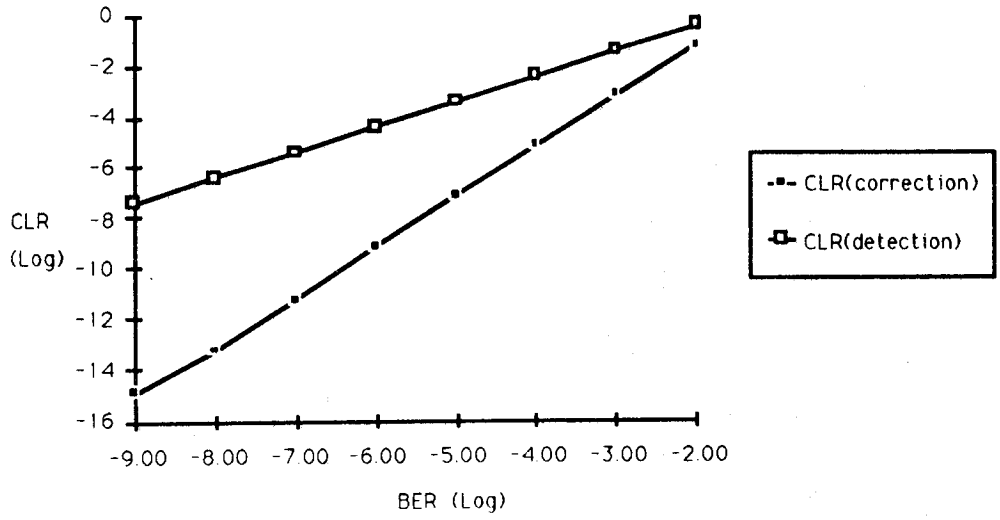

Figure 5-21. Cell Loss Ratio vs. Bit Error Rate for Different Operational Modes

5.4.1.3.2 CLR due to Burst Errors

Before calculating the CLR due to burst errors, the burst error must be defined first. Let s(t) be the sending information bit sequence and r(t) be the receiving information bit sequence. A burst error of length B is defined as

$s(t_0) \neq r(t_0)$ and

$s(t_0+B) \neq r(t_0+B)$ for any t_0.

During the burst length B, s(t) and r(t) agree for at most g-1 consecutive bits where g is the guard space between bursts of errors. In other words, two error bursts are separated by at least g error free bits. The terminology used in this section is introduced as follows.

N_e the number of error bits

N_b the number of error bursts

P_e the bit error rate

P_b the burst error rate, which is the probability that a burst is started at a given bit.

B the average error burst length

d the error density (the percentage of error bits in a burst)

The relationship between N_e and N_b is

$$N_e = d\ B\ N_b$$

Hence the burst error rate can be written as follows.

$$P_b = \frac{P_e}{B\ d}$$

To calculate the CLR, certain assumptions are made. Assume if two bits or more of an error burst falls in the cell header, then the cell is dropped due to uncorrectable errors in the header. Two cases can be considered. First, the average error burst length is ≤ 48 bytes. Secondly, the average error burst length is > 48 bytes.

Consider the first case. All the possibilities that there are two bits or more falling within the header are depicted in Figure 5-22. There are two extreme cases: the last two bits of an error burst coincide with the first two bits of the header and the first two bits of an error burst coincide with the last two bits of the header. Hence, the CLR is derived as

$$CLR = P_b\ \frac{B + H - 4}{424}$$

Consider the second case. Since the average error burst length is larger than 48 bytes, whenever there is an error burst, at least one cell header is corrupted by two or more errors.

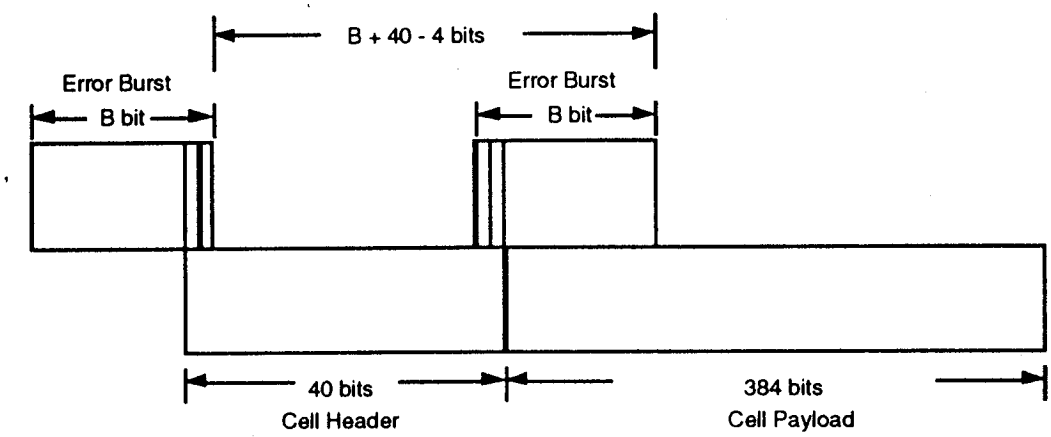

Figure 5-22. Possible Positions of the Error Burst Which Causes the Loss of a Cell

A more accurate analysis of CLR due to burst errors is provided as follows. Assume that the number of bits of an error burst falls within the header is n, the probability that the cell is lost if all these n bits are error free or only one of the n bits is in error. Hence, the probability that a cell is lost given n bits of an error burst fall within the header is

$$1 - ((1-d)^n + n\,d\,(1-d)^{n-1})$$

First assume that B is less than 40. The smallest and largest possible n is 2 and B, respectively. For $2 \le n \le B - 1$, the number of times that n may occur is 2. For $n = B$, the number of times that n may occur is $(40 - B + 1)$. Hence, CLR is expressed as follows.

$$CLR = P_b\{ \sum_{n=2}^{B-1} 2[1 - ((1-d)^n + n\,d\,(1-d)^{n-1})] + (40-B+1)\,[1 - ((1-d)^B + B\,d\,(1-d)^{B-1})]\}$$

Secondly assume that B is equal to or larger than 40.

$$CLR = P_b\{ \sum_{n=2}^{39} 2[1 - ((1-d)^n + n\,d\,(1-d)^{n-1})] + (B-40+1)\,[1 - ((1-d)^{40} + 40\,d\,(1-d)^{39})]\}$$

5.4.1.3.3 CLR of Bit Interleaving

With bit interleaving, a cell header is corrupted if 1) two or more error bursts occur and each burst destroys one bit of the header or 2) the length of a single error burst is larger than the number of cells being interleaved. If the interleaving depth is larger enough, the only possibility for a cell header being corrupted is condition 1. For simplicity, assume that a cell header is corrupted only under condition 1. The probability that three bursts or more occur is much less than the probability that two bursts occur; hence, it is assumed that condition 1 is dominated by the situation that there are two bursts. Assume that number of coinciding bits between two error bursts is n. Then the probability that a cell is lost is given as follows.

> $1 - (1 - d^2)^n$, where $1 - d^2$ is the probability that the two coinciding bits in two bursts are not in error at the same time.

The CLR is given as

$$CLR = P_b^2 \{ \sum_{n=1}^{B-1} 2[1 - (1 - d^2)^n] + 1 - (1 - d^2)^B \}$$

Figure 5-23 shows the effect of burst errors to the CLR with and without bit interleaving.

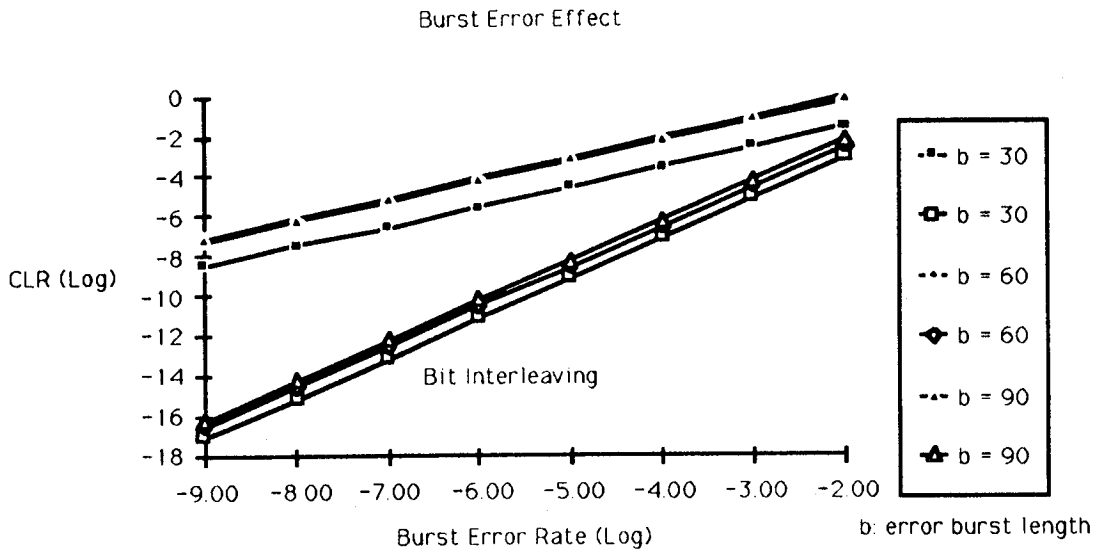

Figure 5-23. Cell Loss Ratio vs. Burst Error Rate for Different Error Burst Length

5.4.2 SVP Error Control

Satellite virtual packets (SVPs) have been proposed in the "ATM Cell Packetization Issues and Alternatives" section to group cells destined for the same down link beam for unified routing, control, and management. The header of the SVP is termed the satellite virtual label (SVL). In this section different coding schemes are proposed to perform error control of the SVP payloads and SVLs. It is noted that a good coding scheme should result in minimal delay, high data throughput, low storage, and low hardware complexity.

5.4.2.1 Coding Options

Two main transmission scenarios are considered depending on whether convolutional coding is used: the first scenario is that the QPSK modulation and convolutional code are optimized for link performance, and the second scenario is only QPSK modulation is used.

The choice of using a block code or a convolutional code largely depends on the link performance. If there are only random errors on the link, then convolutional code is sufficient. If there are only bursty errors on the link, then Reed-Solomon code is sufficient. If there are random errors and bursty errors present at the same time on the link, then concatenated code (inner code is convolutional code and outer code is the Reed-Solomon code) is necessary The other factor should be considered is for convolutional coding, the flush bits are needed at the end of the code word to clear the dependency between two code words. Otherwise, the receiving earth station (or the

space segment) has to have a different decoder for each corresponding sending earth station.

For on-board coding/decoding, several options can be selected. The first is to perform error correction only. The second one is to perform error detection only. In this case, it is feasible to add one more symbol to the code word so that the error detection capability can be enlarged. The third one is to perform error detection and error correction at the same time. For this section, it is assumed that on-board processor performs error correction only.

5.4.2.1.1 Convolutional Code

The first scenario is that the convolutional code and QPSK modulation are used to optimize the line performance (see Figure 10). The advantage of this scenario is that the convolutional decoder can utilize the unquantized outputs from the demodulator (soft decision decoding) and improve the error performance more than using the hard decision decoding. The convolutional code operates on one block of data, which can be viewed as transmission coding. The boundary of the block does not have to be aligned with the boundary of the SVPs. In this case, decoding of the whole block has to be performed on-board. After decoding, the next step is to search the start of the packet. Whenever the packet header is found, the SVP can be routed through the switch. If the convolutional code can achieve the required error performance, then no extra coding is necessary for the SVP. The on-board processor does not perform block code decoding although the block code might still be one of the fields in the SVL. If the SVL contains the CRC code and this code is not decoded on-board, the CRC code has an end-to-end significance. The on-board logic is simplified in this case. In the "RF Transmission System Design" Section, the convolutional code is assumed to be Viterbi, the number of quantization levels is 8 and the constraint length is 7.

If the error performance cannot satisfy the requirement or some services demand a higher error performance, additional coding is necessary to protect the SVP (see Figure 5-24). There are several options for protecting the SVP. The first one is to use a block

code to protect the SVP. The second one is to use a block code (the CRC code) to protect the SVL only. The first option creates a large coding overhead, and the overall channel data utilization is very low. These options are considered below.

In these options, the SVP or the SVL is protected by a block code. The quantized outputs from the convolutional decoder can feed into this block code (such as the R-S code). Since the R-S code can correct symbol errors, the SVP or SVL error performance is good if the error burst is within one symbol. However, if the burst error length is large, then symbol interleaving is necessary. The principle of bit interleaving can be applied to the symbol interleaving. This coding design is termed the concatenated code; the convolutional code is the inner code and the block code is the outer code. The inner code should be chosen such that most of the channel symbol errors are corrected and the outer code simply improves the BER performance to the required level. The channel the inner code seen is the uncoded channel, and the channel the block code seen is called the

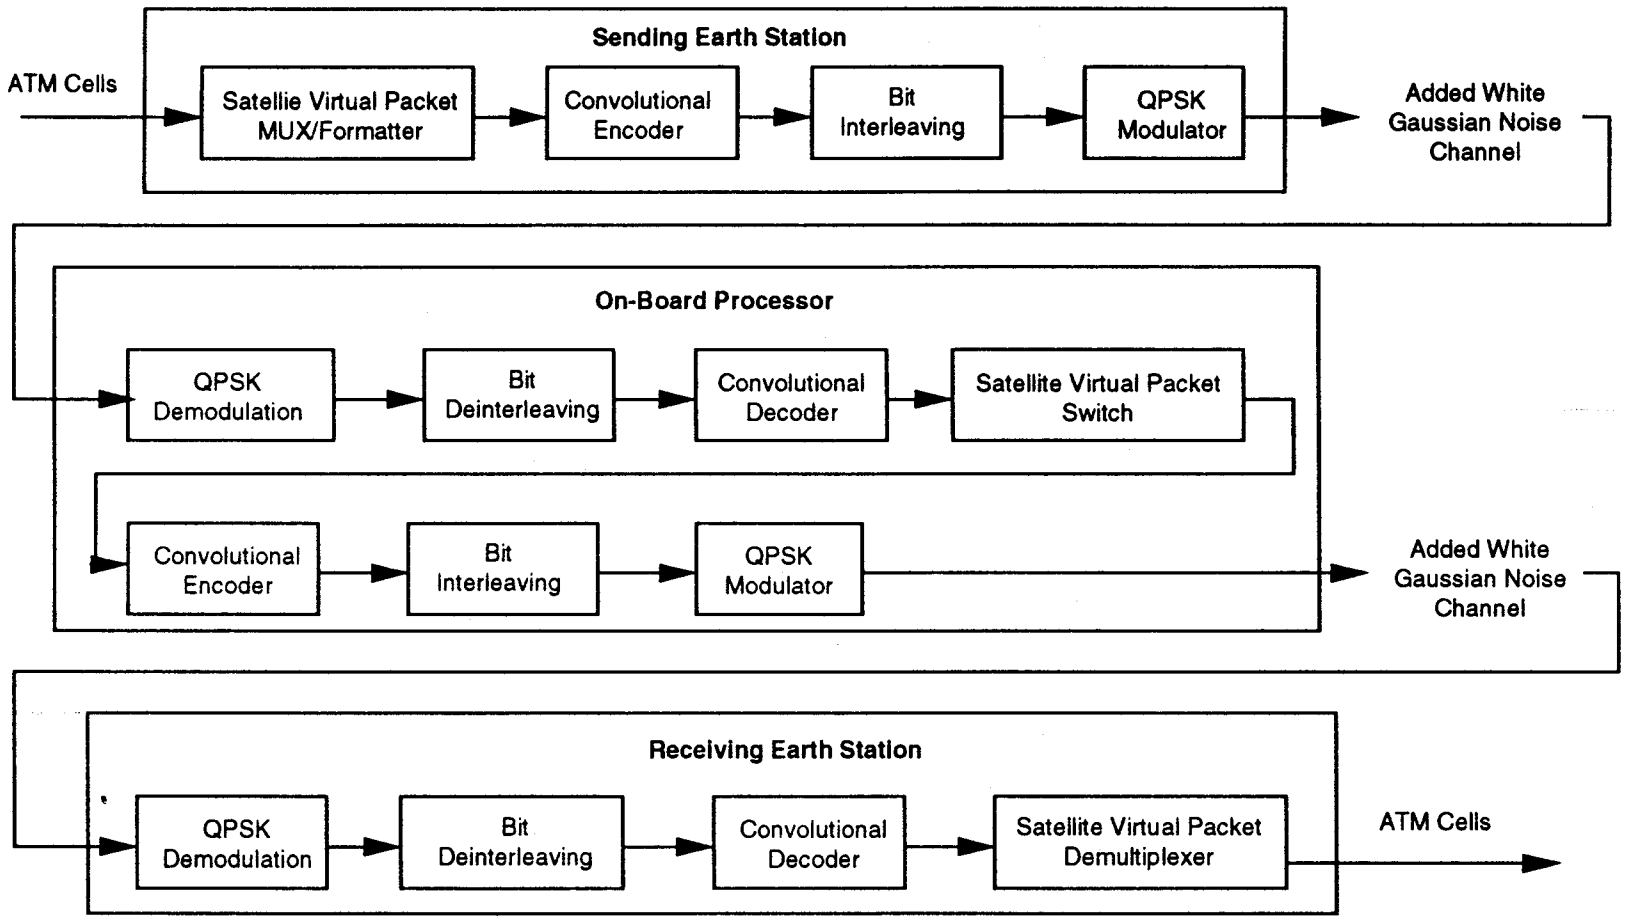

Figure 5-24. Satellite Communication System Model (Convolutional Coding Only)

super channel. It should be noted that one of the objectives of using the concatenated code is to reduce overall implementation complexity compared with using only one code. This scenario is depicted in Figure 5-25, where two decoding processes are necessary on-board the satellite. The first decoder decodes a block of data and the second decoder decodes the SVP or the SVL. Although this scenario has the most complicated logic, it is expected that the link margin can be largely improved. The error performance of this scenario can be calculated in two steps: First, the symbol error probability out of the Viterbi decoder can be calculated assuming the symbol errors are independent due to symbol interleaving. Then the bit error probability out of the block decoder can be calculated using the binomial formula. To have a fair comparison, bit error rate versus E_b/N_0 curve is usually used to show the performance.

5.4.2.1.2 Block Code

In this scenario, no convolutional coding is used instead the block code is chosen for protecting the SVPs (see Figure 5-26). For high speed transmission and for easy hardware implementation of the on-board decoder, the block code is more suitable than the convolutional code. A shortened block code should be chosen to improve the bandwidth efficiency so that no extra bandwidth is wasted if a suitable code length cannot be found.

There are several ways of performing error control of the SVPs using the block code. The first one is to use one block code for the whole SVP. The second one is using a block code for the SVP payload and another block code for the SVL. In this scenario, since coding of the SVP payload and coding of SVL are separate, only SVL decoding is necessary on-board. The transmission SVP payload format is not disturbed. The third one is only the SVL is protected by one block code.

5.4.2.2 The CCSDS Error Control Code

The Reed-Solomon code has been suggested in the consultative committee for space data systems (CCSDS) to protect the virtual channel data unit (VCDU) and to provide different grade of services (GOSs). The minimum VCDU length is 124 bytes and the maximum VCDU length is 1275 bytes. If these numbers are translated into cell size, one VCDU contains at least 3 cells and at most 21 cells (excluding the VCDU primary header and trailer overheads).

Three optional error control fields may be added to the VCDU to support different grade of services (GOSs). Two error control schemes are suggested for the header and one error control scheme is for the payload.

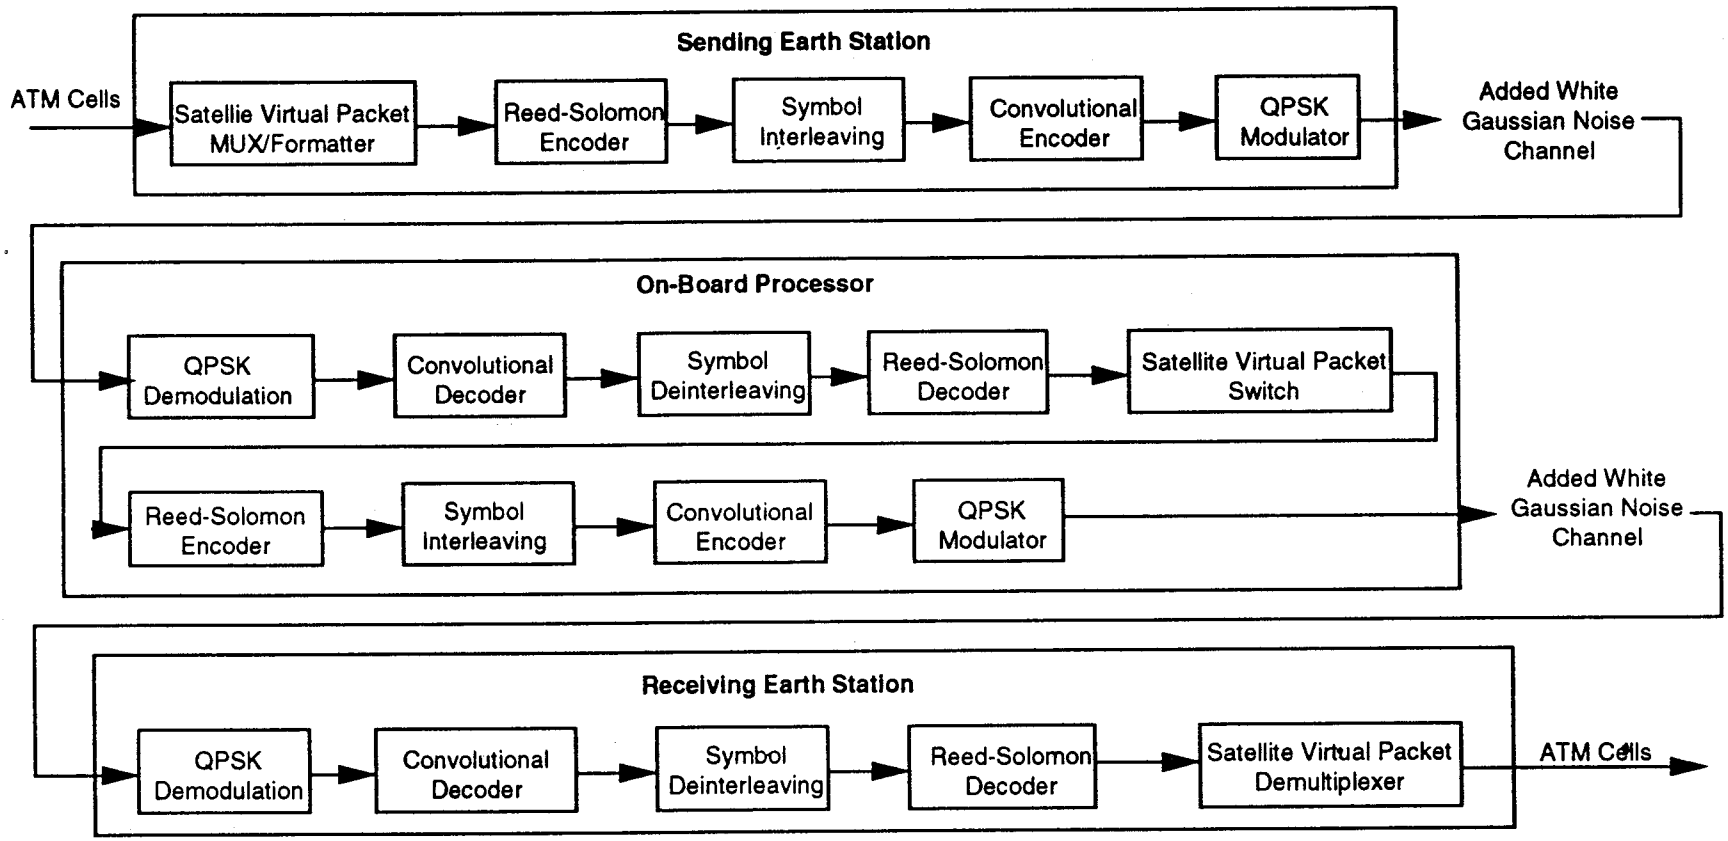

Figure 5-25. Communication System Model (Convolutional + Block Coding)

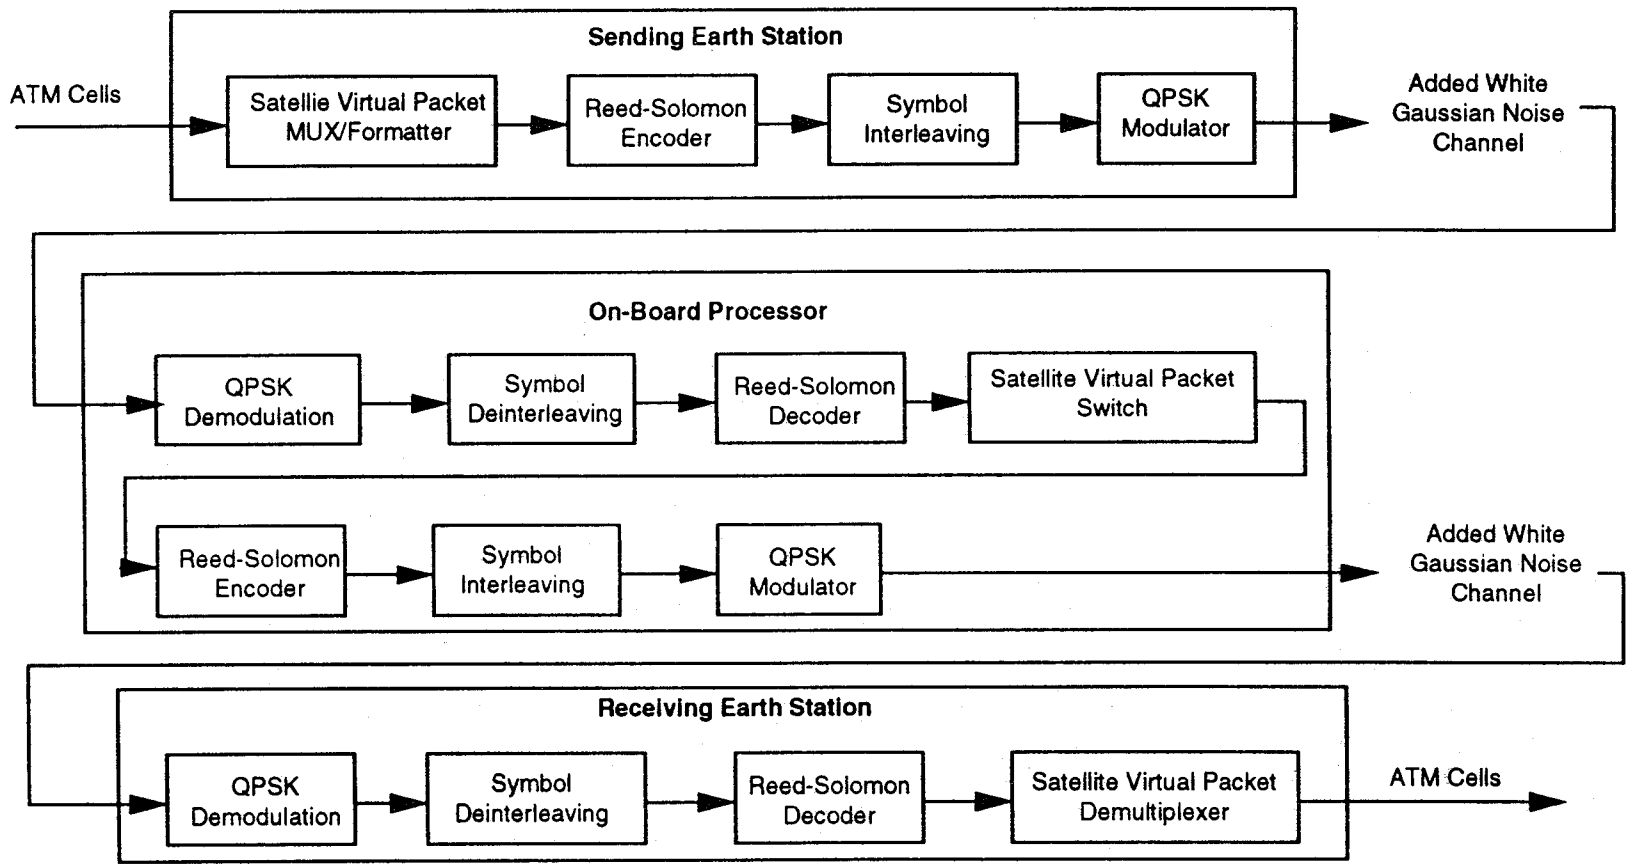

Figure 5-26. Satellite Communication System Model (Block Coding Only)

The Reed-Solomon (R-S) code used for the header error control is the 4-ary (10,6) cyclic code. The number of bits per symbol is 4 and the error correction capability is 4/2 = 2 symbols. The code generator is

$$g(x) = (x+ \alpha)(x + \alpha^7)(x + \alpha^8)(x + \alpha^9) \text{ over } GF(2^4),$$

where $\alpha^4+\alpha+1 = 0$, $\alpha^6 = 1100$, $\alpha^7 = 1011$, $\alpha^8 = 0101$, $\alpha^9 = 1010$.

The coded virtual channel data unit (CVCDU) is formed by appending a block of R-S check symbols at the end of the VCDU. The R-S code used for the CVCDU is the 8-ary (255, 223) code. The number of bits per symbol is 8 and the error correction capability is 32/2 = 16 symbols. If the VCDU is larger than the 223 symbols, then the symbol interleaving method should be used. In the symbol interleaving method, several R-S encoders are used in parallel to generate the codeword. The length of the information block suggested is I*223 symbols, where I ranges from 1 to 5. For ATM cells, the number of cells which can be supported by the R-S code with symbol interleaving method is 4, 8, 12, 16 and 21 for the corresponding I.

If the virtual channel data unit is not protected by the R-S code, then a cyclic redundancy code is mandatory. The number of parity bit is 16 and the generator polynomial is

$$g(x) = x^{16} + x^{12} + x^5 + 1.$$

The main function of this code is to detect the errors within the VCDU.

5.4.2.3 Reed-Solomon Code for SVP, SVP Payload, and SVL

A t-error-correcting R-S code (n,k) with symbols from Galois Field GF (2^m) has the following characteristics:

- block length: $n = 2^m - 1$
- number of parity check digits: $n - k = 2t$
- minimum distance: $d_{min} = 2t + 1$

For different size n and different correcting capability t, the number of cells, which one SVP can contain, is provided in Table 5-5.

The approach used in the CCSDS error control code can also be applied to protect the SVP payload and the SVL which contains routing information. From the discussion above, the size of the check symbol for the CVCDU is 32 byte. It is possible that these redundancy parity bits can be put into one cell. This suggests that the SVP (or SVP payload) parity bits can be put into a separate cell so that the Reed-Solomon coding is optional depending on the QOS and GOS. If the R-S code is used, then the SVP transport bandwidth is decreased by one cell for parity bits. This design has the

Table 5-5. Number of ATM Cells Contained in an FEC Code Block

BLOCK LENGTH (n)	SYMBOL SIZE (m)	CORRECTION CAPABILITY (t)	NUMBER OF CELLS
127	7	2	2
		4,8,16, or 32	1
255	8	2,4,8, or 16	4
		32	3
		64	2
511	9	2,4,8 or 16	10
		32	9
		64	8
1023	10	2	24
		4,8, or 16	23
		32	22
		64	21

flexibility in choosing an optimal trade-off between the bandwidth efficiency and the link performance. One example that error control of the SVP payload is necessary is that cells carries important control information such as timing and clock frequency.

The advantages of on-board SVP, SVP Payload, or SVL error correction is that link margin is increased since the uplink impairments are separated from the downlink impairments. One advantage of the using a block code for the SVL is that the SVP synchronization can be easily achieved using the correlation between the SVL and the parity bits, as in the procedure defined in the ATM cell self-delineation algorithm. If the SVL error control code has an burst error correction capability, then hunting of synchronization becomes quicker since the number of slip bits (within the error correction capability) can be detected. Hence the hunting window for the SVL can be moved by the amount of the slippage at a time instead of moving the window bit by bit.

5.4.2.4 STM Virtual Container Loss

In addition to provision of ATM transmissions, another option for the B-SIDN satellite network is to support the SDH signal transmission. The embedded operation and built-on performance monitoring capability of SDH are the advantages of using this approach. However, due to the random errors and burst errors exhibited in the satellite links, the SDH virtual container may be lost. The virtual container of SDH is lost for two reasons: frame synchronization is not maintained and administrative unit (AU) pointer contains error.

The six-byte framing pattern located in the regenerator section overhead is provided in each STS-1 signal for the framing purpose. The format of framing pattern is

A1A1A1A2A2A2, where A1 is 11110110 and A2 is 00101000. If the framing is not found within four consecutive frames, then out of frame occurs. If out of frame occurs, then the reframing process is invoked. The reframing process is accomplished by examining the 8-bit framing block until one of the two framing patterns is identified. Then the other pattern is tracked. During the reframing process, several frames are lost. The in-framer condition is achieved if two consecutive framing patterns have been detected.

The AU pointer points the start of the virtual container. The pointer value is located in H1 and H2. The value indicates the offset in bytes between the pointer and the first byte of the payload. The pointer action is located in the H3, which has three bytes for frequency justification. This AU pointer allows a dynamic alignment of the payload with the fixed position of the AU pointer. The H1 and H2 values are binary numbers ranging from 0 to 782. The overhead (OH) bytes are not counted in the offset. For example, a pointer value 0 means that the payload starts right after the H3 field. An offset of 261 means the payload starts after the K2 filed (see Figure 12). The frequency difference between the frame rate of the overhead and the payload are adjusted by decrementing or incrementing the pointer along with a positive or negative stuff byte. During the normal operation, the pointer value should remain constant for at least three frames. If the payload rate is less than the overhead rate, then the distance between the starting byte of the payload and the pointer is increased. Hence the pointer value should be increased by 1 and a positive stuff byte should be inserted right after the H3 field. The subsequent pointer values contain the new value and the stuff byte is not needed. If the payload rate is greater than the overhead rate, then the distance between the starting byte of the payload and the pointer is decreased. Hence the pointer value should be decreased by 1 and a negative stuff byte should be inserted at the H3 field. The subsequent pointer values contain the new value and the stuff byte is not needed. It is noted that the AU pointer value should not be adjusted for more than 1 byte between two adjacent frames. Hence if the pointer value of the present frame is deviated from the previous pointer value for more than 1, then the AU pointer contains errors. As a result, the SDH payload is lost. If the error causes the AU pointer value to be changed by exactly one compared to the previous one, then the error cannot be detected. If the SDH payload contains circuit slots, the 1 byte error in the AU pointer value either deletes one byte or adds one byte to the payload. If SDH payload contains cells, cells (not the whole payload) are lost since the header of the cell is not in the right boundary. However, the cell delineation algorithm will achieve cell synchronization within several cells time.

To reduce the loss ratio of the SDH virtual container through the satellite links, the proposed scheme is to put FEC block codes for the framing pattern and the AU-4 pointer. The FEC parity bits might be put in the spare field Z1 and Z2 (6 bytes) of the SOH so that the format of the SDH overhead is not be violated. At the receiving earth station, the receiver performs error correction for the framing pattern and the AU-4 pointer first, and then frame synchronization procedure is performed followed by the AU-4 pointer processing.

The prerequisite that cells or circuit slots can be extracted from the SDH cell payload is that the SDH virtual container is not lost. If the virtual container has been

demultiplexed from the SDH frame, then the circuit slots can be extracted. For ATM cells, the next step is to use the H4 offset indicator in the path overhead to extract cells. If H4 offset contains errors, then several cells are lost since the initial cell boundary cannot be found. However the ATM cell self-delineation state machine will find the cell boundary within a few cells period. To improve the error performance of SDH cell payload and to assure the robustness of the cell delineation capability, two schemes can be used. First, an FEC code is provided for the H4 offset. These check bits can be resided in the 3-byte spare fields of the POH, Z3 through Z5. Second, the H4 offset is used for the initial cell delineation. Then the confirmation of the cell boundary is done using the cell HEC self-delineation method. The advantage of the second scheme is that the confirmation can be done within a few cell periods instead of a few SDH frame periods.

5.5 Impact of Satellite Delay

Because of the window concept used in data communication protocols, the satellite delay affects the throughput performance and the automatic repeat request (ARQ) protocol efficiency. In this section, corrective measures of the above problems are proposed. To improve the performance of data communication protocols, an increase of the window size and the time-out period are two effective approaches. To improve the efficiencies of ARQ protocols, modification of the retransmission strategies is necessary. The approach used for improving the performances of X.25/X.75 protocols through the satellite links have been analyzed by COMSAT [5-4]. Some of the results, which will be discussed below, are also applicable to the ATM protocols.

5.5.1 X.25/X.75

CCITT Recommendations X.25/X.75 are international standardized protocols used for data communication networks. X.25 specifies the physical, link, and network layer protocols for communication between the data terminal equipment (DTE) in the end-user and data circuit-terminating equipment (DCE) in the network. X.75 specifies the protocols for communication between two networks. In other words, X.25 specifies the protocols used in the user network interface (UNI) and X.75 specifies the protocols used in the network node interface (NNI).

The physical layer in X.25 is specified by X.21, which defines the use of a duplex point-to-point synchronous circuit.

The link layer in X.25 provides reliable transmission of information between two adjacent nodes. The basic unit of information at this layer is called a frame, and there are three different types of frames. Information frames are used to send data. Supervisory frames are used for acknowledgement (or negative acknowledgement of data) or to slow the sender for flow control. Unnumbered frames are for control purposes. There are fields associated with these frames: control (such as sequence numbers and acknowledgements), data, and checksum. Note that sequence numbers are used to order frames, to request retransmission if an out-of-sequence frame is

received, or to acknowledge receipt of in-sequence frames. The parameters in this layer depend on the characteristics and quality of the link. One such parameter is the retransmission time defined as the time-out value for the sending side to wait for an acknowledgement.

The network layer of X.25 manages connection between two DTEs. Two forms of connections are supported: switched virtual circuit and permanent virtual circuits. Switched virtual circuits need to have call setup and call teardown procedures. A permanent virtual circuit is like a leased line in which the DTEs can send data to each other at any time without a call setup sequence.

End-to-end connections of virtual circuits are established using the following procedure. The DTE forms a Call Request packet and sends it to the DCE. The network sends the packet to the destination DCE and then to the destination DTE. If the DTE likes to receive the call, it sends back a Call Accepted packets. When the DTE receives the Call Accepted packet, the virtual circuit has been established.

During this procedure, some parameter values may be negotiated. The first one is the packet size, i.e., the maximum number of bytes in a data packet. The second one is the window size, i.e., the number of sequential packets that can be transmitted without waiting for an outstanding acknowledgement.

After the call setup phase, the DTEs use the full-duplex channel to exchange data packets. When either side likes to terminate the connection, it sends a Clear Request packet to the other side. The other side then sends a Clear Confirmation packet back as an acknowledgement to the originating DTE.

X.75 is used to specify the protocols to set up a virtual circuit between DTEs connected to different networks. X.75 is almost identical to X.25.

As mentioned above, several parameters within X.25/X.75 protocols affect the performances through the satellite link, which include the window size, retransmission timer value, and packet size. For efficient operation, appropriate parameters should be set by network managers. The packet formats of X.25/X.75 are shown in Figure 5-27.

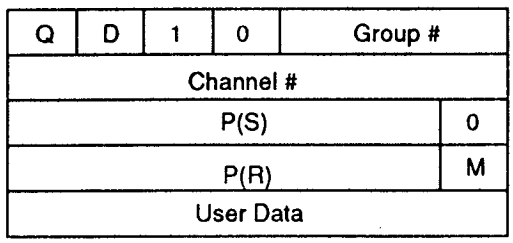

Q: types of data
D: local or end-to-end significance for the ACK
P(S): packet send sequence number
P(R): packet receive sequence number
M: more bit

Figure 5-27. X.25/X.75 Data Packet Format

The window size specifies the maximum number of packets the source can send before receiving an acknowledgement from the destination. The timer value specifies the time-out period for receiving an acknowledgement. The determination of these parameters is based on the characteristics of the satellite links such as propagation delay, transmission rate, random bit errors, and channel fading errors. The automatic repeat request (ARQ) within the link layer protocol also affects the performance through the satellite links. It has been shown that "selective repeat" (positive ACK) or "selective reject" (negative ACK) retransmission schemes outperforms the "go-back-N" scheme. Figure 5-28 shows the comparison of go-back-N and selective retransmission schemes operated in a 45 Mbit/s satellite link and terrestrial link [5-4]. Due to the long satellite propagation delay, the performance of the go-back-N scheme degrades quickly in the satellite environment while the performance of the selective retransmission scheme is not affected. The positive ACK is defined as that the packet has been received at the receiver and no error is detected. In this case, the sender can delete the packet stored in the buffer. The negative ACK is defined as that the packet has been received in the receiver and errors are detected. In this case, the sender has to retransmit the packet. This automatic repeat request (ARQ) protocols will be explained and discussed latter.

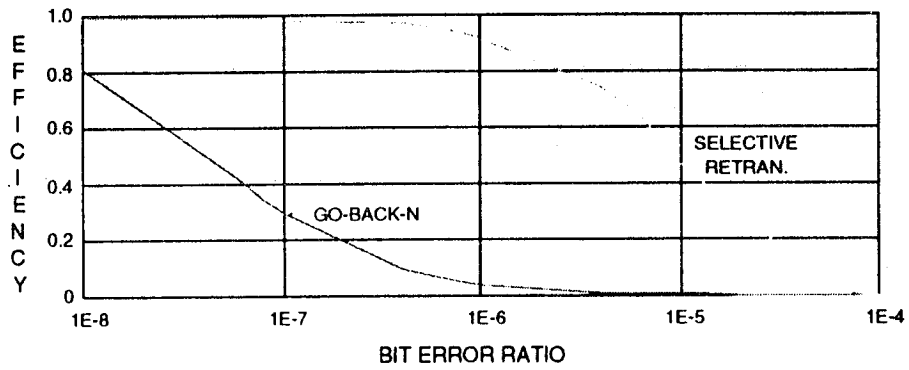

Efficiency of a 45 Mbit/s Satellite Link for a Frame Size of 4096 bytes

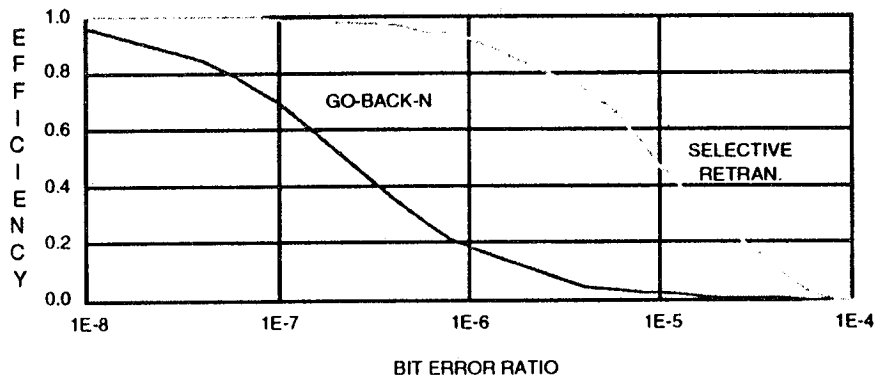

Efficiency of a 45 Mbit/s Terrestrial Link for a Frame Size of 4096 Bytes

Figure 5-28. Performance Comparison of Go-Back-N and Selective Retransmission Schemes

5.5.2 Frame Relay

Frame relay is a new ISDN frame mode bearer service (FMBS) using a new data link protocol and packet-switching technique. This bearer service provides the order preserving bidirectional transfer of service data units (layer 2 frames) from one UNI to another UNI. Routing through the network is based on the data link connection identifier (DLCI). DLCI only has a local significance, which means retranslation of the identifier is required at each switch node. Frame relay is an end-user service and it is a specification for ISDN UNI (primary rate interface). The feature of frame relay is not to use the complete link layer protocol as used in X.25. X.25 was designed to provide error-free packet delivery service in high error-rate links environment. Frame relay is designed to provide frame delivery service in low error-rate links environment; hence, it can eliminate many functions provided by X.25. The difference of frame relay protocol and X.25 protocol is shown in Figure 5-29. Frame relay network simplifies the switch node processing by moving the flow control and error recovery to the end-users in order to minimize processing overhead and maximize data transfer. Frame relay multiplexing is done in the link layer and the signal is out-of-band while X.25 multiplexing is performed at the packet layer and the signal is in-band. These differences result in the following advantages of frame relay over X.25. First, the frame can be operated in high speed up to 2.048 Mbit/s. Secondly, the nodal transit delay within the frame relay network as compared with the X.25 network is significantly minimized.

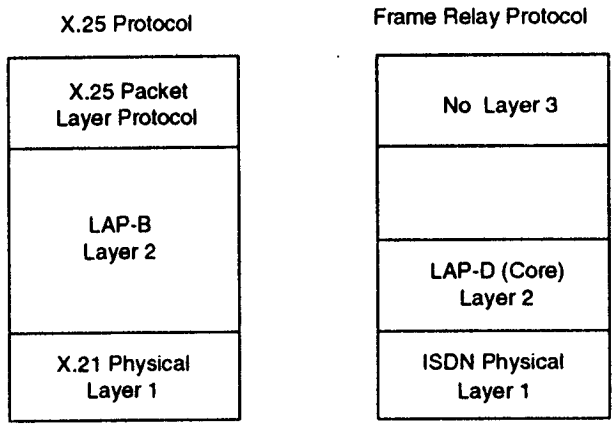

Figure 5-29. Difference Between X.25 and Frame Relay Protocols

The link layer can be divided into two sublayers: the core function and the data link control function. The core function includes the following:

- frame delimiting, frame alignment, and transparency
- frame multiplexing/demultiplexing using the address field

- inspection of the frame to ensure it consists of an integer number of
- bytes prior to zero bit insertion or following zero bit extraction
- verification of the frame size
- detection of transmission errors.

The data link control function includes such as flow control and error recovery. Frame relay protocol within a network is based on the core function of the ISDN D-channel Link Access Procedure (LAPD) The end user employs both core function and data link control function. Frame relay uses the D-channel for signaling and the D- or B-channel for information transfer.

The frame structure as shown in Figure 5-30 consists of the opening flag, address field, control field, information field, frame check sequence (FCS), and the closing flag. The address field consists of the address field extension bit (E/A), command/response field bit (C/R) and data link connection identifier (DLCI). The DLCI field is used when frame relay is operated on B channels. DLCI provides the multiplexing function. If frame relay is used for D channel, then DLCI0 is used for the service access point identifier (SAPI). The SAPI is to identify different services. For example, SAPI = 1 is for frame relay and SAPI = 16 is for X.25. The DLCI1 is used for the terminal endpoint identifier (TEI). The control field identifies the type of frames. There are three kinds of frames: numbered information transfer (I format), supervisory functions (S format) and unnumbered information transfers (U format), which are essentially the same as the high-level data link control (HDLC) protocol. At each network node, frame relay protocol only performs routing based on the data link connection identifier (DLCI) and ignores the control field. Flow control and error recovery is executed on an end-to-end basis.

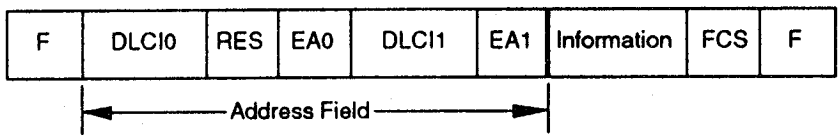

F: flag
DLCI0: high order bits of data link connection identifier
DLCI1: low order bits of data link connection identifier
RES: reserved bit for command/response
EA0, EA1: address field extension bits
FCS: frame check sequence

Figure 5-30. Frame Relay/Frame Switching Frame Format

According to the frame format, the maximum number (k) of sequentially numbered I frames that may be outstanding at any given instant shall not exceed 127. The default value for a 64-kbit/s channel is 7. For satellite applications, during the call setup time, the value of k should be negotiated to use a large value to improve the performance. As shown in [5-5], the value of k should be 40 for a 128-byte frame size. Otherwise, the

transmitter at the earth station has to constantly wait for acknowledgements if the window size is not large enough. (Frame switching protocol is based on both the core function and the data link control function. The effect of the satellite delay on frame switching is basically the same as the frame relay; hence, frame switching protocol is not discussed here.)

As discussed in [5-5], the frame relay congestion control can be achieved using two methods: congestion-avoidance and congestion recovery. The congestion-avoidance mechanism is to use the explicit congestion notification to control the rate of the transmitter at the user side when the network is in the on-set of congestion. The explicit congestion notification scheme is to send an explicit congestion notification to the source and/or the destination to throttle the information rate, depending on whether the transmitter is source-controlled or destination-controlled. The effectiveness of the explicit notification scheme due to the long satellite delay depends on the transmission speed. For a high speed satellite network, the explicit scheme cannot effectively represent the queueing status since many packets have been transmitted or received during the long propagation delay.

If throttling of the transmitter rate is not effective and the network is forced to drop frames. There is an option to use a discard eligibility (DE) indicator bit for low-priority frames (if the bit is set to 1) to relieve network congestion. In case of network congestion, the low-priority frame will be discarded first.

The recovery scheme is used when the network suffers severe congestion. Congestion recovery is to dynamically adjust the window size of the number of unacknowledged outstanding frames based on either the implicit scheme: expired timers, detection of frame loss, or the explicit notification schemes. In general, the window size is reduced by a preset factor and then the window size is incrementted till its original size.

Note that the error performance through the satellite link has to be ensured so that the QOS of frame relay is not affected. This subject is discussed in the "BER and QOS" section.

A general comparison among X.25, Frame Relay, and ATM technologies is provided in Table 5-6 [5-6].

Table 5-6. Comparison of X.25, Frame Relay, and ATM for Various Network Functions

FUNCTION	X.25	FRAME RELAY	ATM
Network Delay	larger	low	low
Error Tolerance	very good	poor	poor
Error Processing	link by link	end points	end points
Link Speed	low	medium	high
Types of Traffic (handling capability)	data	voice, data	voice, data, video

5.5.3 ATM Adaptation Layer Services

The ATM adaptation layer (AAL) supports four classes of services according to the following parameters (see Figure 5-31): time relationship between source and destination, bit rate (constant or variable), and connection mode (connection oriented or connectionless). These four classes are:

- Class 1: circuit emulation
- Class 2: variable bit rate video
- Class 3: connection oriented data
- Class 4: connectionless data transfer

The services which are affected by the satellite delay, and the necessary corrective measures for these services are described as follows.

5.5.3.1 AAL Class 1 Circuit Emulation

The AAL Class 1 circuit emulation services are characterized by timing relationship between source and destination, constant bit rate, and connection-oriented mode. Real time services such as video and telephony which generate a periodic cell stream require time transparency of the signal through the satellite B-ISDN. For circuit switching, time transparency of the signal is achieved using the periodic slot concept. For ATM switching, routing is based on the label (ATM header) concept. Although all cells within the same connection can be delivered to the destination, the periodic nature of the circuit traffic is distorted by the stochastic queueing delay over the satellite B-ISDN. Time transparency of the signal within the satellite B-ISDN is achieved by putting constraint on the queueing delay so that the delay developed within the satellite network is bounded. If the queueing delay is bounded, then at the receiving end, the transmit timing frequency can be recovered and an elastic buffer can reconstruct the signal and suppress the stochastic delay. Timing recovery schemes at the end receiver is not discussed here since the satellite B-ISDN is performed as a transport network. However, the methods used to improve the performance of timing recovery are introduced.

There are two major factors which affect the performance of timing recovery at the receiving end: the cell loss ratio and cell delay jitter. For information cells, the cell loss has a minimal effect on the receiving end if the sequence number (SN) at the adaptation layer (AAL) is utilized. The lost cell can be detected and appropriate actions can be taken. Since the SN only has 4 bits, loss of more than 16 cells with the same VPI/VCI should be avoided so that undetected loss of cells does not occur. This puts a constraint on the number of cells with the same VPI/VCI that the satellite virtual packet (SVP) can carry, which is discussed in the "ATM Cell Packetization Issues and Alternatives". For

	Circuit Emulation	Variable Bit Rate Video	Connection Oriented Data	Connectionless Data Transfer
Timing Relationship Between Source And Destination	Required		Not Required	
Bit Rate	Constant	Variable		
Connection Mode	Connection-Oriented			Connectionless

Figure 5-31. AAL Service Classification

control cells, the cell loss could mean loss of timing information or loss of frame synchronization. Hence, improvement of cell loss ratio for the control cell is necessary.

There are three ways to reduce the cell delay jitter (CDJ) and cell loss ratio (CLR). The first one is using the priority concept. This scheme has been fully studied in Appendix C. In a sense the high-priority cells are transmitted and served first in a switching node, and the results show that the improvement of CLR and CDJ is significant. However, the high-priority cells might monopolize the buffer or the transmission channel. The second scheme is to use an error control scheme. The possible error control schemes are discussed in the "Protection of ATM Cell Header from Burst Error" section. Nevertheless, this scheme only improves the CLR. The third one is to reserve cells in the SVP for the control cells as in the reserved slot concept used in circuit switching; the result is that the cell delay jitter is bounded. This is an effective scheme as long as the cells do come in to the earth station periodically so that the reserved cells in the SVP are not wasted.

5.5.3.2 AAL Classes 3 and 4 Assured Operation

AAL classes 3 and 4 services are essentially corresponding to the two services provided by the network layer: the virtual circuit and datagram services. The corrective measures of the ARQ protocol used in these services through the satellite links are are discussed below.

The ARQ protocol considered here is only the continuous ARQ, i.e., the stop-and-wait ARQ is not considered. In this case, the transmitter sends the packets to the receiver continuously and receives the ACK packets continuously. There are two ways of implementing the continuous ARQ: the go-back-N and the selective repeat. For the go-back-N ARQ, the transmitter will back up to the packet in error and resend the packet plus the following packets after an negative-ACK packet has been received. Hence, the receiver only has to store one packet at a time. For the selective repeat ARQ, the transmitter only sends those packets which are acknowledged negatively. As mentioned before, because of the satellite long end-to-end delay, the go-back-N scheme degrades the throughput and the delay. Hence the selective repeat is always proposed for the satellite applications. However, using the selective repeat scheme requires more buffering at the receiver and hardware at both transmitter and receiver than the go-back-N scheme.

There are several objectives of implementing an efficient ARQ protocol.

- Excessive retransmission avoidance
- Single mode receiver operation
- Round trip delay and network delay variance insensitivity
- Maximum window size

- High throughput and low delay

Note the excessive retransmission is not necessary a drawback. This is because if the satellite link performs poorly under a fading condition, then excessive retransmission is the alternative to improve the link performance. In this case, the excessive retransmission becomes the repetition strategy, which has been proposed to combat high channel bit errors.

One of the recent proposals submitted to T1 meeting and ANSI meeting from COMSAT [5-4] [5-7] implements an ARQ protocol to accomplish the above objectives. This protocol is based on two main procedures:

a. The transmitter sends a poll packet to poll the receiver periodically for state information. Every poll packet contains a sequence number.

b. The receiver generates a response packet whenever a poll packet is received. The response packet contains the following information: the received poll packet sequence number, the data packet sequence numbers (those have been received), and the flow control window size.

The advantage of this protocol is that only the transmitter needs to have a timer to generate the poll packet, and the receiver does not have any timer. However, whenever a timer is used in the protocol, the insensitivity of the protocol to round trip delay and variance does not exist anymore.

Another advantage of this protocol is that there is no excessive retransmission. The avoidance of excessive retransmission is achieved using the poll sequence number. This is explained using an example below. Assume, at the transmitter, a data packet with sequence number n has just been sent to the receiver and a poll packet with sequence number m has also just been sent out after that. Assume the transmitter receives a response packet which states that data packet n has not been received. Two cases can happen: If the response packet is generated in response to the poll packet m, this is the normal situation and the transmitter should retransmit the data packet n. If the response packet is generated in response to another poll packet, for example, the previous one $m-1$, then the transmitter should not retransmit data packet n since the poll packet $m-1$ was generated before the data packet n was generated. In order to distinguish the timing relationship between the poll packets and the data packets, the data packets stored in the transmitter (waiting for an ACK) all keep the current value of the poll packet sequence number. Hence, by comparing the poll packet sequence number in the response packet and the poll packet sequence number in the data packet, the transmitter only retransmits those data packets whose poll numbers are smaller than the poll number in the response packet.

The disadvantages of this scheme are that the poll packets waste the bandwidth of the link and complicate the protocol design. The loss of several poll packets makes the response packet grow very large. For the ATM application, the response packet has to be sent in two or more cells.

The flow control scheme used in the COMSAT proposal [5-4] is the explicit feedback control scheme. This scheme is achieved using a credit field in the status packet based on the queue status of the receiver. The credit field is basically the window size for the transmitter. The receiver either decreases or increases the window size for the transmitter based on the congestion level. In this protocol, the range of window size is between 0 and (maximum sequence number - 1). However, this scheme is effective in low speed applications. For high speed ATM application, more study is necessary to perceive the effectiveness.

The COMSAT ARQ protocol throughput and delay performance curves compared with high-level data link control (HDLC) and ideal selective retransmission protocols are shown in Figure 5-32 [5-7], where w is the buffer size provided at the receiver side.

5.5.4 SDH Data Communication Channel Delay Compensation

The synchronous digital hierarchy (SDH) contains standardized overhead bits for operation maintenance, communication, and performance monitoring functions. The delay compensation technique is proposed for one of the communication fields,i.e., the data communication channel.

There are twelve bytes available for the data communication channel (DCC). Four bytes (192 kbit/s) are used for the regenerator section (alarm, maintenance, control, monitor, and administration between section regenerator terminating equipments). Eight bytes (576 K bit/s) are used for the multiplexer section (communication between cross-connects and digital loop carrier elements). To terminate the DCC for delay compensation, the earth station must have packet-handling capability. According to the CCITT Recommendation and Bellcore document, DCC at the section overhead of SDH is based on the high-level data link control (HDLC) protocol. The balanced link access procedure (LAPB) of X.25 is based on the HDLC protocol. The LAPD protocol is based on LAPB and HDLC. The HDLC frame structure is shown in Figure 5-33. The frame has the following fields: beginning flag (8 bits), address, control (8 or 16 bits), data, frame check sequence, and end flag. The control field defines three types of frames: Information frames, Supervisory frames, and Unnumbered frames. The delay compensation techniques used for X.25 and frame relay can also be applied to HDLC, which will not be repeated here.

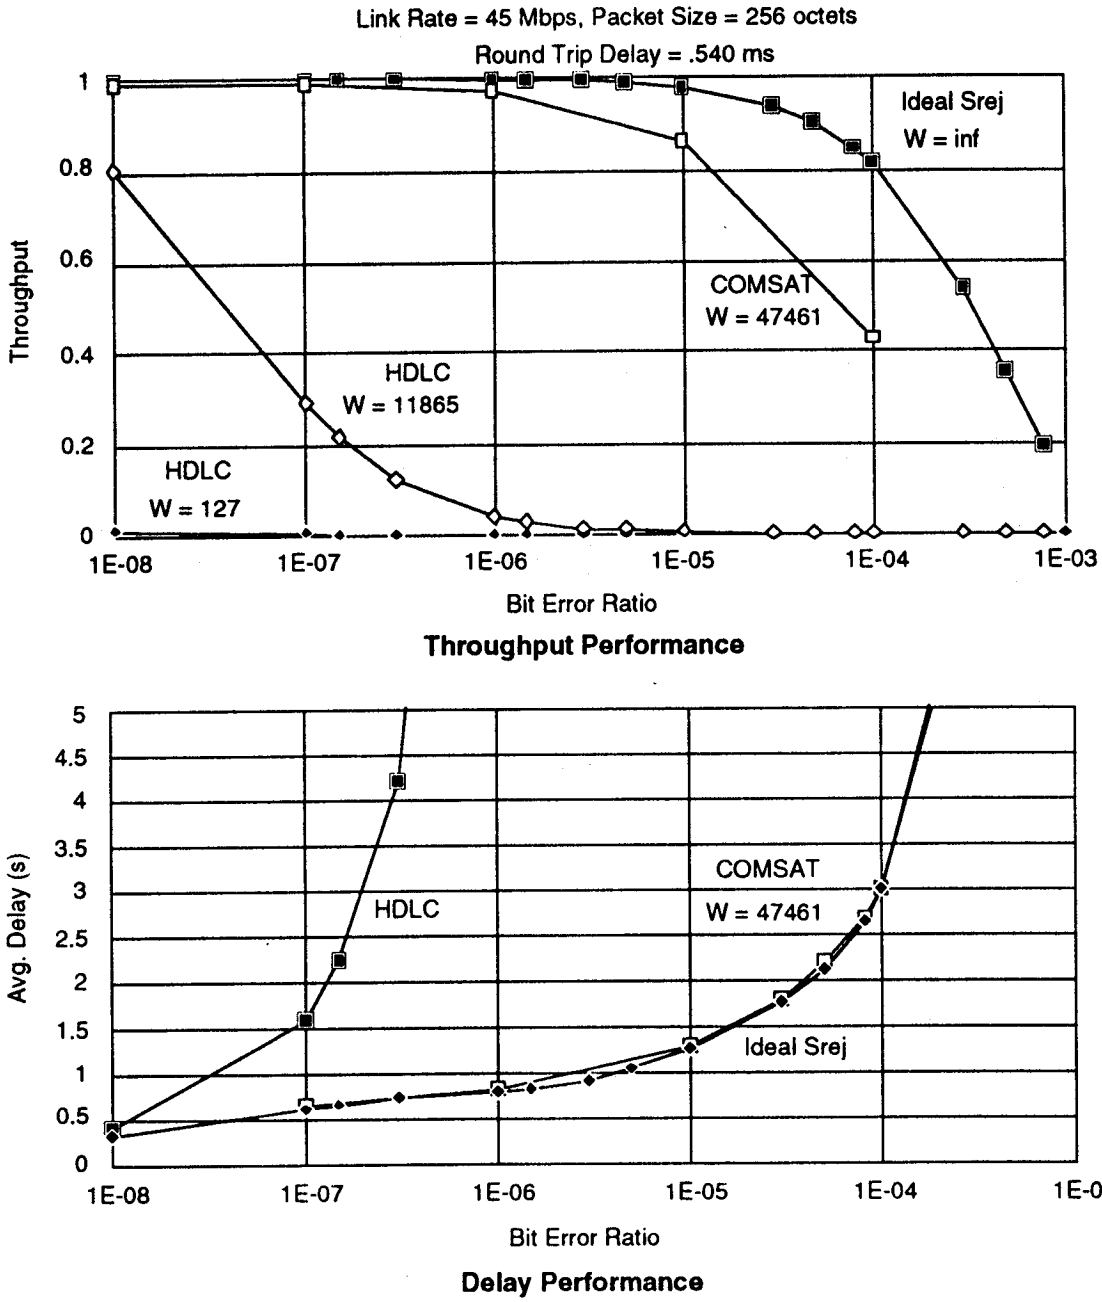

Figure 5-32. Performance Comparison of COMSAT Selective Retransmission Protocol and HDLC Protocol in the Satellite Environment

FCS: frame check sequence

Figure 5-33. HDLC Frame Format

5.6 References

[5-1] Y. Yamamoto and T. Wright, "Error Performance in Evolving Digital Networks Including ISDNs," IEEE Communications Magazine, pp. 12-18, April 1989.

[5-2] CCITT Study Group XVIII - Report R4 1989.

[5-3] A. E. Eckberg, B. T. Doshi, and R. Zoccolillo, "Controlling Congestion in B-ISDN/ATM: Issues and Strategies", IEEE Communications Magazine, pp. 64-70, Sep. 1991.

[5-4] S. Jarrar and A. Agarwal, "Elements of Procedure for AAL Type 4 Assured Operation with Selective Retransmission," T1S1.2/91-045.

[5-5] D. M. Chitre, "On-Board Functions for the Next Generation Services and Networks," Space Communications Technology Conference: On-Board Processing and Switching, 1991.

[5-6] B. Bhushan, "Frame Relay, Fast Packet, and Packet Switching - Convergence or Coexistness," Data Communications, pp. 51-54, Dec. 1990.

[5-7] A. K. Agarwal, "Contribution to ANSI X3S3.4 Comments and Proposals on Data Link Control Protocol," COMSAT Technical Note NTTR-081, April 1991.

Section 6

Fast Packet Switch Architectures

In a self-routing fast packet switch, the path between each source-destination (input-output) pair is determined by the binary representation of a destination address. By prepending a path header, a packet can be sent to the destination without the need of a central controller. Many architectures have been proposed to implement a self-routing fast packet switch. In this section, the design principles of different fast packet switching architectures for point-to-point and multicast connections are described.

General characteristics of fast packet switching are:

a. It provides a high-speed packet transfer path between any input line and any output line.

b. Packet transfer protocol and distributed routing functions are performed by a hardware-oriented self-routing architecture.

c. Packet transfer is performed according to the routing information contained in the packet header with a high degree of parallelism.

d. The switching architecture supports a wide range of services with various bit rates.

e. The switch does not process the information in the packets, nor does it retransmit the packets that are lost due to transmission errors or switch buffer overflow.

f. The switch includes an output contention control mechanism.

g. The switch employs a building block architecture to allow for modular expansion.

High throughput, minimum delay, and high reliability are the three main goals in the design of a fast packet switch. These design issues will be further discussed in various sections of this report.

This section is organized as follows. The first subsection introduces the most commonly used 2 x 2 self-routing switching element, banyan network, 2 x 2 sorting element, and

batcher sorting network. Also described in this subsection are blocking type switch structures and packet transfer protocols used in the fast packet switch.

Subsection 6.2 presents an approach to classify different point-to-point fast packet switches and discusses methods for resolving blocking and output contention problems and for increasing switch performance.

Subsection 6.3 investigates design techniques for multicast fast packet switches and addresses advantages and disadvantage of each technique.

Subsection 6.4 presents the strategies of using priority control to guarantee the qaulity of service (QOS) for different classes of services in a fast packet switch. Implementation issues regarding to priority control are also described.

6.1 Fast Packet Switching

6.1.1 Self-Routing Switching Element and Banyan Network

A banyan network is in the category of multistage interconnection networks [6-1]. It can be constructed using any size of switching elements. If the size of the switching elements in the banyan network is a D x D switching element, the number of switching elements at each stages is N/D, and the number of stages is Log_D N. The banyan network is a unique path network in which there is only one path between any input-output pair. The banyan network is topologically equivalent to many other multistage interconnection networks, such as baseline, omega, flip and shuffle networks.

A 2 x 2 switching element has four allowed states: straight, exchange, lower broadcast, and upper broadcast (see Figure 6-1). For a point-to-point banyan network, only the straight and exchange states are used, and each switching element needs to check only one bit of routing tag to route the packet. The lower broadcast and upper broadcast states are the basic principles that a banyan network can perform a multicast function; the multicast banyan network will be discussed in a latter subsection. If the corresponding routing bit is zero, the data will be sent to the upper link of the that element; otherwise, to the lower link (see Figure 6-2). For easy hardware implementation, the switching element at stage 1 checks bit 1 of the routing tag. The switching element at stage k checks bit k of the routing tag, where $1 \leq k \leq \text{Log}_2$ N. Following this bit representation, the leftmost bit of the routing tag is the least significant bit and the rightmost bit is the most significant bit. An example of routing two packets in an 8 x 8 banyan network is shown in Figure 6-3.

6.1.2 Sorting Element and Sorting Network

The batcher sorting network is in the category of bitonic sorting networks which produce sorted outputs from circular bitonic inputs [6-2]. A bitonic list is a list which monotonically increases from the beginning to the i-th element and then monotonically decreases from the i-th element to the end. A circular bitonic list is created from joining

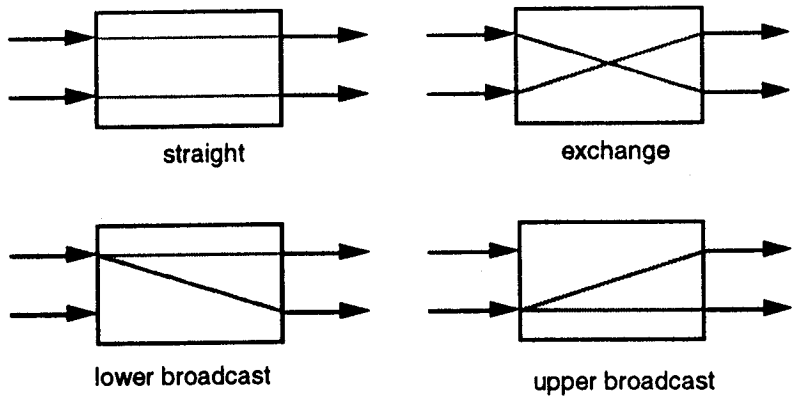

Figure 6-1. Four States of a 2 x 2 Self-Routing Switching Element

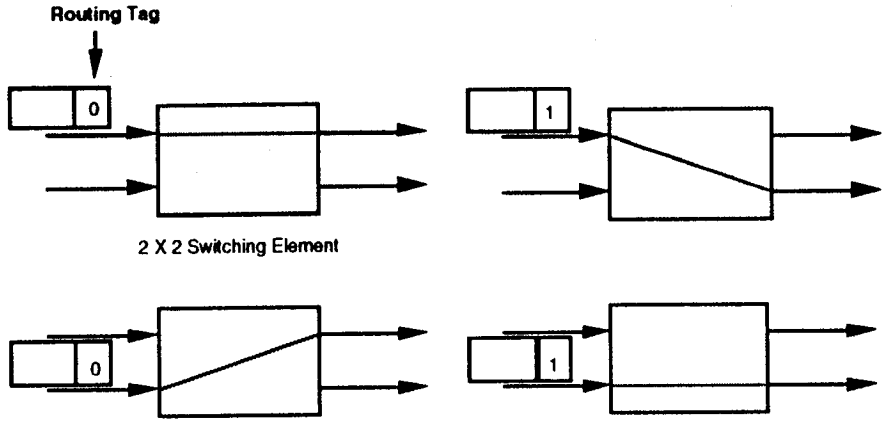

Figure 6-2. Destination Tag Routing Algorithm

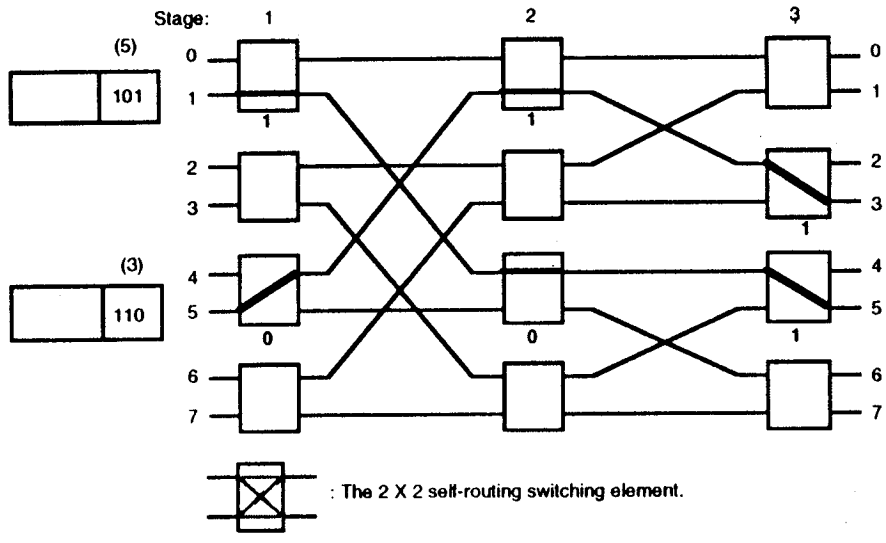

Figure 6-3. The 8 x 8 Banyan Network

the beginning and the end of a bitonic list, and then breaking the circular structure into a linear structure at any desired point.

The sorting network has a similar property as a banyan network, i.e., a large network is constructed recursively using a smaller network. An N x N batcher sorting network has $\frac{1}{2}$ Log_2 N (Log_2 N + 1) stages and each stages consists of $\frac{N}{2}$ sorting elements. The function of a sorting element is shown in Figure 6-4. An 8 x 8 sorting network is shown in Figure 6-5.

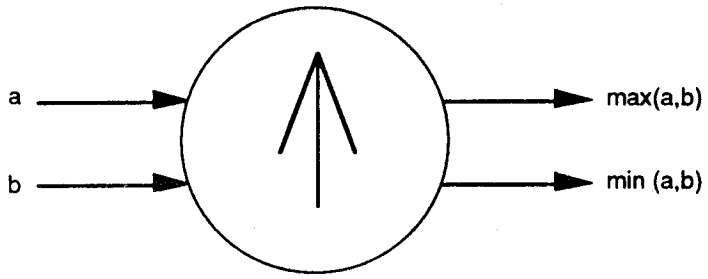

Figure 6-4. A 2 x 2 Sorting Element

6.1.3 Blocking Types

In a fast packet switch, there are three possible types of blocking a packet may encounter: internal blocking, output contention, and head of line blocking. Not all blocking types exist in every switching architecture. However, the output blocking problem is an unavoidable situation of packet switching; therefore, it exists in every switching architecture.

6.1.3.1 Internal Blocking (IB)

The first type of blocking shown in Figure 6-6 is internal blocking which occurs within the switching fabric. Basically, this happens because a link only can serve one packet at a time. If two packets contend for the same link, then only one is allowed to use the link and the other one is blocked. There are many ways to tackle this internal blocking problem; these methods will be discussed in subsection 2.

6.1.3.2 Output Contention (OC)

The second type of blocking is output contention shown in Figure 6-7 which occurs at the output ports of the switch. Due to the statistical nature of packet switched data, several packets from different output ports may be destined to the same output port at the same time. The key difference between internal blocking and output contention is that output contention is an unavoidable situation in the packet switching environment. But with a careful design of the switching architecture, internal blocking can be

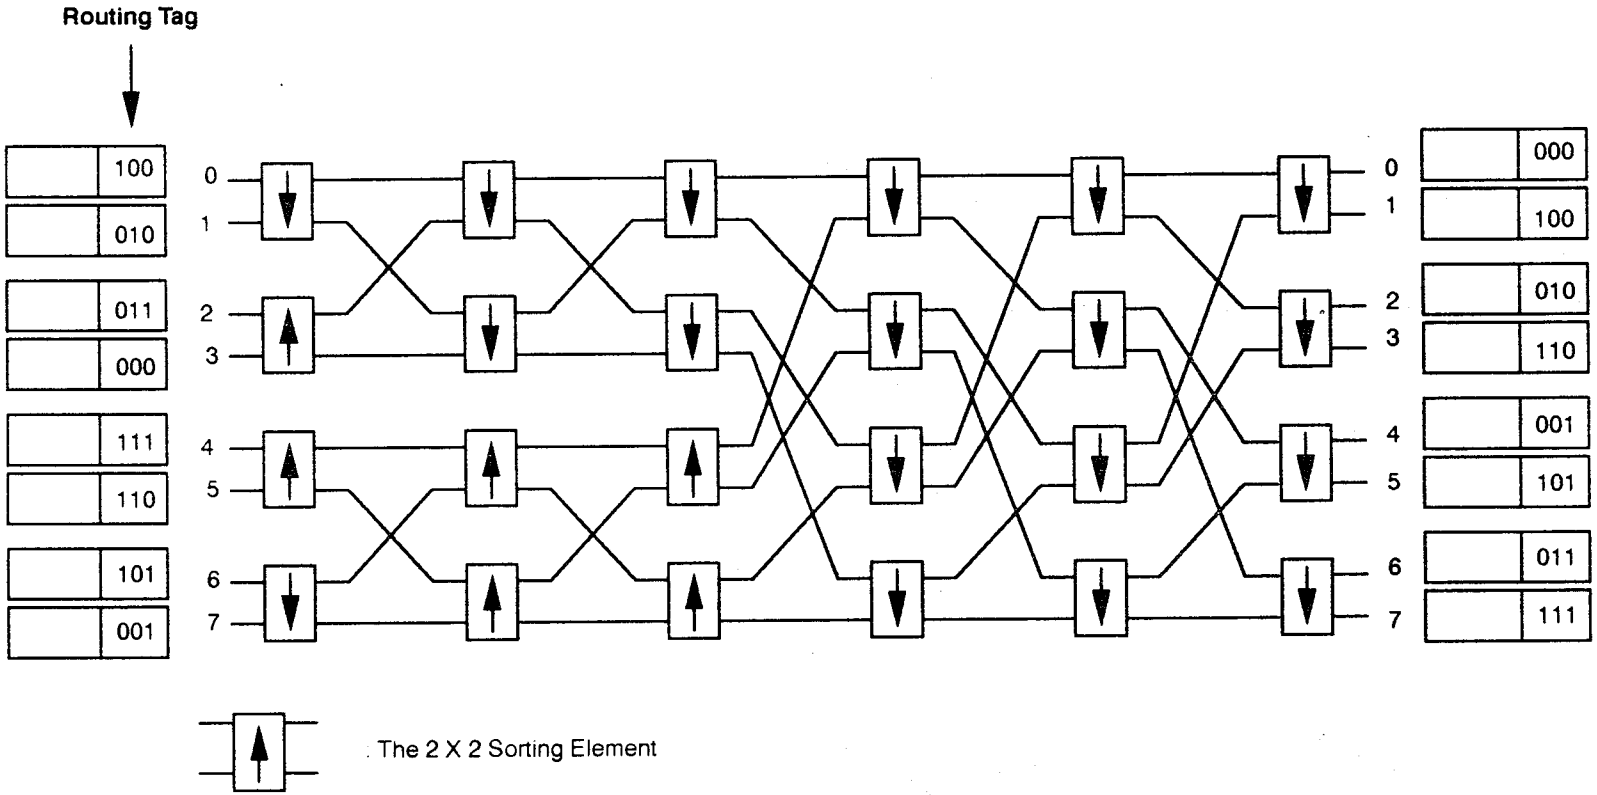

Figure 6-5. 8 x 8 Sorting Network

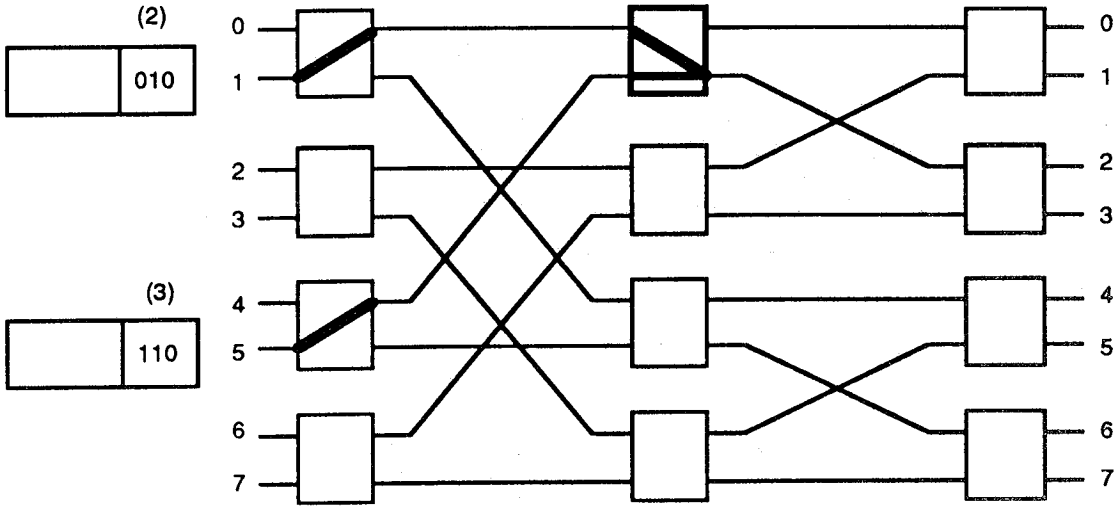

Figure 6-6. Internal Blocking

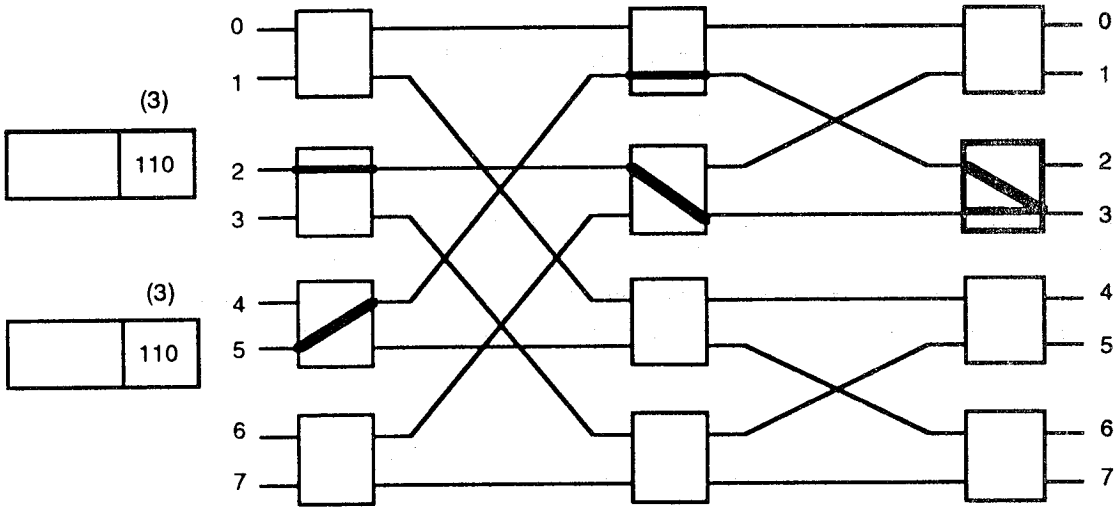

Figure 6-7. Output Contention

avoided. There are many ways of resolving the output contention problem; these will be discussed in subsection 2.

6.1.3.3 Head of Line Blocking (HOL)

The third type of blocking is head of line blocking, shown in Figure 6-8, which occurs at the input port queue or at the switching element's buffer within the switch. Evidently if the switch does not employ the input buffering method or switching element buffering method, there is no HOL blocking. This blocking is a side effect of the results of the previous two blocking types. Assume one packet at the head of queue cannot be transmitted due to internal blocking or output contention. Then this blocked packet

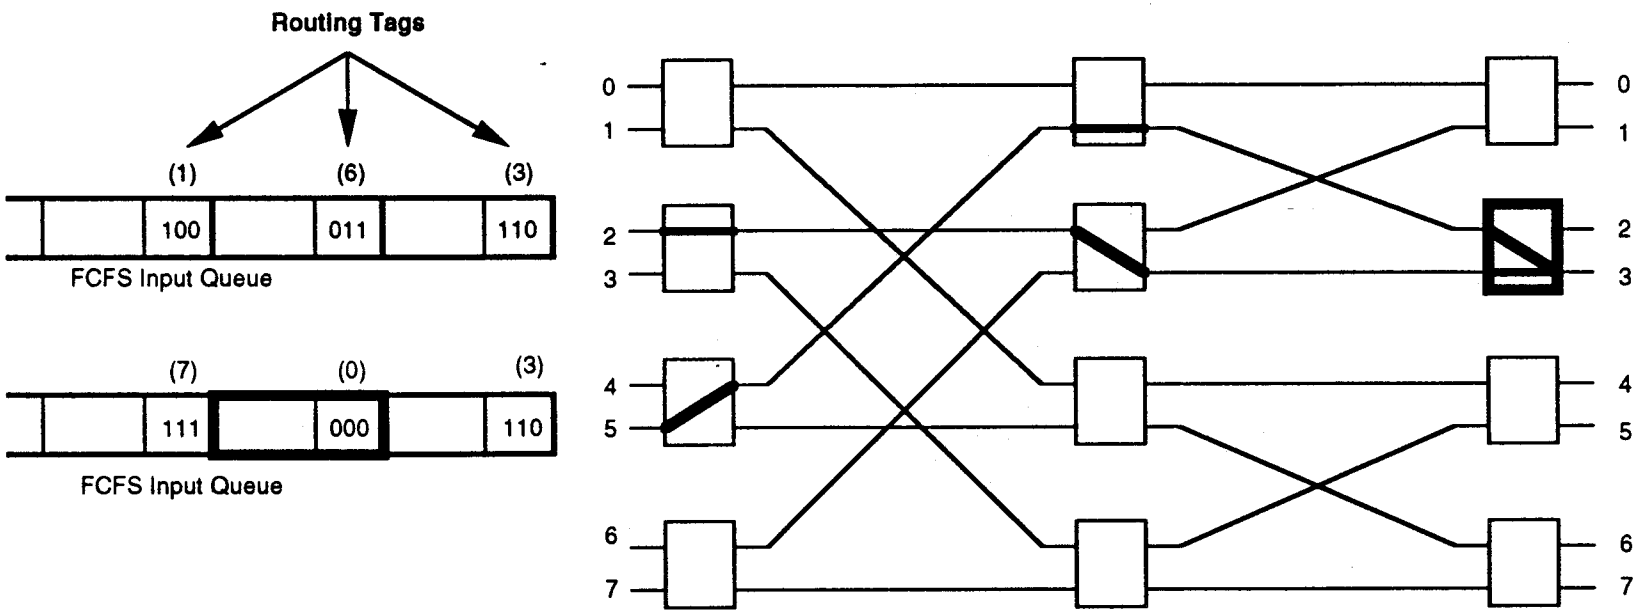

Figure 6-8. Head of Line Blocking

hinders the delivery of the next packet in the queue due to the first come first serve (FCFS) nature of the queue, even though the next packet can be transmitted to the destination without any blocking. As shown in Figure 6-8, the packet with destination 000 at input port 5 is blocked by the packet with destination 110, although the packet with routing tag 000 can be transmitted to the destination without any blocking at the current slot time.

6.1.4 Packet Transfer Protocol

Since a packet might encounter blocking in a packet switch and hence lost, the switch has to use the store-and-forward (input buffering) or forward-and-store (output buffering) property to guarantee the delivery of the packet or reduce the blocking problems so that the packet loss ratio is within a tolerable range.

There are three approaches to implement the store-and-forward protocol. The first one tries to set up a path between the input and the destined output; it has a setup phase and followed by a forwarding phase if the setup phase is successful. The second one has a packet forwarding phase and followed by a retransmission phase if the forwarding phase is not successful. The third one is almost the same as the first one except the protocol is used between every two stages of the switch such as a buffered banyan switch.

The store-and-forward protocol mentioned above uses the input buffering approach. If no input buffering is used in the packet transfer protocol, then either output buffering has to be used to implement the forward-and-store protocol or special operations are necessary to impose on the switch to guarantee that the packet loss ratio satisfies the requirement. This will be discussed in the last two subsections.

6.1.4.1 The Setup Phase + Forwarding Phase Protocol

In this protocol, a packet will not be sent from the input port buffer unless a path between the input-and-output pair has been reserved for the packet. There are two methods to implement this path reservation depending on whether the switching fabric is nonblocking or not. If the switching fabric is blocking, then the setup packet method is used; if the switching fabric is nonblocking, then either the setup packet scheme or the output port reservation scheme can be used.

(a) The Setup Packet Scheme for Blocking Switching Fabric

The procedure of this protocol is shown in Figure 6-9. The input port sends a small setup packet and attempts to reserve a path between the input port and the destined output port. The setup packet consists of only the routing tag. If the output port receives the setup packet, the output port sends an acknowledgement (ACK) back to the original input port. After the path has been successfully set up, the input port can release the packet and send it to the output port. If the input port does not receive an ACK within three routing tag's unit time (two tag's time for the round trip delay time

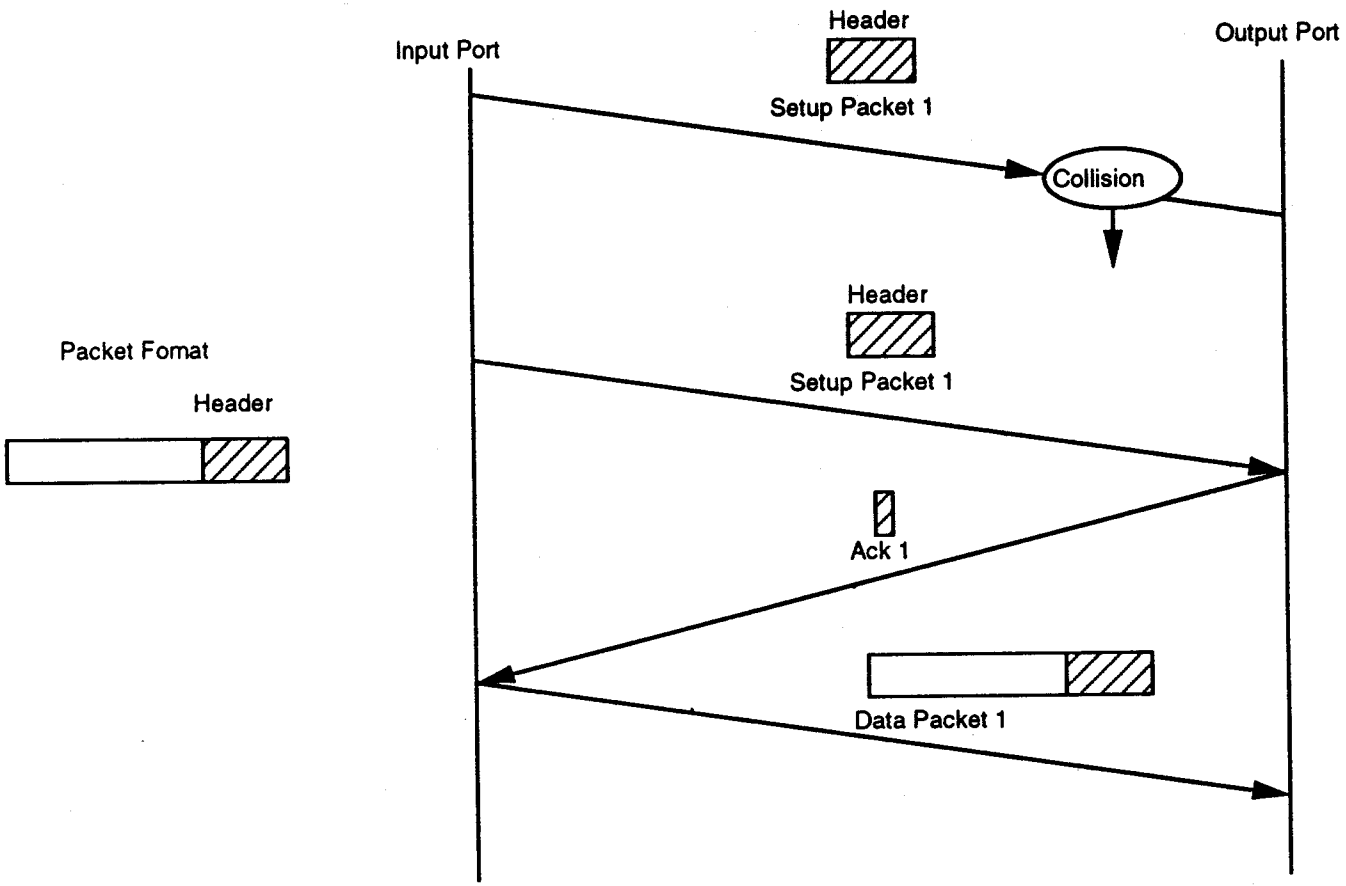

- Resolve switching fabric internal blocking and output port contention at the same time
- It can be used for blocking and nonblocking switching fabric
- The switching element is bidirectional

Figure 6-9. Setup Phase + Forwarding Phase Protocol

and one tag's unit time for the transmission time), then the input port sends the setup packet again, and the whole procedure is repeated. From the above discussion, the switch needs to have bidirectional connection capability. This method can operate in the minislot mode, where the length of the minislot is the setup time (three routing tag's unit time).

(b) The Reservation Scheme for the Nonblocking Switching Fabric

Since the switching fabric is nonblocking (there is no internal blocking), if the output port is reserved, the path has been reserved. One of the methods to reserve the output port is to use the ring reservation scheme [6-3].

Basically, the ring reservation scheme uses the token ring principle to resolve output contention. The input ring connects all the input ports of the switch, and a stream of tokens, where each token is for each output port, are sent through the input ports (see Figure 6-10). The function of the ring is to perform the output reservation for each input port. At the beginning of every slot time, the ring module sends a stream of tokens and passes these tokens to all the input ports. The input port searches the right token according to the destination routing tag of the current packet. If the token for the corresponding routing tag is on the stream, then the token is removed so that no other input port can transmit a packet to the same output port at the same slot time. After the token stream has passed through all the input ports, the input ports that have reserved a token can transmit the packet at the beginning of the next slot time. In implementation, only one bit is necessary for one token. Value 1 represents there is a token and value 0 no token. To assure fairness among the input ports of accessing the tokens, several ways can be employed. The first is at different slot time, the stream will be started at different input port. The second is to send this stream form the beginning of the input ports and from the end the input ports alternatively.

6.1.4.2 The Forwarding Phase + Retransmission Phase Protocol

The procedure of this protocol is shown in Figure 6-11. This procedure is only suitable for the slotted mode operation. First the input port stores a copy of the packet in the buffer. Then the input port sends the whole packet to the destined output port. When the output port receives a packet, an ACK is sent back to the originating input port. When the input port receives an ACK, the input port discards the packet and processes the next arriving packet. If the ACK does not come back within two routing tag's time plus one packet length's time, then the packet is sent again and the whole procedure repeats. The switch needs to have bidirectional connection capability. The length of the slot of is the packet length's time plus two routing tag's time.

6.1.4.3 Store-and-Forward Protocol Between Stages

This protocol is specially used for the multistage interconnection networks with buffered switching elements. In this protocol, a packet in one switching element will not be sent to the next stage's switching element unless there is no blocking at the output link and

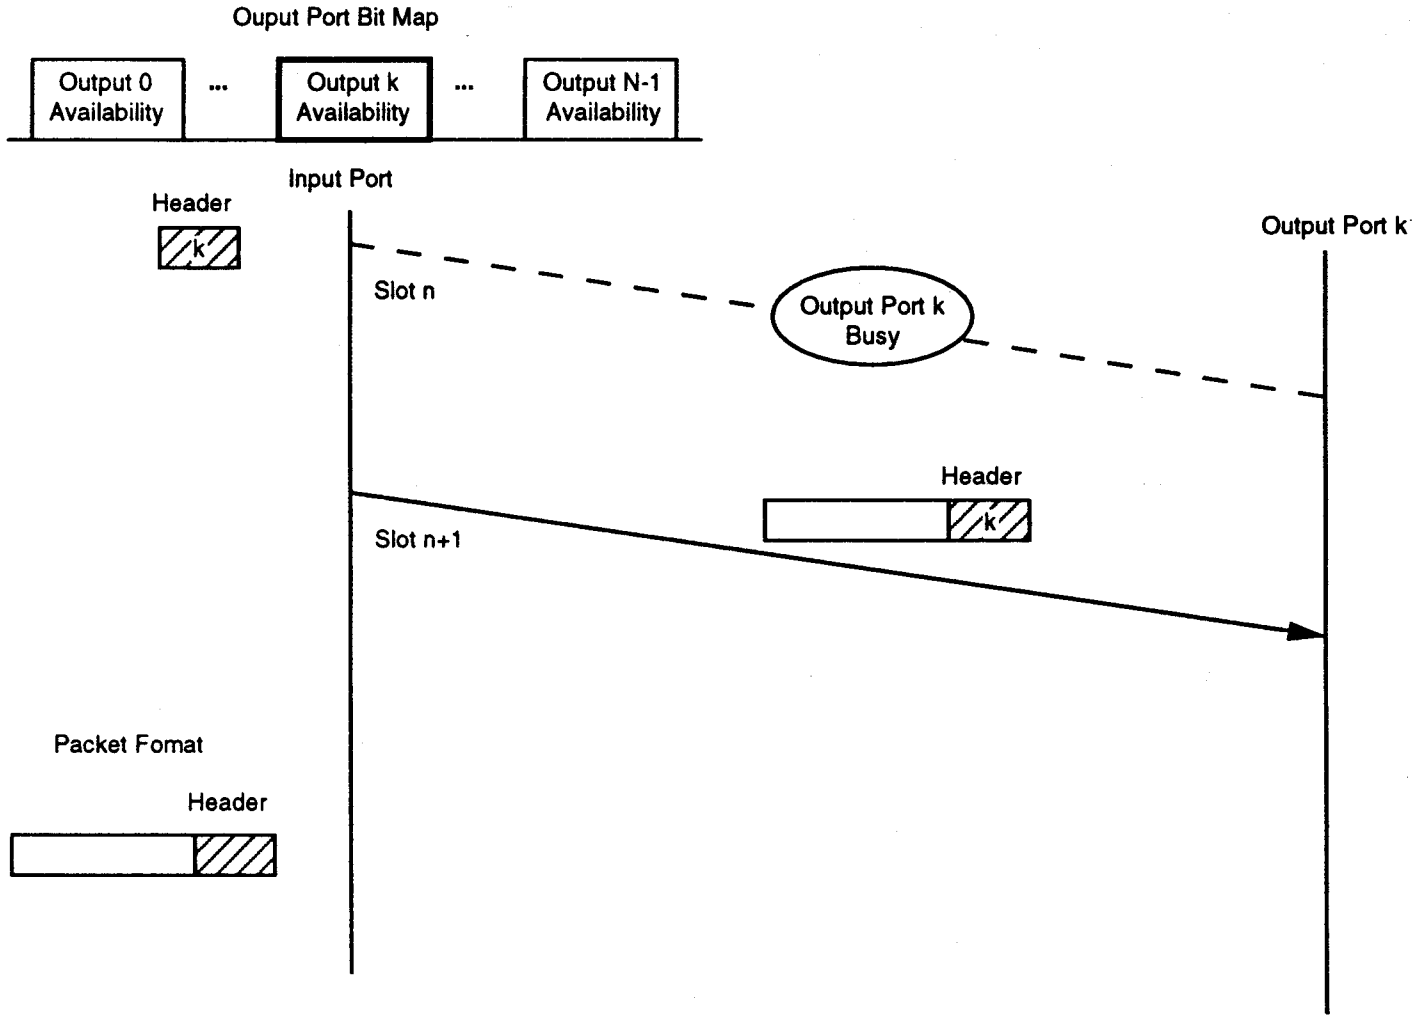

- Resolve output port contention
- It can be used only for nonblocking switching fabric and input buffering

Figure 6-10. Output Port Reservation Scheme

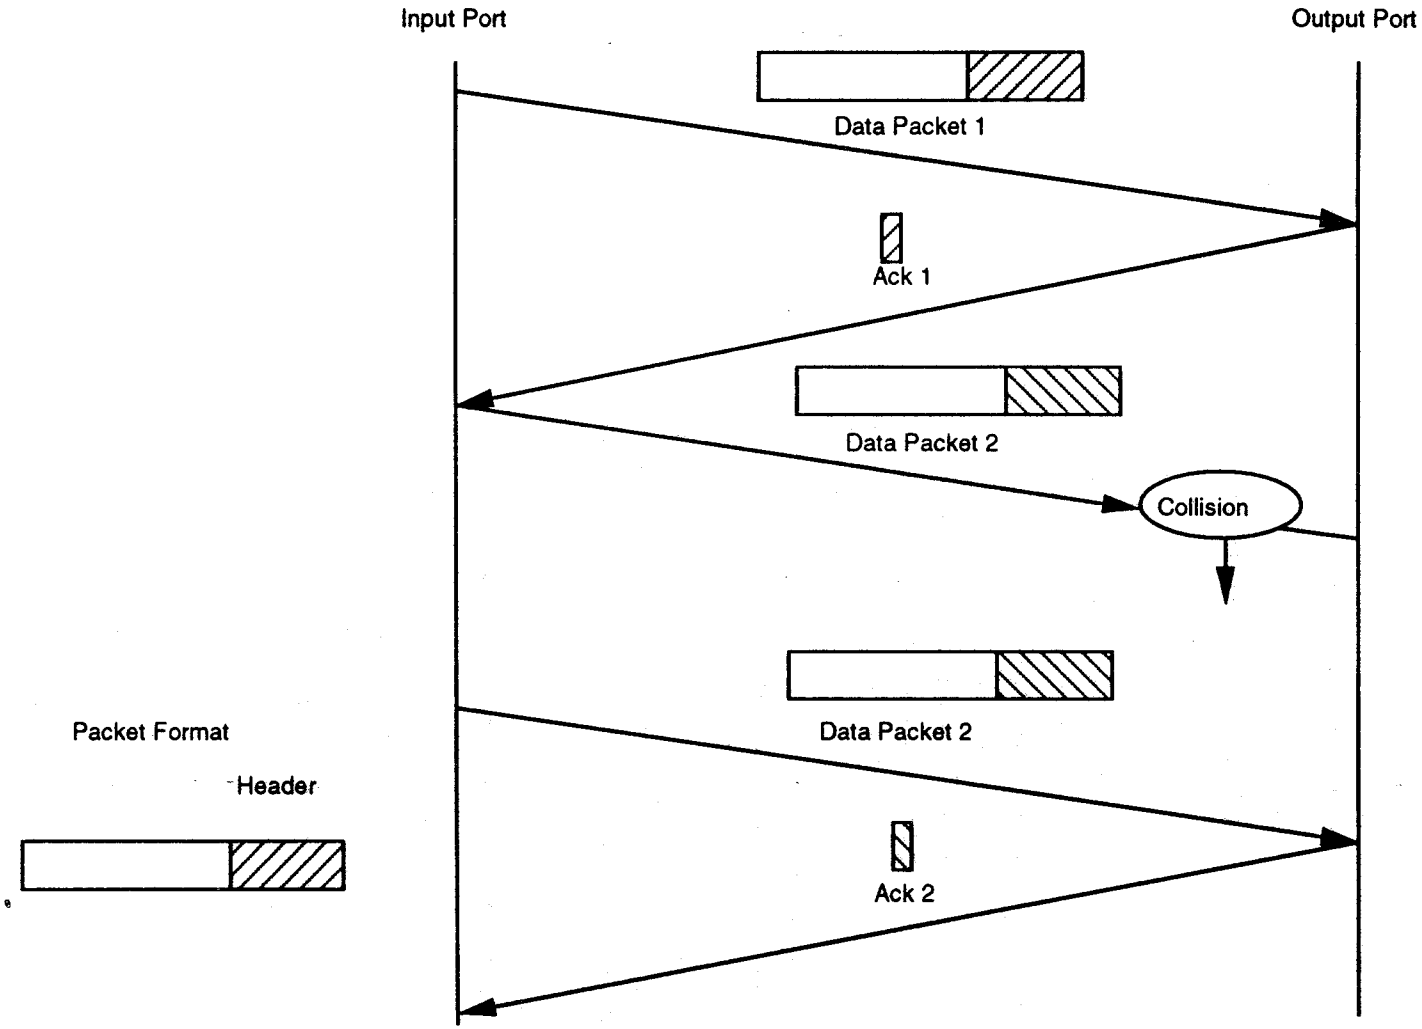

Figure 6-11. Forwarding Phase + Retransmission Phase Protocol

there is an available buffer space in the next stage's switching element. In a sense, the scheme is similar to the first scheme: the setup phase + forwarding phase protocol. The difference is in this scheme the store-and-forward has to be performed between every two stages in the switch, and, in terms of implementation, control signals are used for handshaking between stages instead of using setup packets.

6.1.4.4 Forward-and-Store Protocol with Output Queueing

In this protocol, the buffering is performed at the output ports not at the input ports. The forwarding phase is used to transmit the arriving packets to different output ports since the packets are not stored at the input ports. In this scenario, more than one packets are allowed to arrive to one output port at the same time; hence, multiple receivers and output buffering are required to store the packets (see Figure 6-12). Since there is no path setup phase or retransmission phase, a packet will be lost if the packet encounters any blocking in the switch. Evidently the blocking switching fabric is not suitable for this protocol. The feasible candidate is a nonblocking switching fabric with output buffering. Since the switching fabric is nonblocking, there is no internal blocking problem and since the buffers are located at the output ports, there is no head of line blocking. To resolve output port contention, the output ports need to have the capability of receiving more than one packets and the switching fabric has the capability of delivering more than one packets to the output port within one slot time. The are two basic design rules for the switching fabric to have the capability of delivering more than one packets to the output port. The first one is to provide multiple paths between each input port and output port pair. The second one is to operate the switch at a higher speed. In either case, there is a multiple-input port buffer at each output port so that multiple packets can be stored in the output buffers if they all arrive within the same slot time.

6.1.4.5 Rerouting Protocol

In this protocol, the operation of the switching element is modified so that when two packets request the same link, one is granted and the other one is routed to the wrong link. The packet which is routed to the wrong link has a chance to retry the transfer at next slot using another switching fabric. This will be discussed in the cascaded banyan network case. The other scenario is specific to the sorted-banyan-based network. After the sorting procedure, all the packets destined to the same output port are adjacent to each other. One packet is granted for transmission and the other packets are routed to the reentry input ports for retry at the next slot time. In this case, only one switching fabric is required for both new and retried packets.

6.2 Classification of Fast Packet Switches

The fast packet switching architectures were extensively surveyed by Tobagi [6-3] and Ahmadi [6-4]. The classification methods used in both papers are summarized as follows.

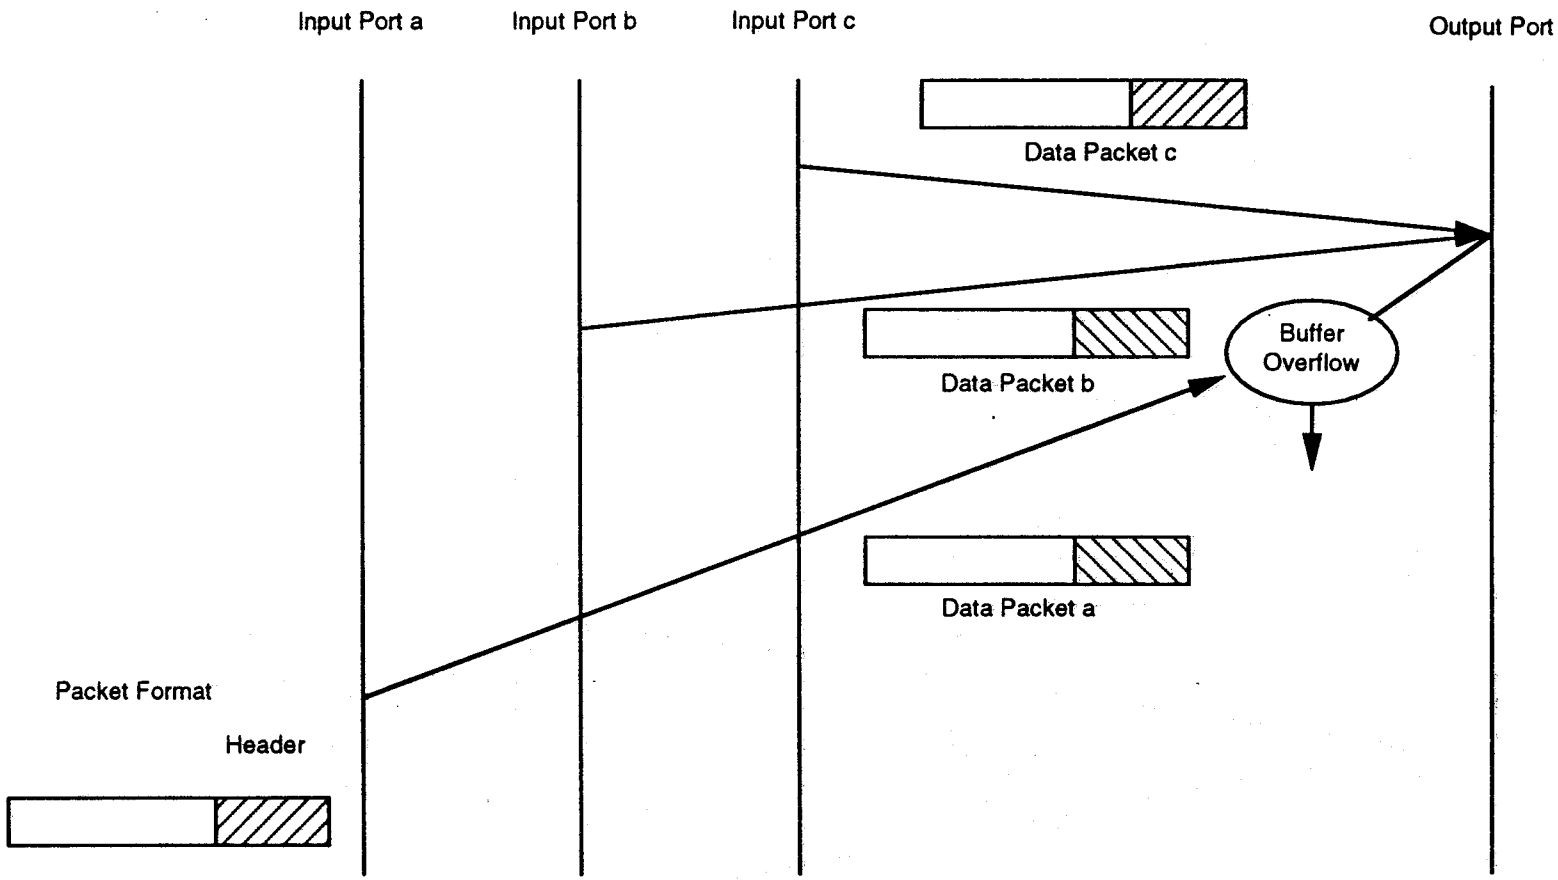

Figure 6-12. Forward-and-Store Protocol with Output Queueing

(a) Tobagi's Classification

Tobagi [6-3] classified the fast packet switching architectures into three categories: the shared-memory, the shared-medium, and the space-division types.

(b) Ahmadi's Classification

Ahmadi and Denzel [6-4] classified the fast packet switching architectures into the following categories: the unbuffered banyan and buffered banyan-based fabrics, the sort-banyan-based fabrics, the fabrics with disjoint-path topology and output queueing, the crossbar-based fabrics, the time division fabrics with common packet memory, and the fabrics with shared medium. In this context, Tobagi's classification has a broader sense than the approach used by Ahmadi and Denzel. For example, the unbuffered banyan and buffered banyan-based fabrics, the sort-banyan-based fabrics, the fabrics with disjoint-path topology and output queueing, the crossbar-based fabrics all can be categorized under the space-division type.

In the discussion following, the classification of the fast packet switching architectures follows the Tobagi's approach (see Figure 6-13). Under space-division switching network category, there are two main classes to be considered based on the topology: the banyan type network and the N^2 disjoint-path switching network. There are three classes of switching networks based on the buffering approach: input buffering, internal buffering, and output buffering. There are two classes of switching networks based on the blocking property of the switching fabric.

At each architecture, the blocking type and the performance improvement method, the buffering scheme, and the packet transfer protocol are discussed and examined.

6.2.1 Space Division Network

Since a packet may be blocked within the packet switch, the throughput and transfer delay performance degrade. Many schemes were proposed to improve the switch performance by easing the blocking problems. Note that to effectively bring up the switch performance, the method should tackle the three types of blocking all together so that no one becomes the bottleneck of the switch performance.

6.2.1.1 Banyan Type Switch

Three main classes are considered in this subsection: the blocking buffered banyan switching fabric (internal buffering), blocking unbuffered banyan switching fabric (input buffering), and nonblocking unbuffered banyan switching fabric (input buffering or/and output buffering).

Due to the internal blocking and output contention problems of a blocking network or the output contention problem of a nonblocking network, a switch needs to store the packets in the buffer. Mainly there are two places where buffering can be implemented: the switching fabric level buffering or switching element level buffering. The switching

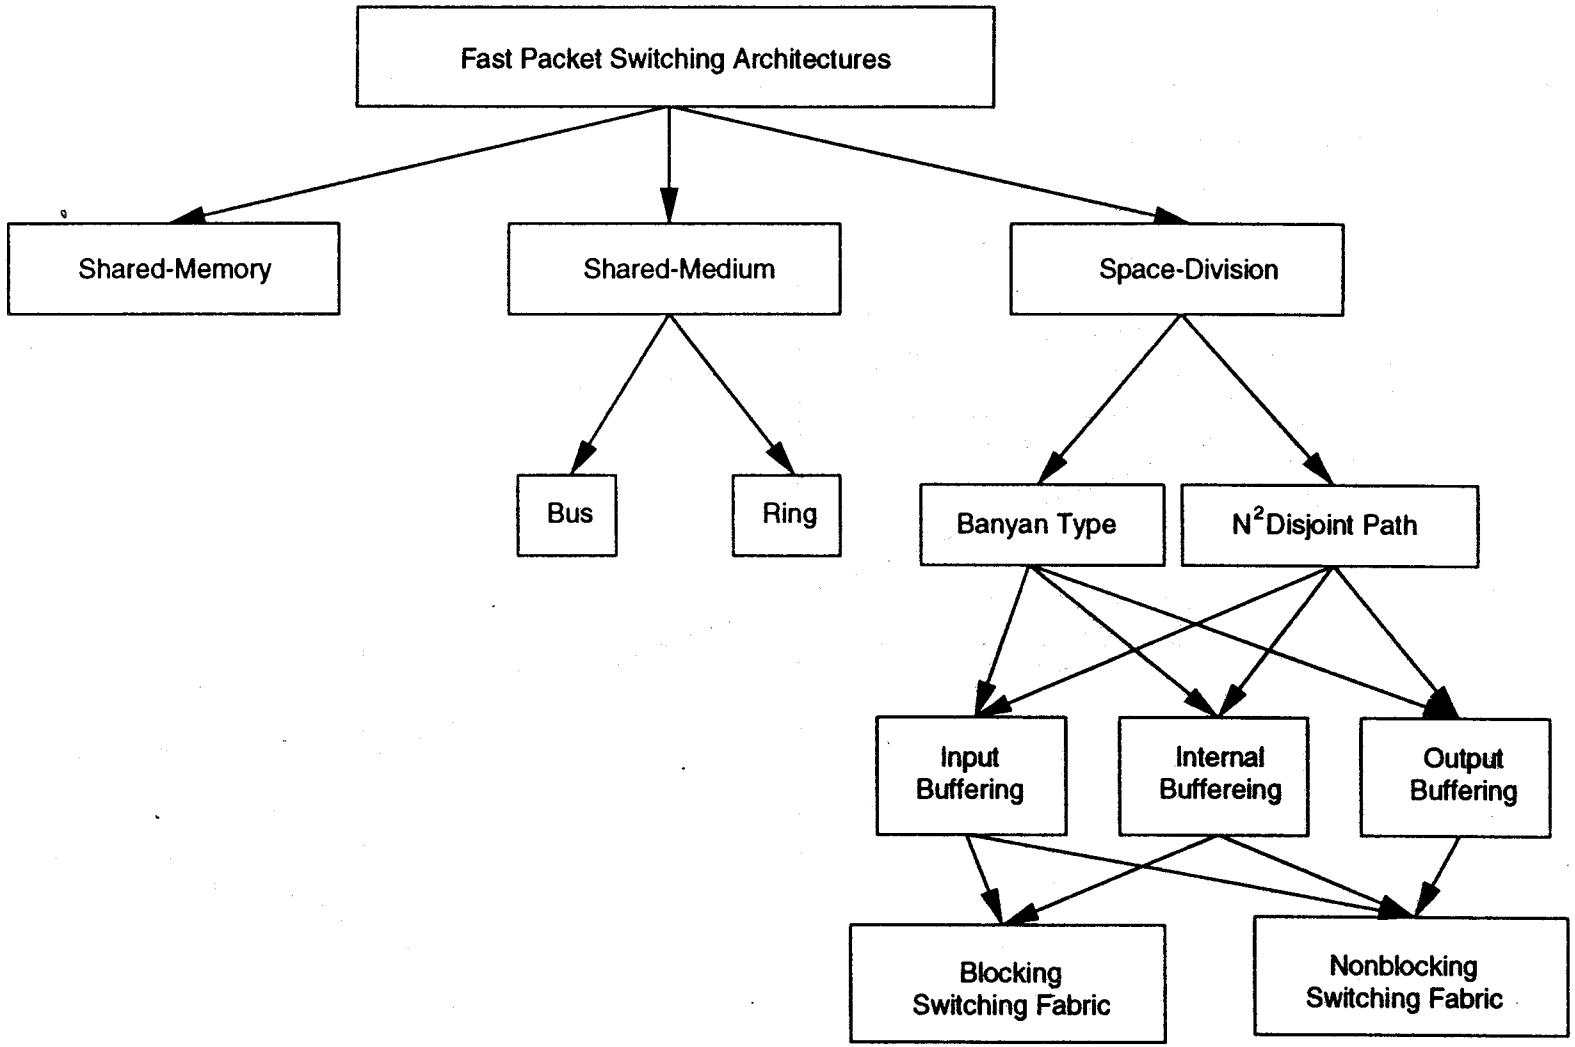

Figure 6-13. Fast Packet Switching Architectures Classification

element level buffering approach will be discussed first, followed by the switching fabric level buffering approach.

6.2.1.1.1 Blocking Buffered Banyan Switching Fabric with Internal Buffering

In this approach, the buffers are implemented within the switching element (see Figure 6-14). In general, the throughput of this approach is higher than that of the unbuffered switching fabric with input buffering approach. However, it has the following disadvantages. The first is the long transfer delay and large delay jitter due to buffering at every stage. The second is the hardware complexity of each switching element is higher. The third is fault-detection of the switching fabric is difficult. The fourth is that a packet will be transmitted out of sequence if alternate paths exist in the switching fabric.

The methods used to improve the switch throughput include the following:

- increased switch speed [IB,OC]
- number of banyan networks in parallel (IB,OC)
- switching element with a larger size (IB)
- distributive (randomization) network (IB)
- non-FIFO queue at each switching element (HOL)
- a separate queue for each output link at each switching element (HOL)

(a) Increased Switch Speed

This approach is to operate the switch at a higher speed than the input lines so that the internal blocking and output contention problem can be relieved. Assume that the link speed is S_l, the switch speed is S_f, and the packet size is N_p, then link slot time is N_p/S_l and the switch slot time is N_p/S_f. The effect of increasing switch speed is that $N_p/S_f < N_p/S_l$. Hence, the switch can process more than N packets within one link slot time, where N is the size of the switch.

(b) Multiple Switching Fabrics in Parallel

By using multiple banyan networks in parallel, the utilization of each switching fabric is reduced to ρ/p, where ρ is the utilization of each incoming line and p is the number of copies stacked in parallel. Since the utilization is reduced, the internal blocking problem is also reduced. This approach will be discussed in more detail in the unbuffered switching fabric subsection.

(c) Larger Switching Element

As mentioned before, a banyan network can be constructed using a larger switching element. If a banyan network is built on D x D switching elements, the routing tag has

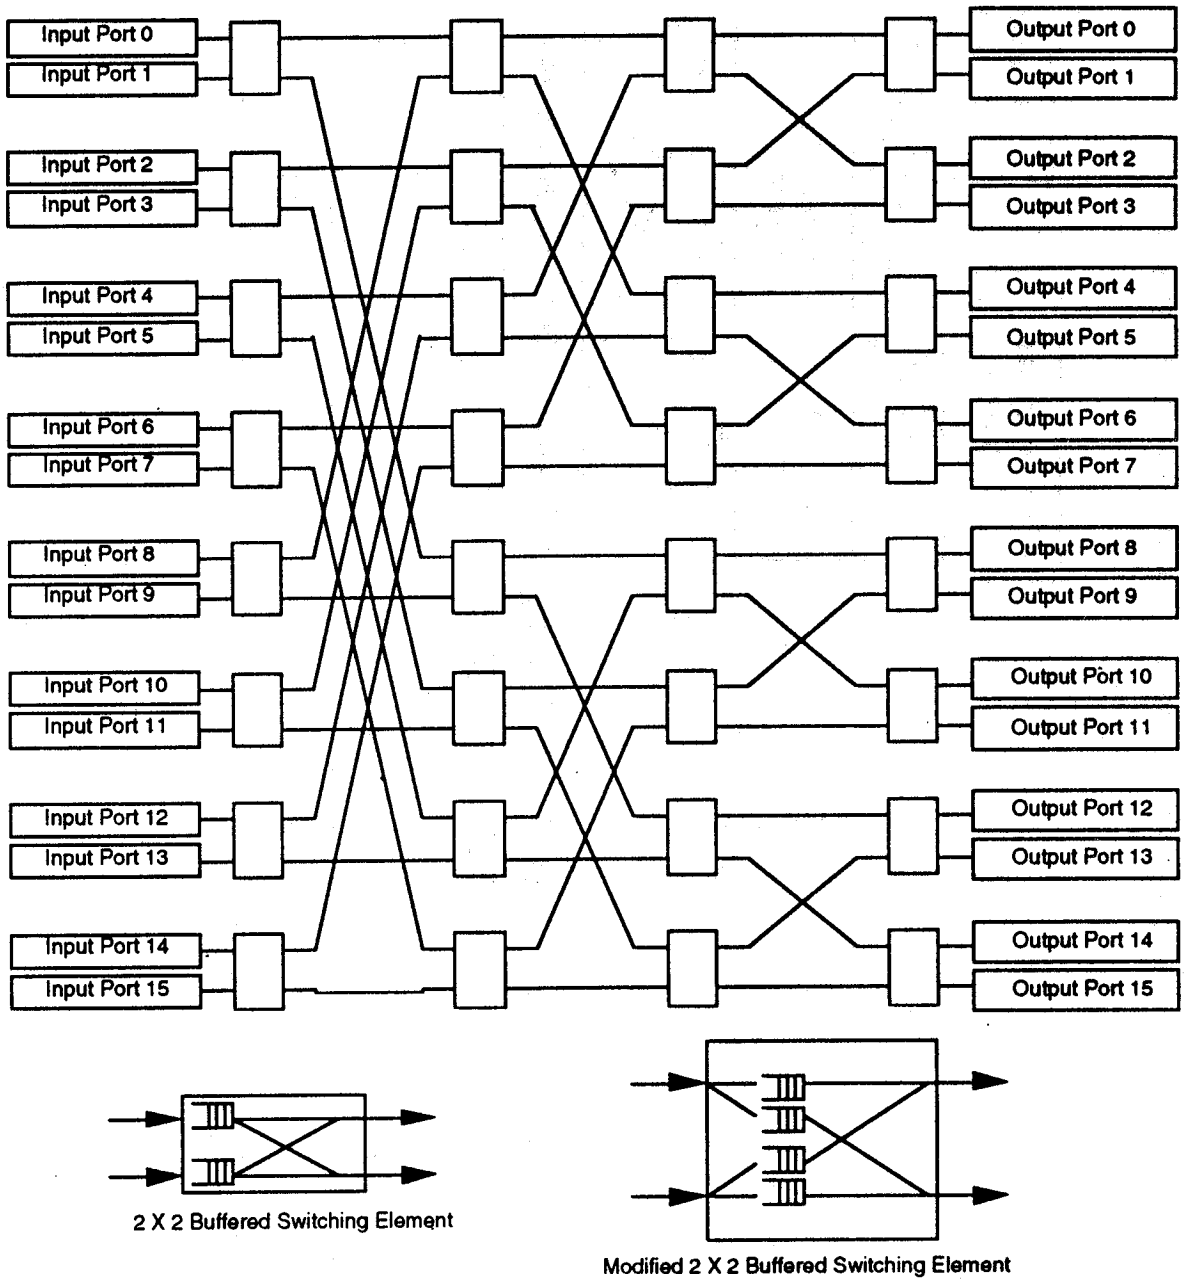

Figure 6-14. 16 x 16 Buffered-Banyan Switch with Base-2 Switching Element

to represented in base D. The advantage of using a larger switching element is that the places which internal blocking can occur are less. As can be seen from Figures 14 and 15, the number of places internal blocking can occur for the base-2 16 x 16 network is 3 while the other one for the base-4 16 x 16 network is only one. It is noted the buffer size for each switching element has to be increased correspondingly according to the switching element size.

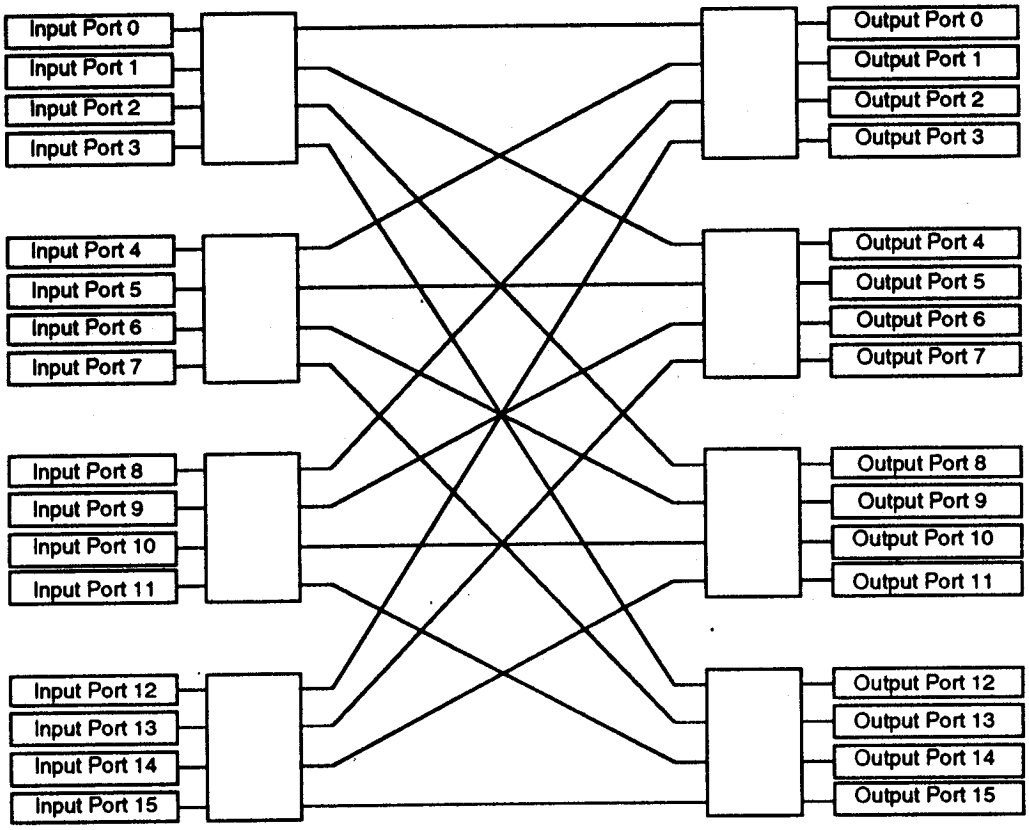

Figure 6-15. 16 x 16 Buffered-Banyan Switch with Base-4 Switching Element

(d) Distribution Network

If the incoming traffic pattern is irregular such that congestion always occurs at certain locations (called community of interest), the packet loss ratio will be intolerable. A distribution network is to randomize the input traffic pattern so that the congestion does not occur at the same place of the switching fabric each time (see Figure 6-16). A banyan network can be used a distribution network. During the call set up phase, the switching elements in the distribution network randomly select the routing state such as straight or exchange. The distribution network sends the packets to the outputs without regarding to the routing tags. During the transfer phase, the distribution network behaves as a routing network and a fixed path is always used for the packets from the same input line to guarantee in-sequence delivery.

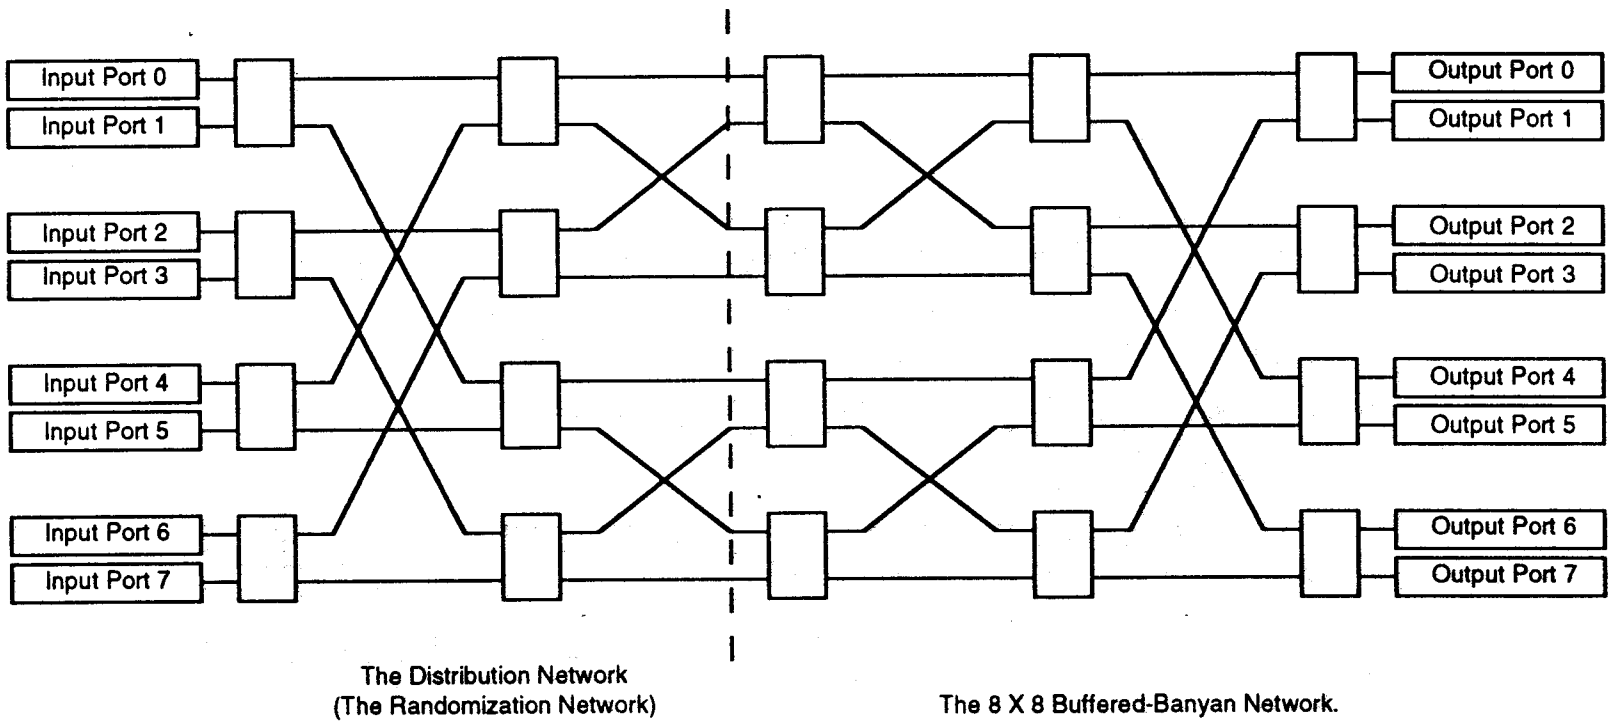

Figure 6-16. Distribution Network Approach

(e) Non-FIFO Queue at Each Switching Element

Since the packet in one switching element cannot move forward unless there is no output link blocking and there is an empty space in the next stage's buffer, the HOL blocking problem also occurs at the buffered switching element. A modification of the queue service discipline can be incorporated into the buffer design to reduce HOL blocking such that if the HOL packet cannot be transmitted due to output link blocking the packets at the back of the HOL packet can also be examined.

(f) A Separate Queue for Each Output Link at Each Switching Element

Another approach to resolve the HOL blocking problem is to use a separate queue for each output link at each switching element. To have a separate queue for each output link at a 2 x 2 switching element, the switching element needs to have 4 internal buffers (see Figure 6-14). Packets from each input line are stored in the buffers according to the destination routing bit corresponding to a particular stage, i.e. the output link the packet is destined to. In this way, the head of line blocking problem due to output link blocking at each switching element is completely resolved.

6.2.1.1.2 Blocking Unbuffered Banyan Switch

In this class, the switching fabric has an internal blocking problem. The buffering of packets is either performed at the input port, at the output port, or at both. The techniques used to improve the switch throughput for the slotted mode operation include the following:

- increased switch speed with output buffering (IB, OC)
- number of banyan networks in parallel (IB, OC)
- switching element with a larger size (IB)
- number of transmitters at each input port (HOL)
- number of receivers at each output port (OC)

The techniques used to improve the switch throughput for minislotted mode operation include the following:

- same as the slotted-mode operation
- path setup strategy for p parallel banyan networks with p transmitters at each input port and p receivers at each output port (IB, OC,HOL)
    - broadcast one setup packet to p copies (flooding)
    - send p different setup packets to p copies (random)

- send p set packets sequentially starting from the first transmitter (searching)
    - the setup packet at the first transmitter can try at most p copies in p minislots
    - the setup packet at the second transmitter can try at most p-1 copies in p-1 minislots; and so on.

Since all the approached used in the slotted mode operation can be used for the minislotted mode operation, only the techniques used in the minislotted mode operation will be discussed.

As mentioned above, in this class of switching networks, buffering is performed at the switching fabric level. The switching fabric level buffering can be performed either at the input port or/and at the output port. The place of the buffers is determined based on the throughput requirement, the operation procedure, the packet transfer protocol, and the hardware complexity of the switch.

(a) Input Buffering

In this scheme, the packets are buffered at the input ports. Since the switching fabric is blocking, the packet transfer protocols can use the setup phase + forwarding phase or the forwarding phase + retransmission phase protocols. It has been shown that the throughput of a switch with FIFO input queue cannot exceed 58% due to HOL blocking. The techniques used to relieve this problem have been discussed before.

(b) Output Buffering

Output buffering is required if the switch can deliver more than one packets to the output port during one link packet slot time. The output port has a multiple-input port buffer which can receive more than one packets at the same time so that the output contention problem can be relived. If the number of arriving packets exceeds the capacity of the multiple-input port buffer, packets will be lost. In this sense, pure output buffering is usually used for a nonblocking switching fabric. This scenario is deferred to the next subsection. There is one special case in the blocking buffered banyan switches, the cascaded banyan switch, that uses output buffering approach. This case will be studied latter.

(c) Input Buffering Plus Output Buffering

As mentioned above, for the output buffering approach if the number of arriving packets exceeds the capacity of the multiple-input port buffer, packets will be lost. The loss can be avoided if the transmission of packets are scheduled at the input port queues so that the number of delivered packets will not exceed the capacity of the multiple-input port buffer at the output.

In this scheme, packet transmission is scheduled at the input port queues. To effectively improve the switch throughput, either the switch has to operate at a faster speed than the link speed or multiple switching fabrics are stacked in parallel so that more than one packets can be processed in one link packet slot time (packet size/link speed). Therefore, output buffering is required to hold the packets.

Several examples of blocking buffered banyan switch are introduced as follows.

6.2.1.1.2.1 Unique Path Unbuffered-Banyan Network

The internal blocking in the unique path unbuffered-banyan is very sever such that the throughput of the switch is very low. Hence, the unique path unbuffered-banyan network is not suitable for the switching fabric unless the topology has been modified or switching elements contain buffers as mentioned before.

6.2.1.1.2.2 Multipath Unbuffered-Banyan Network

In this approach, multiple copies of unbuffered-banyan networks are stacked to create multiple paths between any input-output pair or a larger size switching fabric is used (see Figures 17 and 18). The immediate result is that internal blocking is relieved since the link utilization of each copy is reduced to 1/p of that of a single copy, where p is the number of copies stacked.

For a blocking switching fabric, input queueing is a must to perform the packet transfer protocol. There are many ways of operating the switch based on the number of transmitters, the number of receivers, the number of queues at each input port, and whether output buffering is used.

The simplest scheme is described first. In this scheme there are p copies of networks stacked in parallel and there is only one transmitter at the input port, one receiver at the output port, and no output buffering. During the path setup time, the setup packet searches the idle copy from the p copies to setup the path. if the setup packet finds the path in one copy is busy, then the setup packet uses another copy to search for the path. Note that the setup packet not only reserves a path within the switching fabric but also reserves the output port so that the packets destined to the same output port will be scheduled at different transmission slot time. Hence, the setup packet scheme resolves internal blocking and output blocking at the same time.

The second scheme has one TX at the input port and multiple receivers at the output port. During the path setup time, the input port broadcasts the setup packet to all the copies at the same time. The output port selects one of the duplicated setup packets that reach the output port and sends an ack back to the original input port. The reason the flooding technique can improve the switch throughput is that internal link blocking resolution method is implemented differently at different copies. For example, if packet a and packet b are collides at stage 1 of the switching fabric, copy 1 allows packet a to pass through while copy 2 allows packet b to pass through.

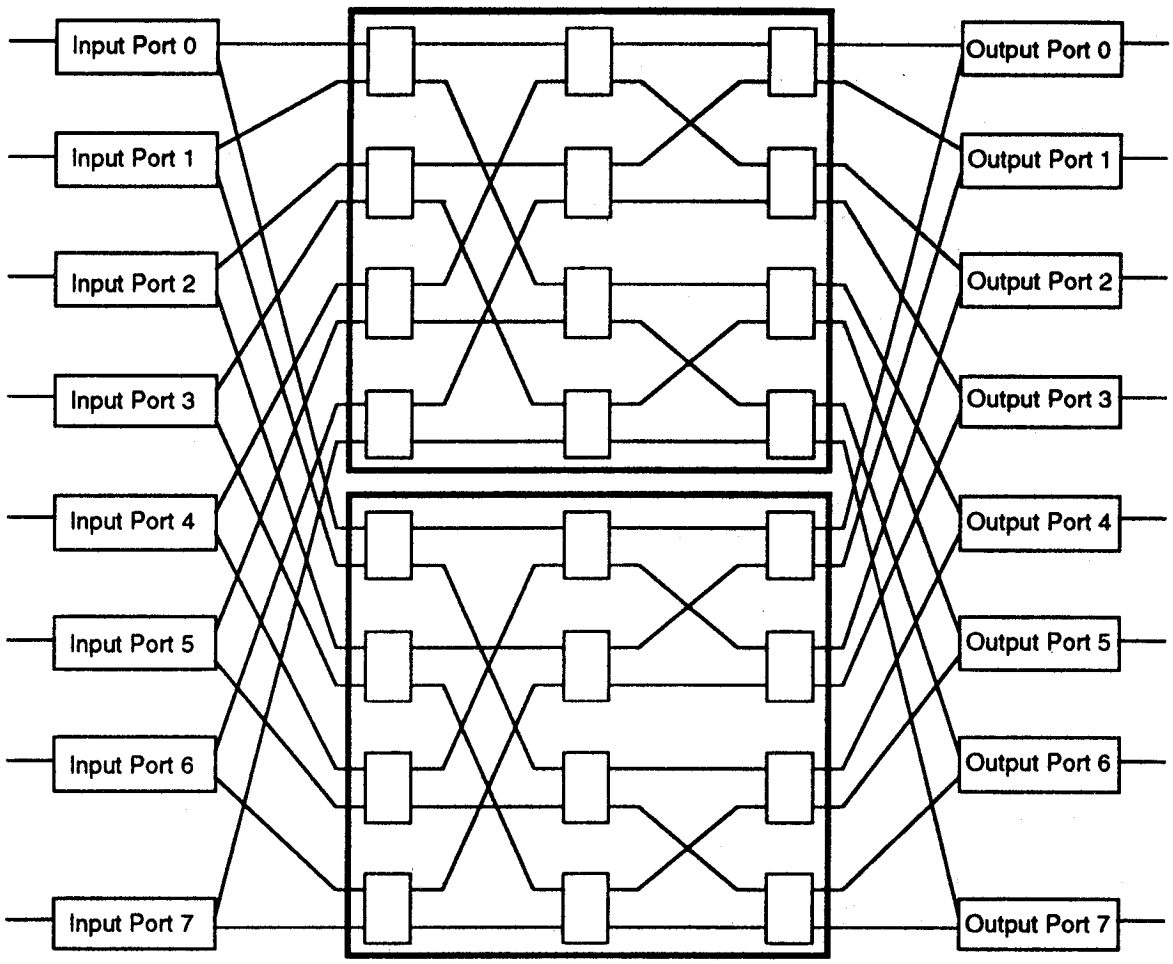

- Internal Blocking Resolution: 2 disjoint paths exist between each input-output pair.
- Output Contention Resolution: 2 packets can arrive to one output port at the same time.

Figure 6-17. 8 x 8 Double-Banyan Blocking Network

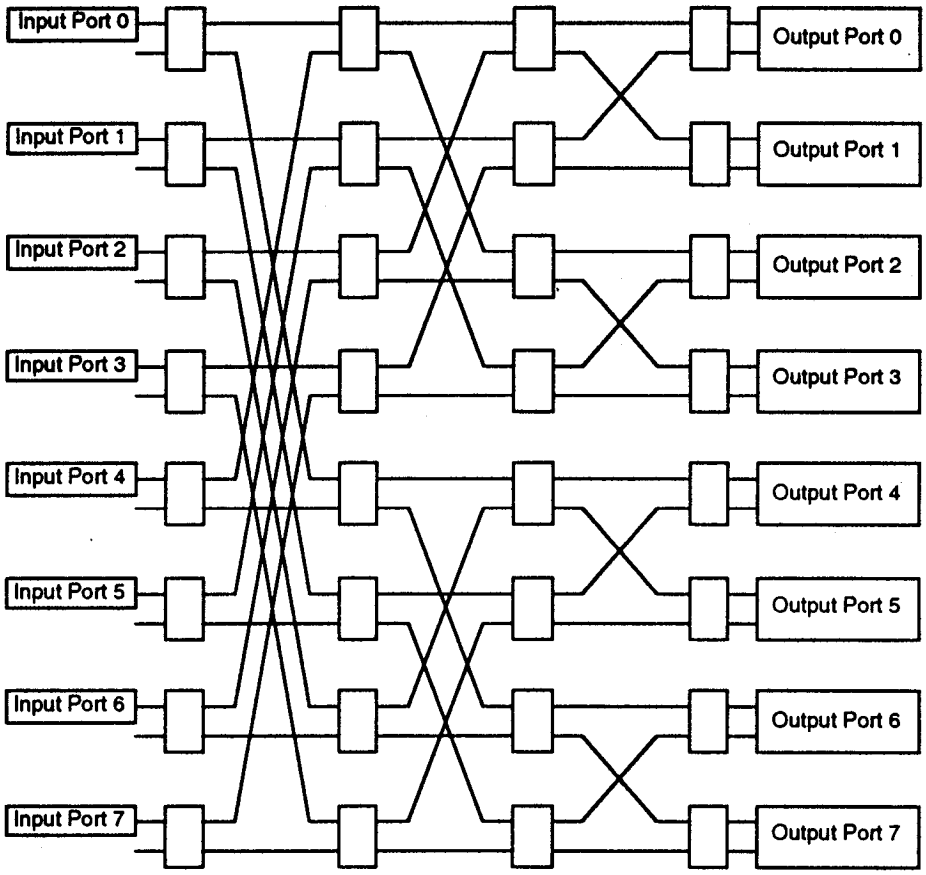

• Output Contention Resolution: 2 packets can arrive to one output port at the same time.

Figure 6-18. 8 x 8 Switch using a 16 x 16 Banyan Blocking Network

The HOL blocking problem associated with the input queueing can be relived by using multiple transmitters at the input port. The number of transmitters is the same as the number of copies. These schemes are introduced below.

The third approach is to use multiple transmitters at the input port, one receiver at the output port, and no output buffering. In the path setup phase, if the setup packet in the first transmitter is blocked in the first copy, the other packet with a different destination in the second transmitter can use the second copy to set up the path; and so on. This effectively reduces the effect of HOL blocking.

The fourth approach is to use multiple transmitters at the input port, multiple receivers at the output port, and output buffering. In the path setup phase, the setup packets at different transmitters are sent to different copies randomly at the same time. Since the output port has multiple receivers which can receive more than one packets at the same time, output buffering is required for this approach.

The fifth approach is also to use multiple transmitters at the input port, multiple receivers at the output port, and output buffering. There are p minislots reserved for the packet setup phase, where p is the number of copies. At the first minislot, the packet at the first transmitter tries to set up a path using the first copy. If the packet encounters blocking either at the switching fabric or at the output port, the packet uses the second copy to set up a path at the second minislot. If the packet successfully set up a path at the first minislot, then the packet at the second transmitter can use the second copy to setup a path; and so on. In this sequential searching algorithm, the maximum number of reserved minislots to setup a path for the packet at the first transmitter is p. The maximum number of reserved minislots to setup a path for the packet at the second transmitter is p-1; and so on. Note that corresponding to each minislot of the setup phase, a different copy is used for setting up the path for a packet.

6.2.1.1.2.3 Cascaded-Banyan Network

In this approach, multiple banyan networks are cascaded in series [6-4] as shown in Figure 6-19. The packet transfer protocol uses forward-and-store. Since no input buffering is provided, the operation of the self-routing switching element has been modified so that when two packets request for the same output link, one is granted for access and the other one is routed to the wrong link. For a fair operation, the switching element randomly selects one for access the output link. Since the banyan network is a unique path network, the packet which has been routed to the wrong link cannot reach its destination in this copy. Before this packet leaves the switching element, it is marked as a rerouted packet so that it will not affect the normal packet transfer at the following stages. At the output ports of the banyan network, a demultiplexer is employed to extract the unmarked packet which arrives at the right destination, and send it to the output buffer. The marked packets which arrive at the wrong destinations will be unmarked first and be sent to the second banyan network, and the same procedure is repeated. The number of banyan networks required to be cascaded in series depends on the traffic loading and the switch size. Since the utilization of the

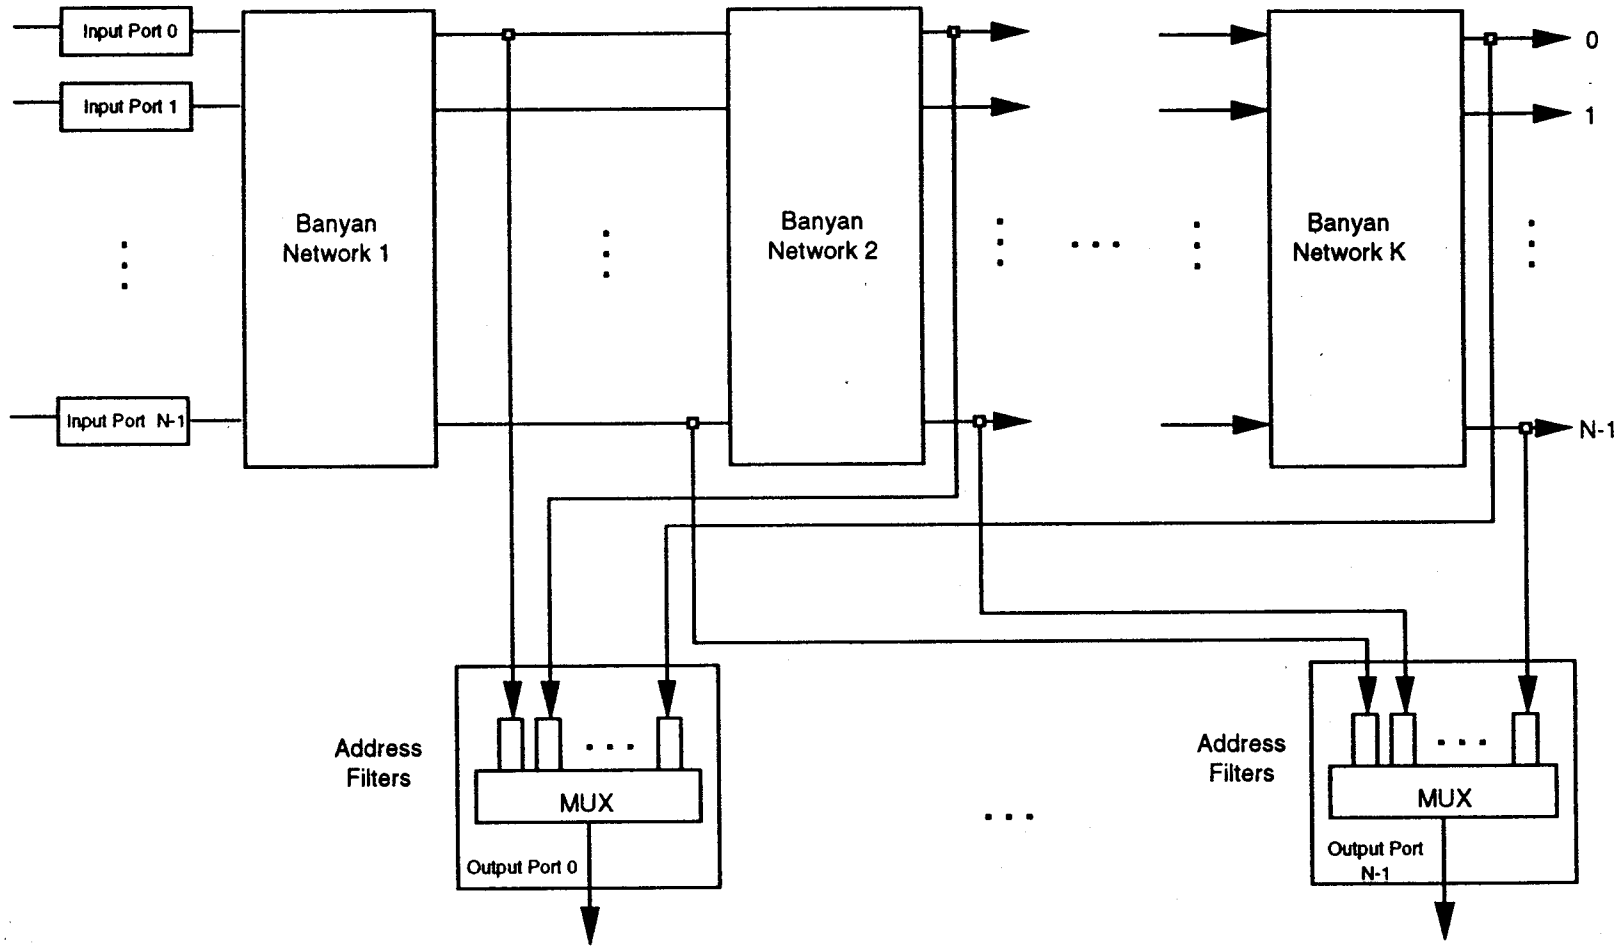

Figure 6-19. Cascaded Banyan Network

link decreases on the successive banyan networks, the blocking problem is also reduced on the successive banyan networks. If the number of stages is large enough, it is guaranteed that the packet loss ratio is vert small.

The rerouting procedure used in every switching element solves the internal blocking and output blocking all at the same time, which is a very attractive feature. However, since rerouting packets may arrive to the output port at any instant, the packets may be transmitted out of sequence. A resequencing buffer is required at the output port.

6.2.1.1.3 Nonblocking Unbuffered Banyan Switch (Sorted-Banyan-Based Network)

One of the important properties of the banyan network is if the incoming packets are arranged either in ascending or descending orders and there is no empty line between any two active lines, there is no internal blocking within the banyan network. An active line means that there is a packet waiting to be transmitted. A way of arranging the arriving packets in a descending and concentration order is to use a batcher sorting network. Although the sorted-banyan-based network as shown in Figure 6-20 is internally nonblocking, the output blocking problem and HOL problem (if input buffering is used) still exist. The techniques used to improve the switch throughput include the following:

- increased switch speed with output buffering (OC)
- number of sorted-banyan-based networks in parallel (OC)
- number of receivers at each output port (OC)
- number of transmitters at each input port (HOL)
- non-FIFO queue with a larger checking depth at each input port (HOL)

One example of designing an unbuffered banyan switch with a nonblocking switching fabric is introduced as follows.

A parallel switch configuration has been proposed and studied mostly due to high throughput, medium complexity, fault-tolerance with graceful degradation, high reliability, and a variety of ways of operating the switch [6-6].

In Figure 6-21, two transmitters are provided at the input port and two receivers are at the output port. Basically, with this scheme the switch throughput is doubled compared to the switch with only one switching fabric.

The are several ways of resolving the output contention problem and head of line blocking problem using this configuration.

The first three methods used the setup packet protocol that have been discussed in the previous subsection. The first one is the random scheme that the input port transmits two setup packets to two different copies randomly at the same slot time. The second

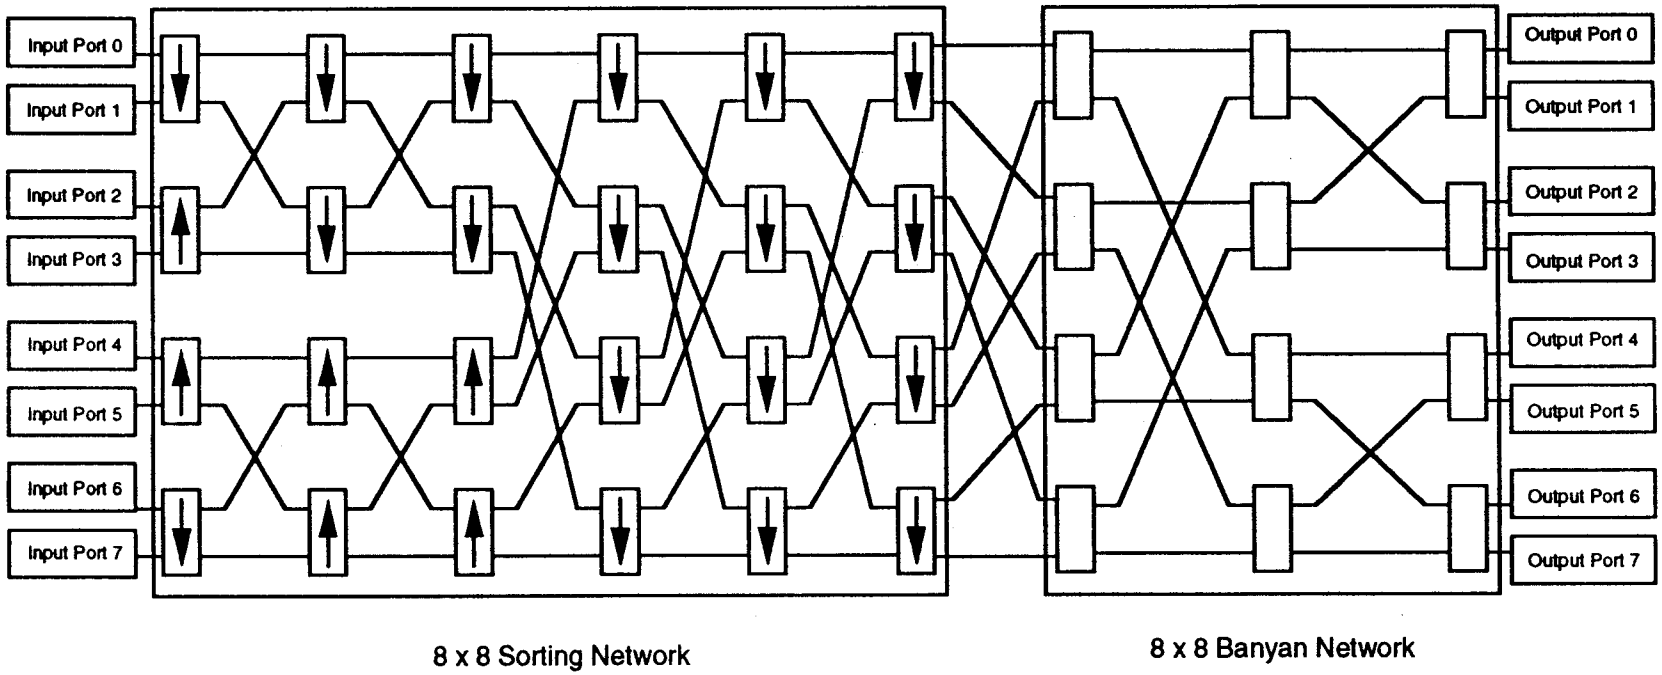

Figure 6-20. Sorted-Banyan-Based Network (Point-to-Point Nonblocking Network)

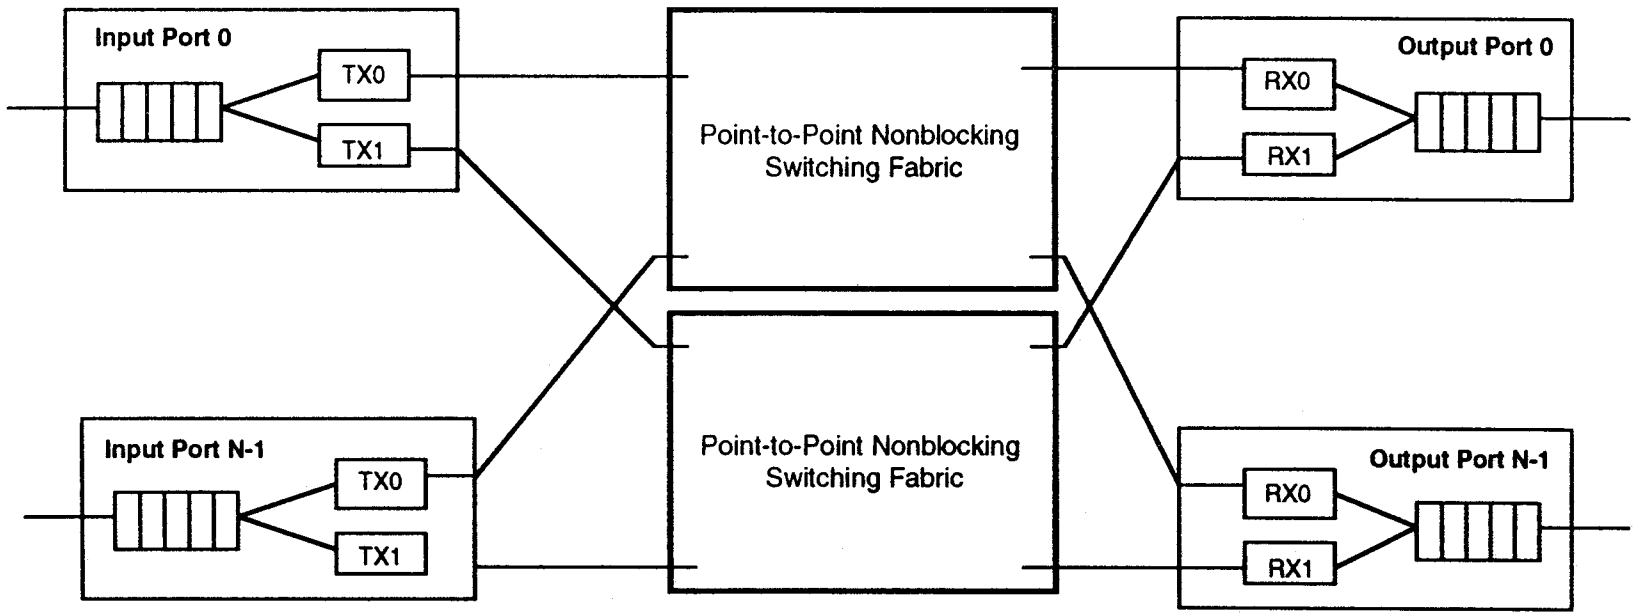

• Output Contention Resolution and Head of Line Blocking Resolution

Figure 6-21. Parallel Switches

one is the sequential searching scheme. There are two minislots reserved for path setup phase. The packet at the first transmitter has at most two minislots to try to set up a path while the packet at the second transmitter has at most one minislot to setup up a path.

The other approach is to apply output port reservation schemes to the input ports. The output port contention resolution (or output port reservation) schemes are introduced as follows.

(a) Multiple Queues at Each Input Port

The number of queues at each input port is the same as the number of output ports (see Figure 6-22). The scheduling algorithm examines every queue in every input to determine the transmission sequence so that output contention can be resolved. The scheduling algorithm is accomplished using a controller. In this sense, this approach is close to the output queueing method (which will be discussed latter), which has been shown to achieve the best performance.

(b) Output Port Reservation with Tokens

The scheme is to reserve the output ports using the reservation scheme mentioned before so that each time the packets presented to the sorting network all have distinct destination addresses.

To improve the throughput of the switch, a non-FIFO input queue with the windowing scheme is used (see Figure 6-23). In this scheme if the first packet is blocked due to output blocking, the scheduling algorithm also examines (searches) the packets on the back of the first packet. This scheme is also referred to input queue by-pass [6-6]. The number of packets examined each time depends on the preset window size or the checking depth. If one of the packets within the checking depth has a chance to be transmitted, this packet will be transmitted first. In this sense, the FCFS input queue has a checking depth 1 while a non-FIFO input queue has a checking depth greater than 1. Theoretically, if the checking depth is infinite, the throughput of the switch can reach 1. However, in practical, the checking depth is finite and less than O(10).

This token scheme mentioned above is a very powerful technique in terms of flexibility, easy implementation, and priority control. For example, it can also be applied to the parallel switches. In this case, the number of tokens is 2N, where one pair of tokens is used for one output port. Both the transmitters at each input port can search the tokens according to the destination routing tags of the current packets.

(c) Output port Reservation with an Arbitrator

The other way of resolving the output contention problem is to use a bidirectional sorting network [6-7] (see Figure 6-24). The sorting network has the property of arranging the arriving packets based on their destination addresses either in ascending or descending orders. At the beginning of every slot, the input ports send a setup packet containing the destination addresses (or routing tag) to the bidirectional sorting

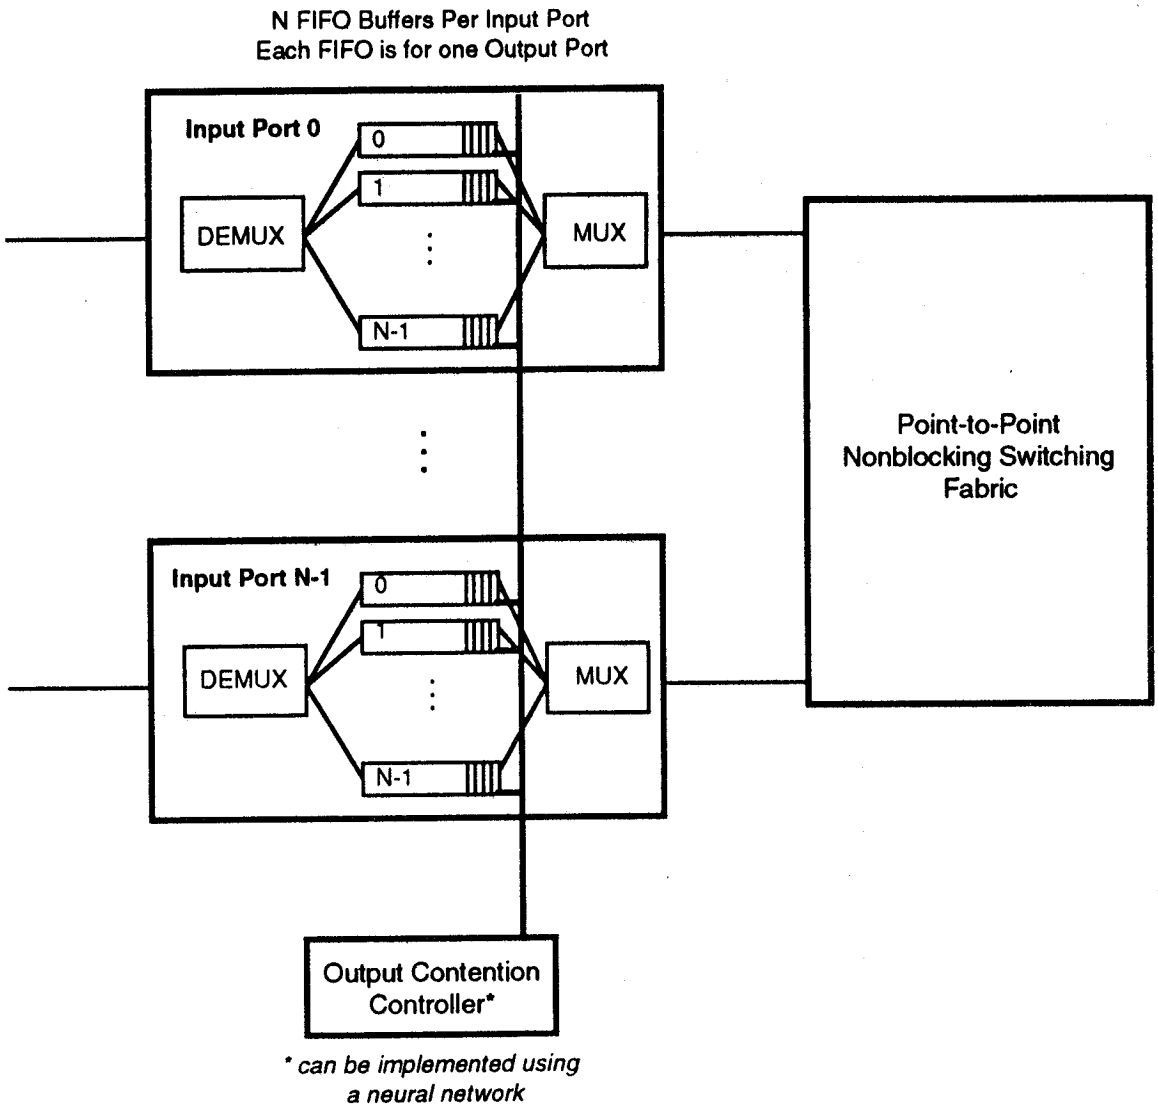

Figure 6-22. N Parallel FIFOs Per Input Port Structure to Resolve Head of Line Blocking

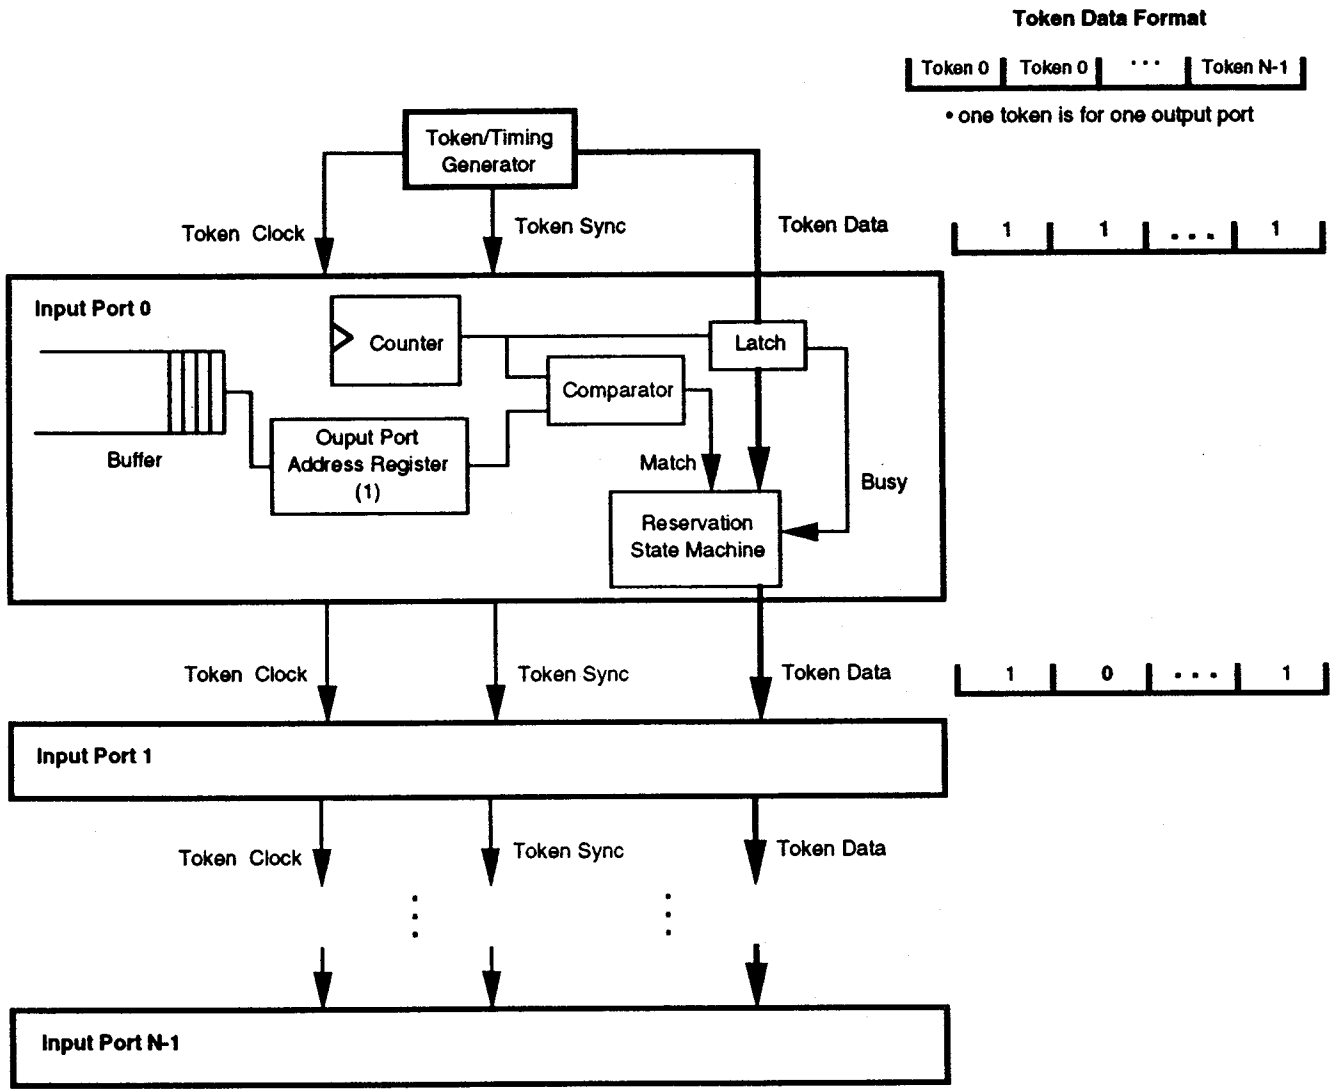

Figure 6-23. Output Port Reservation Scheme with Tokens

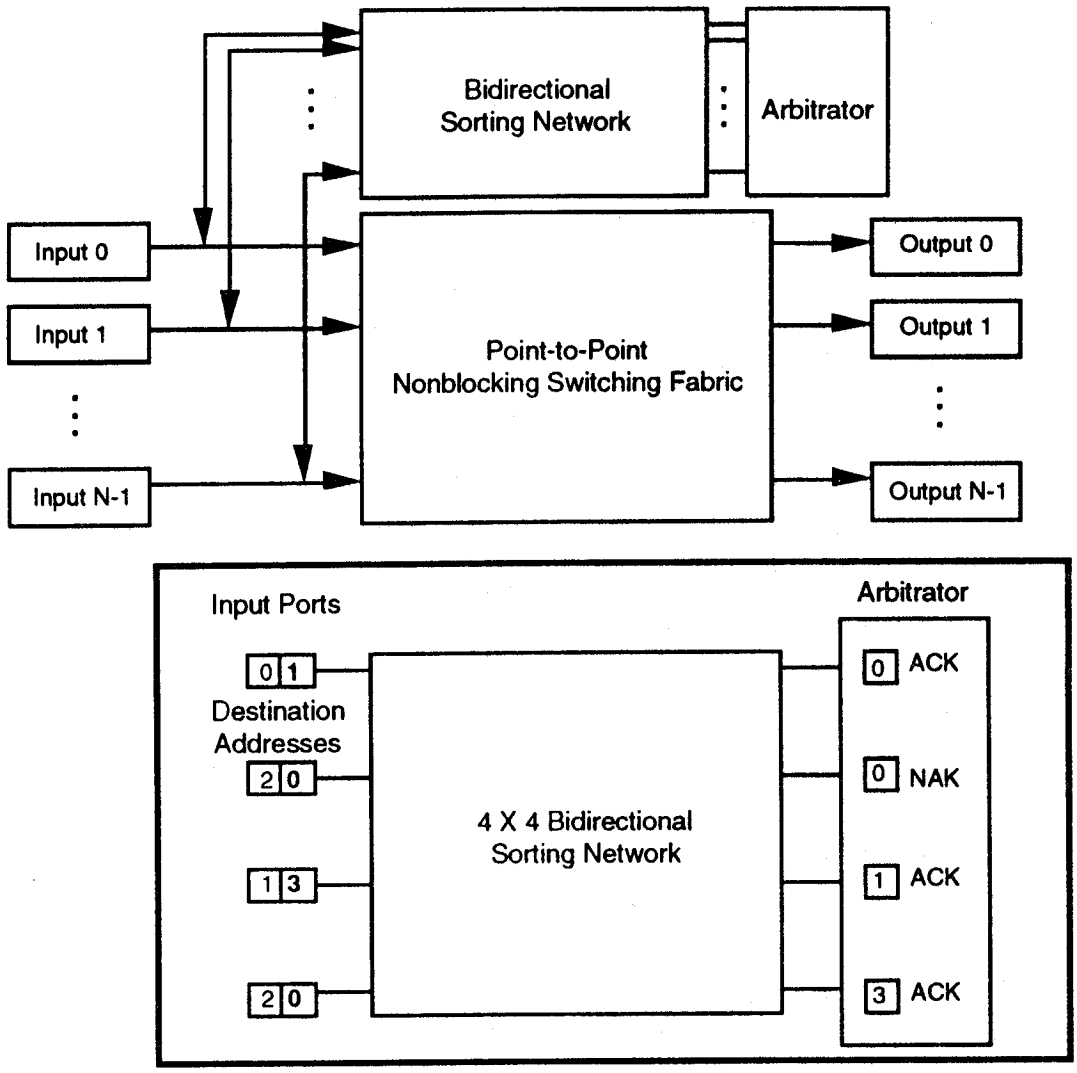

Figure 6-24. Output Contention Resolution Using a Sorting Network

network, and finally reaches the arbitrator. All the setup packets that have the same destination addresses will be adjacent to each other. A distinctive set of destination addresses can be selected easily using an array of comparators within the arbitrator. The arbitrator sends ACKs and NACKs through the bidirectional sorting network to the input ports to report the arbitration result. For the input ports whose packets have been selected for transmission at the next slot time receive an ACK. For the input ports whose packets have not been selected for transmission at the next slot time receive an NACK.

To improve the throughput of the switch, a large checking depth is required to increase the throughput (the average number of packets which can be transmitted at one slot time). To examine more depths, the input port which has received an ACK will send the same routing tag again. To guarantee that the packets which have wined the arbitration at the previous run still win the arbitration at this run, the routing tags of

these packets will be prepended with a priority bit so that these routing tags always win the arbitration. The input port which has received an NACK will send the routing tag of the packet behind the HOL packet.

(d) Output Port Reservation using an Output Contention Resolution Device

The output reservation is accomplished using an output contention resolution device [6-8]. Within this device, there are two arraies of registers A and B, where the number of registers in each array is N and the size of register is the size of the routing tag, that are used to hold the routing tags from all the input ports (see Figure 6-25). There is another array of bit registers R to hold the reservation result. If R_i is 0 at the end of output port reservation process, then input port i can transmit the current packet at the next slot. Initially, all the routing tags from the input ports will be loaded into the array A and array B; hence, the contents of array A and array B are exactly the same. All the bits in array R are 0. To reserve the output ports, the routing tags between A and B have to be compared with each other so that a distinctive set of routing tags can be selected. This operation is achieved by fixing array A and rotating array B. After each rotation, the contents of array A and array B are compared. If $A_i \neq B_i$, there is no action. If $A_i = B_i$, then one routing tag will be selected for transmission. Now the problem is which routing tag should be selected. To resolve this problem, another array of priority bit registers P is used.

Initially, all the bits in array P are all 0. Starting from the first rotation cycle, a bit 1 is loaded into P_0 (see Figure 6-26). In this situation A_0 has the routing tag from input port 0 and B_0 has the routing tag from input port N-1. If $A_0 = B_0$ and now $P_0 = 1$, R_0 remain 0. This means the routing tag at A_0 wins the arbitration. A_1 has the routing tag from input port 1 and B_1 has the routing tag from input port 0. if $A_1 = B_1$ and now $P_1 = 0$, R_1 is changed to 1. This means the routing tag at A_1 loses the arbitration. It can be seen that if $P_i = 1$, it means that the input port number at B_i is larger than the input port number at A_i. if $P_i = 0$, it means that the input port number at B_i is smaller than the input port number at A_i. It can be observed the arbitration rule for $A_i = B_i$ situation is that whoever holds the routing tag from a input port of a smaller number wins the arbitration. This means the priority is given from top to down of the input ports. This priority is implemented using the priority bit register P.

At the second rotation cycle, P_0 and P_1 all have bit 1. The comparison is performed between array A and array B following the same procedure mentioned above.

In summary,

- after every rotation, the contents of A_i and B_i are compared.
  - if $A_i \neq B_i$, no action.
  - if $A_i = B_i$,
    - if $P_i = 0$, $R_i = 1$
    - if $P_i = 1$, no action.
- after N-1 rotations, all the input port i with $R_i = 0$ can be transmitted.

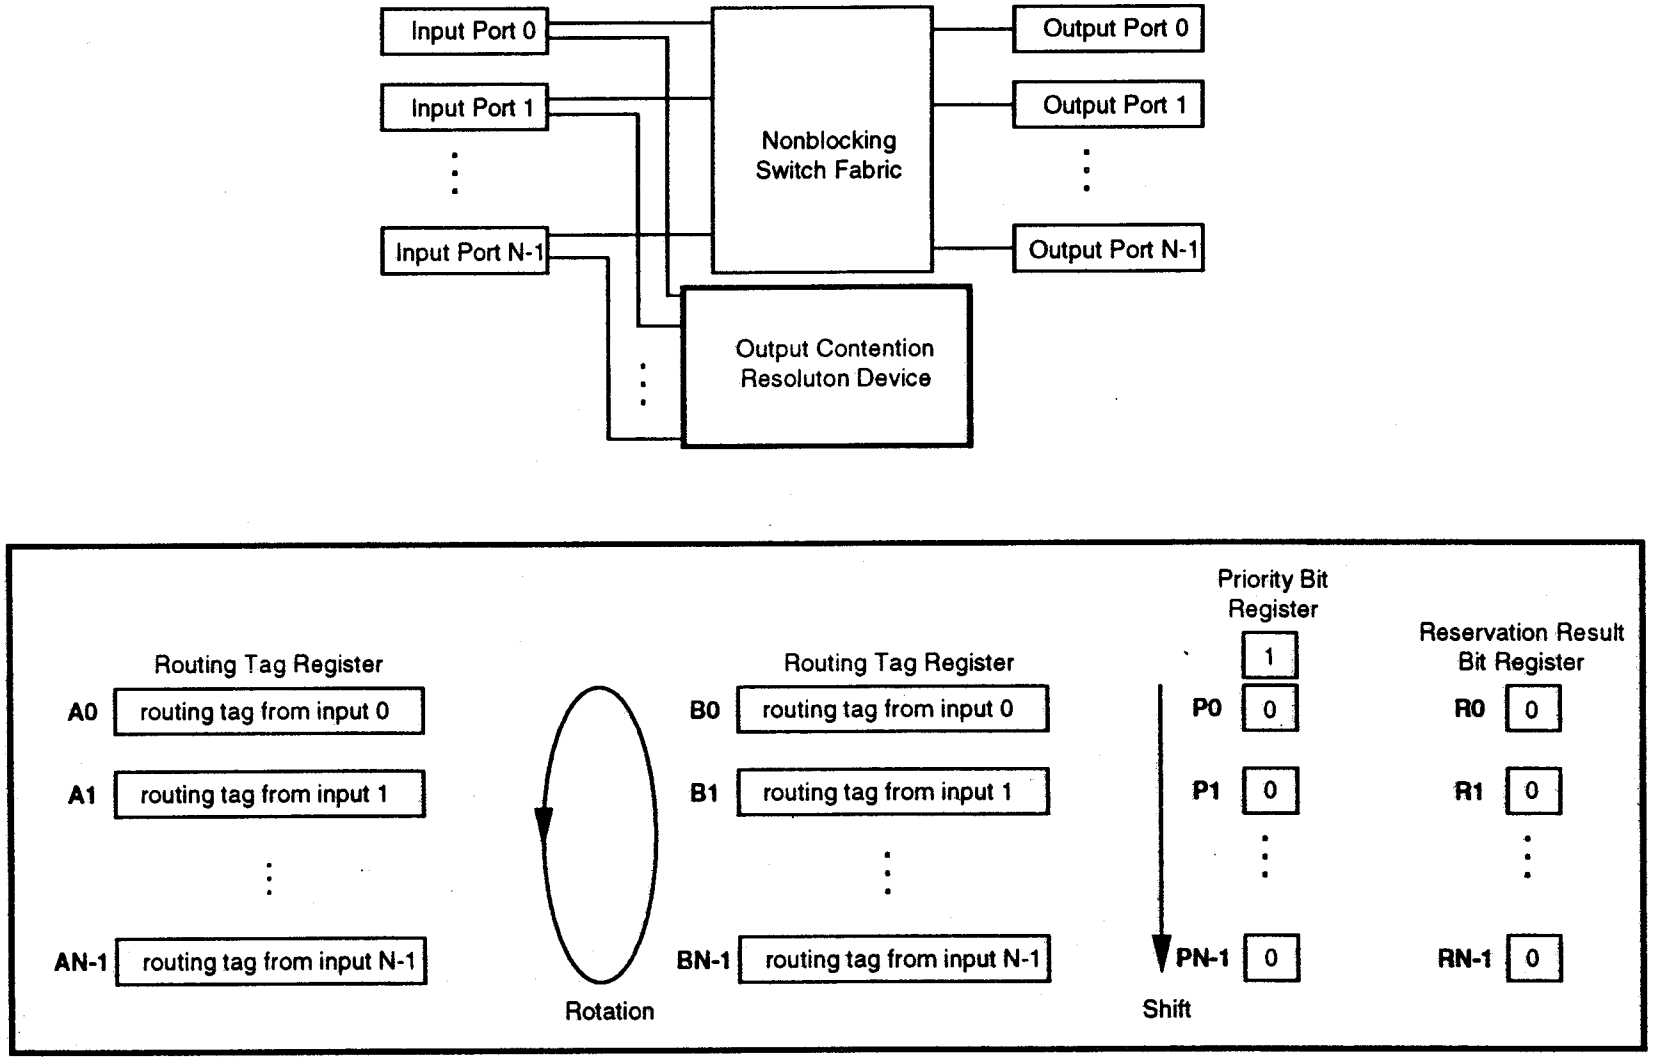

Figure 6-25. Output Contention Resolution Device

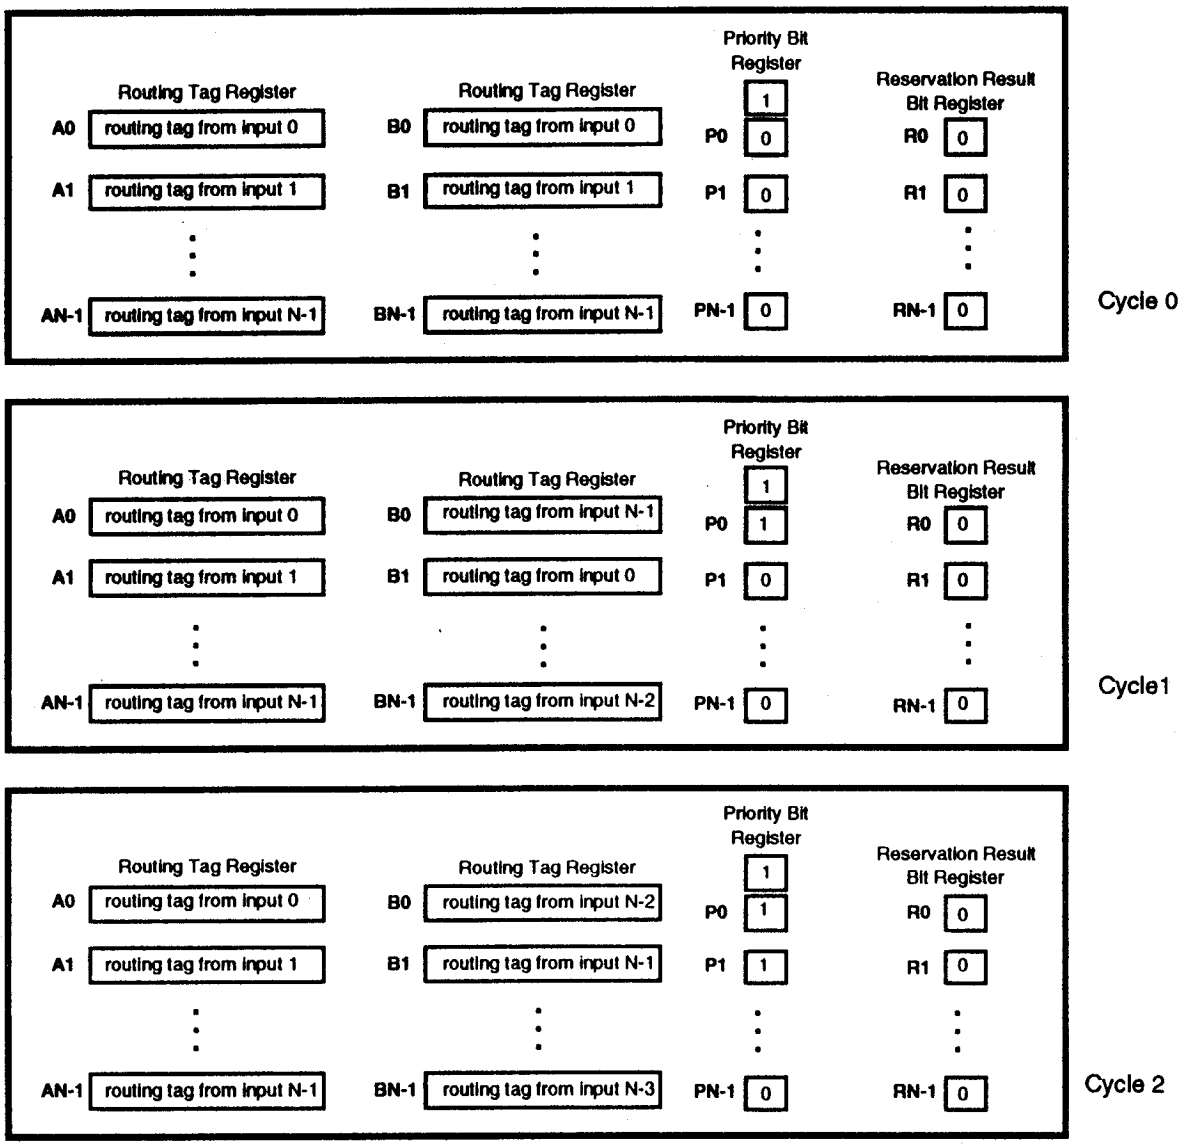

Figure 6-26. Output Contention Resolution Device

To have a fair access to an output port, the priority bit P_0 can be loaded to a different P_i at the beginning of arbitration. To check more depth into the input buffer, extra registers are required [6-8].

(e) Sorted-Banyan-Based Network with Reentry Network

The scheme is to use a trap network after the sorting network [6-9] (see Figure 6-27). The trap network resolves output contention by marking the packets with distinct output addresses. In implementation, the trap network is implemented using an array of comparators. The function of the concentrator is to send the marked packets to the banyan network so that packets presented to the banyan network all have distinct destination addresses. The concentrator sends the trapped packets back to the reentry

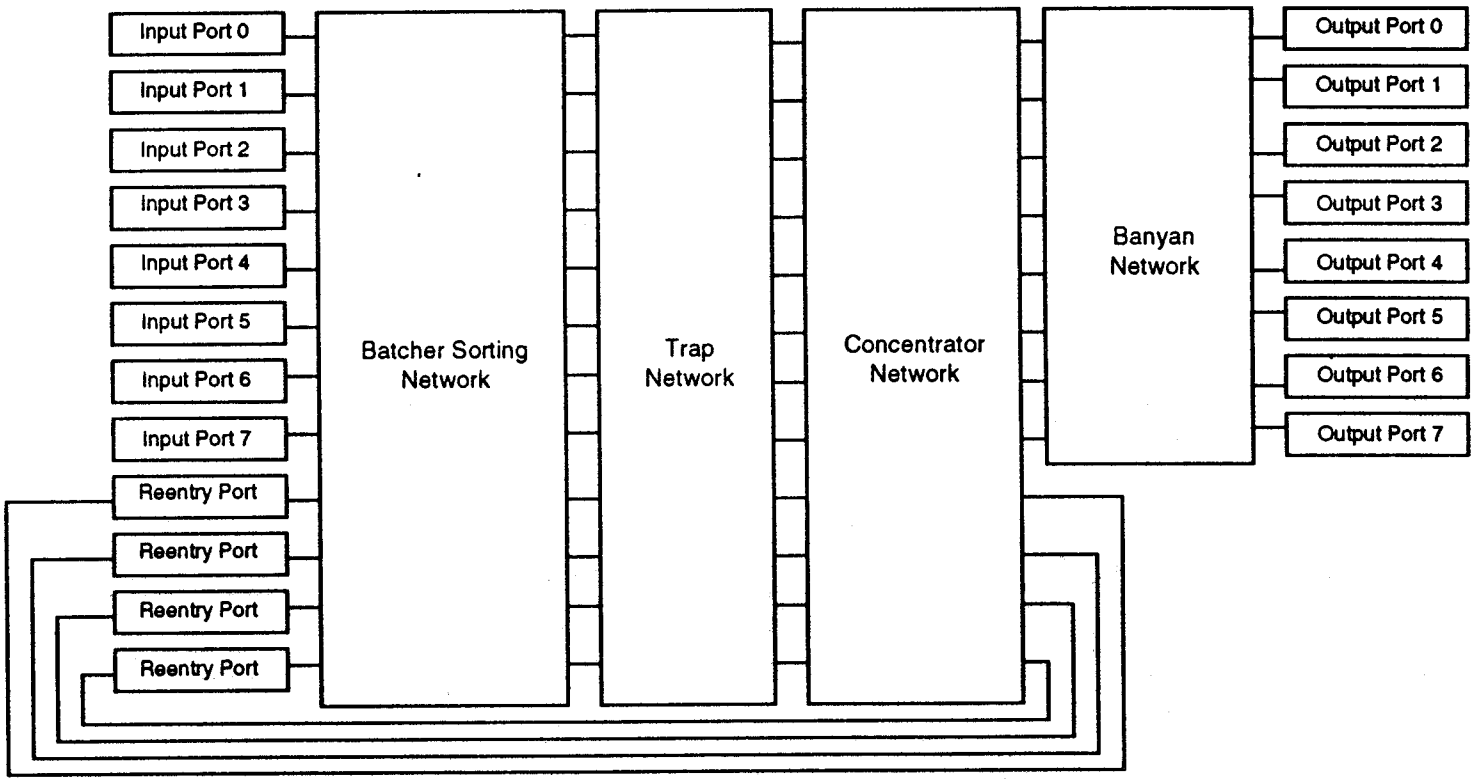

• Trap network resolves output contention by marking packets with distinct output addressess.

Figure 6-27. Sorted-Banyan-Based Network with Reentry Network

inputs of the sorting network. The packets in the reentry port are retransmitted during the next time slot. The size of the sorting network is larger than the size of the switch to accommodate the reentry ports. If the number of trapped packets is larger than the number of reentry ports, the packet will be lost. Also the packets may be delivered out-of-sequence due the trapped packets are sent back to the reentry ports not the original input ports. These retransmitted packets have to be given a priority higher than that of the new packets when conflict occurs at the output port; otherwise, there are chances that packets are transmitted out of sequence.

(f) Sorted-Banyan-Based Network with 3-Phase Algorithm

The output contention resolution algorithm is divided into three phases [6-10] (see Figure 6-28). At Phase 1 the input ports send setup packets to the trap network to resolve output contention, where the setup packet contains the source address and the destination address. At Phase 2 the trap network marked the setup packets with distinct destination addresses. All the marked setup packets will be sent an ACK packet back to the originating input ports, where the ACK packet contains the source addresses. To achieve this function, the outputs of the trap network and the corresponding input ports are connected. All the ACKs are sent to the input ports from the trap network first. The sorted-banyan-based network route these ACK packets using the source addresses to the corresponding output ports. The input port and the corresponding output port are also connected together. Hence these ACK packets are sent from the output ports to the corresponding input ports. At Phase 3 the input ports that receive ACK packets send the data packets prepended with the destination routing tags to the output ports.

(g) Sorted-Banyan-Based Network with Output Queueing

In this approach, the buffers are located at the output ports (see Figure 6-29). Note the prerequisite of using this approach is that there is no internal blocking within the switching fabric (as in the sorted-banyan-based network case); otherwise, the packet cannot reach the output port. To resolve the output contention problem so that the multiple packets destined to the same output port at the same time can be stored, the output port needs to have a multiple-input port buffer (a buffer with multiple receivers) and either the switching fabric has to operate at a speed faster than the line speed or there are multiple paths between each input and output pair. In the example shown in Figure 6-29, there are multiple banyan networks in the sorted-banyan-based network [6-11]. Therefore, the output port can receive K packets at the same time. The trap network operation is modified such that it will mark up to K packets which have the same destinations and send these packets to K different banyan networks. Note that the output contention problem is not completely resolved since it is possible there are more than K packets from different input ports destined to the same output port. These overflowed packets will be sent to the reentry ports.

If the switch speed is N times faster than the line speed, then all the packets which are destined to the same output port during the same slot can be buffered at the output port. Therefore, there is no output contention problem. However, the buffer

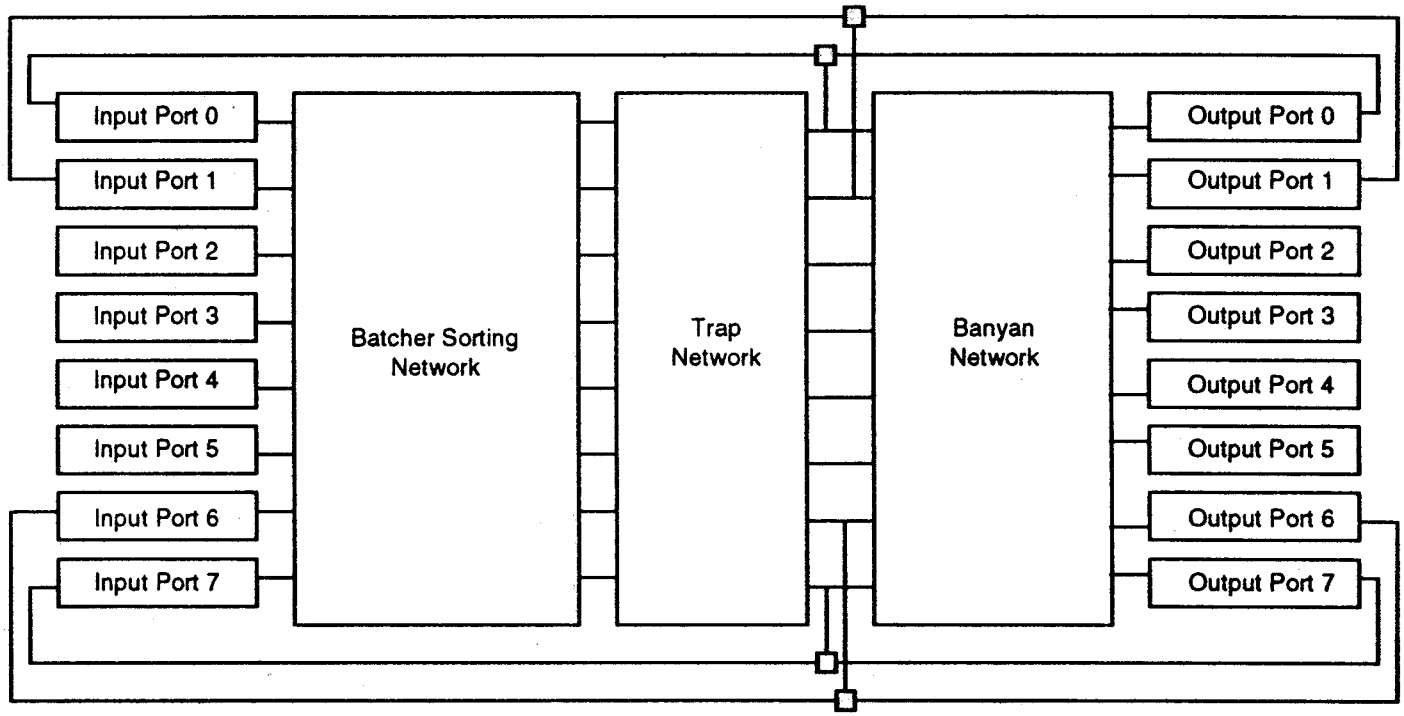

- Phase 1: Input ports send setup packet (source address + destination address) to the trap network
- Phase 2: Trap network resolves output contention by sending ack packets with distinct destination addressess back to the input ports and these ack packets will be routed to the original input ports using the source addresses
- Phase 3: The acknowledged input ports send the data packets to the output ports.

Figure 6-28. Sorted-Banyan-Based Network with 3-Phase Algorithm

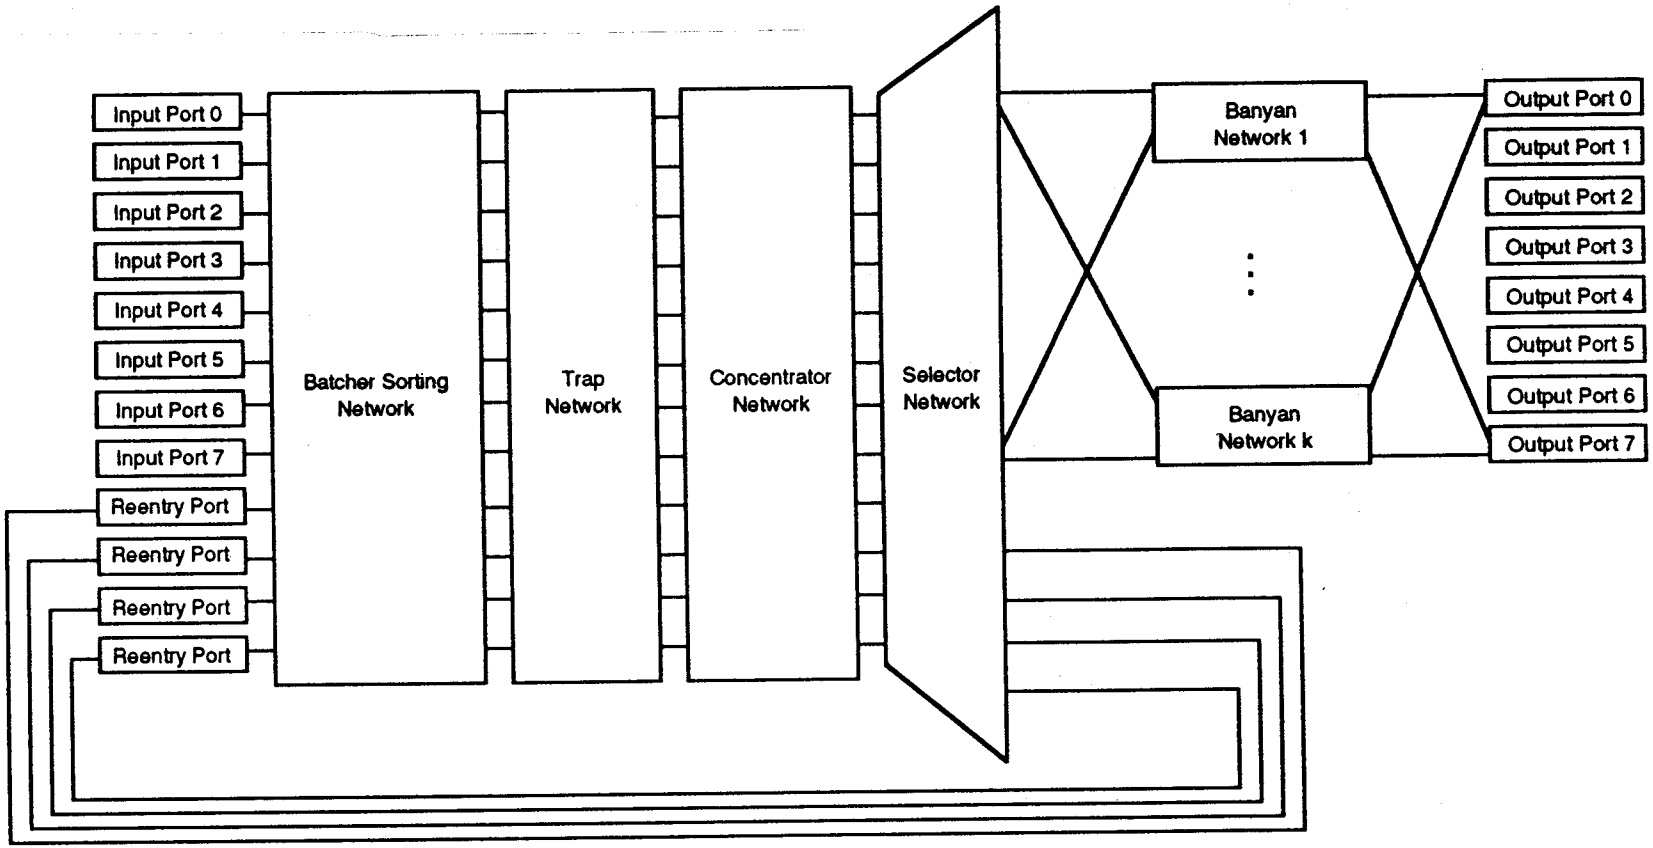

- Trap network resolves output contention.
- k parallel banyan networks provide k independent paths.
- The output port can receive k packets simultaneouly.

Figure 6-29. Sorted-Banyan-Based Network with Output Queueing

requirement at the output port will be very large if the link utilization is near the switch throughput according to queueing theory. Besides this approach is not feasible for high-speed application due to the switch speed has to be N times faster than the line speed.

6.2.1.2 N^2 Disjoint Paths among Inputs and Outputs

In this class of switching networks, there is a disjoint path between each input and output pair. The total number of paths within the switching fabric is N^2. Since there is a disjoint path between each input and output pair, the switching fabric is point-to-point nonblocking. Note that these switching networks are also point-to-multipoint nonblocking. However, since the format of the point-to-point routing tag is different from the point-to-multipoint routing tag, the implementation of the switch such as the address filter design or the switching element design is different for point-to-point and multicast cases even though the switching architectures remain the same.

6.2.1.2.1 Knockout Nonblocking Switching Fabric with Output Buffering

The knockout switch shown in Figure 6-30 uses the bus approach to interconnect the inputs and outputs [6-6]. There are N broadcast buses, one from each input port, in the switch and there are N filters at each bus interface of the output port. The total number of filters for the switch is N^2.

Since there is a disjoint path between any input-output pair in this topology, there is no internal blocking. The packet transfer protocol uses the forward-and-store scheme with output buffering; hence, there is no HOL blocking. The N filters at each output port performs as N receivers which can receive N arriving packets at the same time. After the N receivers, there is one output buffer which performs as a statistical multiplexer. The amount of buffering required at each output port depends on the packet loss ratio requirement.

6.2.1.2.2 Knockout Nonblocking Switching Fabric with Input Buffering

The disjoint path topology in the knockout switch forms a nonblocking switching fabric. Therefore, all the schemes used for the nonblocking sorted-banyan-based switching fabric with input buffering can also be applied to the knockout switching fabric (see Figure 6-31). Note that the input buffering scheme requires less hardware compared with the output buffering scheme.

6.2.1.2.3 Nonblocking Tree Network with Input Buffering

The tree network consists of two portions: the splitter and the combiner (see Figure 6-32). There are two ways of implementing the tree network. The first one uses the selective routing procedure, i.e., there are two control registers at each switching

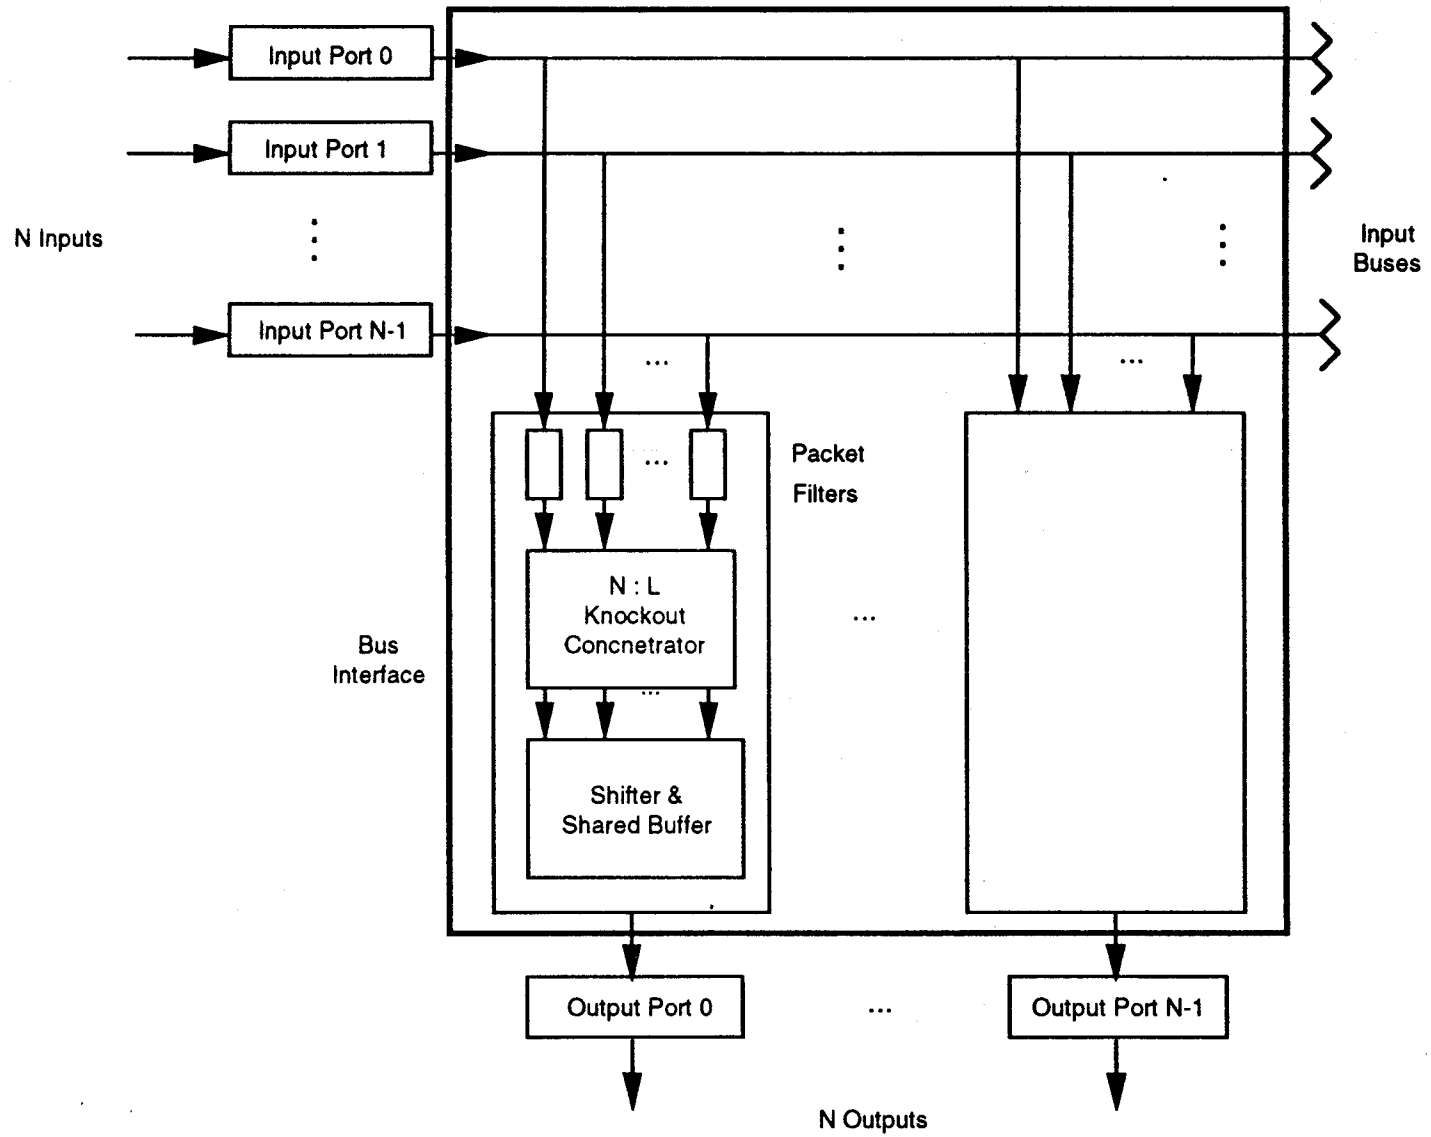

Figure 6-30. N X N Knockout Switch with Ouput Queueing

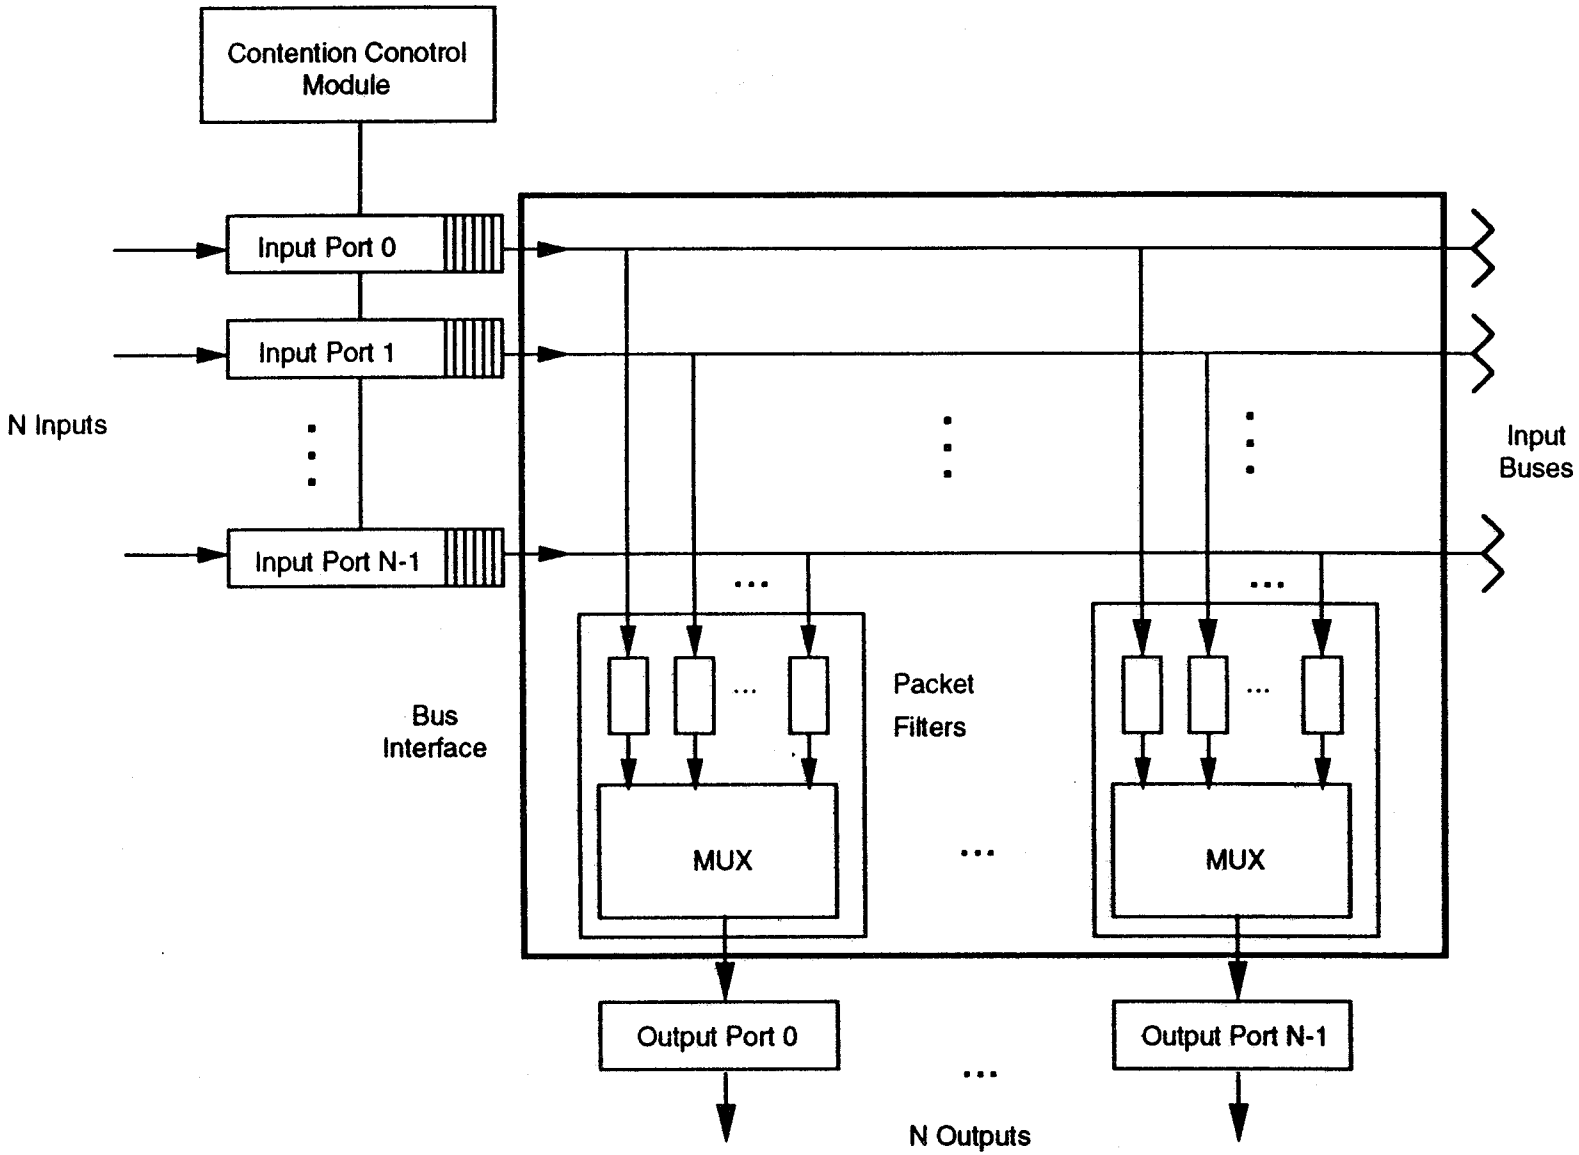

Figure 6-31. N x N Knockout Switch with Input Queueing

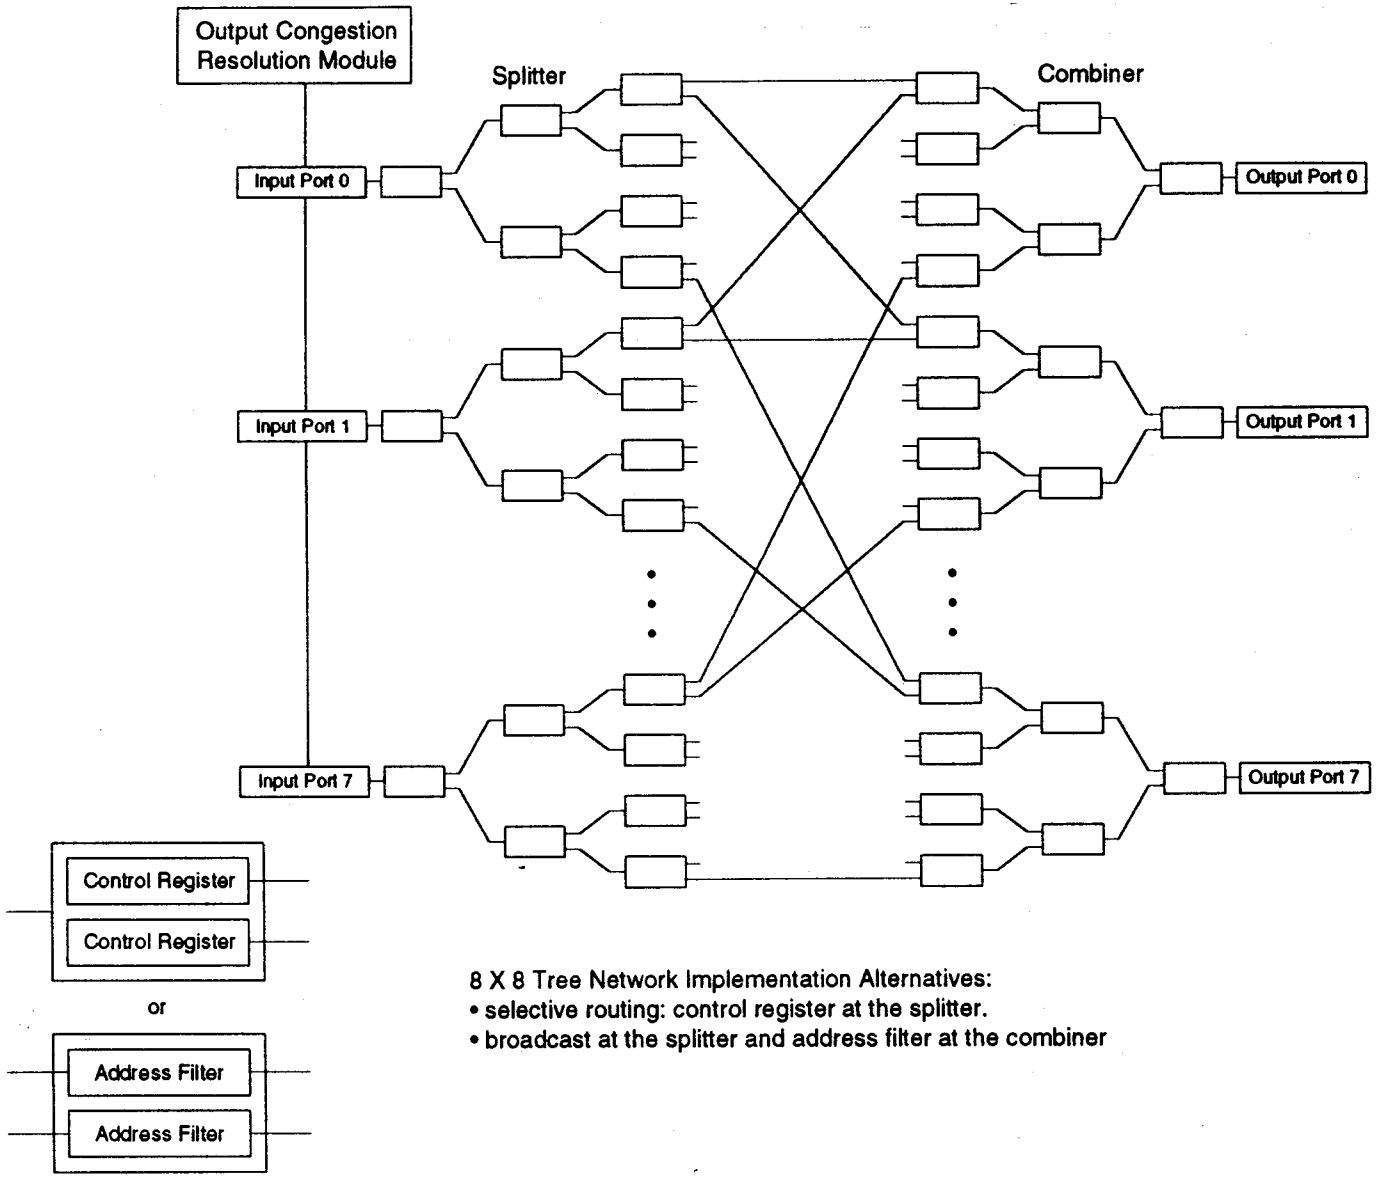

Figure 6-32. Tree Switch With Input Buffering

element of the splitter. The routing tag of the packet is compared with the control register to determine which output link the packet is to be switched to. In this scheme, the combiner can be passive, i.e., it is a simple multiplexer. The second scheme uses two address filters at each switching element of the first stage of the combiner. The packet coming to each input port is broadcasted to all the output links of the splitter. The filters at the combiner select only the packet with the correct destination address.

6.2.1.2.4 Self-Routing Crossbar Network with Input Buffering

A 2 x 2 switching element can be viewed as a crossbar. Due to the square growth feature of the number of crosspoints, a large crossbar switch is not cost effective; hence, the crossbar switch is often used as a building block for a larger switch such as the multistage interconnection network.

In this study, only the self-routing crossbar as shown in Figure 6-33 is considered. At each crosspoint, an address filter is implemented to extract the packet whose routing tag matches the output port address. The crossbar is internally nonblocking; however, it still suffers the output blocking problem. To resolve the output contention problem, input buffering with output port reservation scheme can be used.

The fast packet switch comparison for different buffering approaches and different switching fabrics is shown in Figure 6-34. Four possible schemes to improve the switch's performance are illustrated in this comparison table: larger checking depth (d), speedup factor (S), multiple parallel copies (p), and larger switching element size (D), where speedup factor is defined as the switch speed/link speed. The applicable schemes for each class are marked in the table. The possible blocking types and the possible packet transfer protocols for each class are also shown in the table. Format conversion is possible if the output port has a queue. Finally, the throughput, the fault-diagnosis complexity, and hardware complexity are also compared in this table. Note that the switch throughput and packet transfer delay can be improved by using a larger checking depth, a larger speedup factor, more parallel copies, and/or larger switching element size.

6.2.2 Shared-Memory Switch

In this approach, all the packets from different input lines are multiplexed into one TDM packet stream. The speed of the TDM stream is the sum of the incoming rates. These packets are stored sequentially in the common data memory, and the memory addresses of these packets are written sequentially into the control memory (see Figure 6-35). The self-routing addresses of the packets are sent to the address filters and activate the corresponding pointer array. The address of the control memory is written into the pointer array according to the self-routing address. All the packet control memory addresses whose packets go to the same output port are grouped into one array. The TDM output stream is formed by reading the packet out of the data memory using address obtained from the control memory while the address of the control memory is

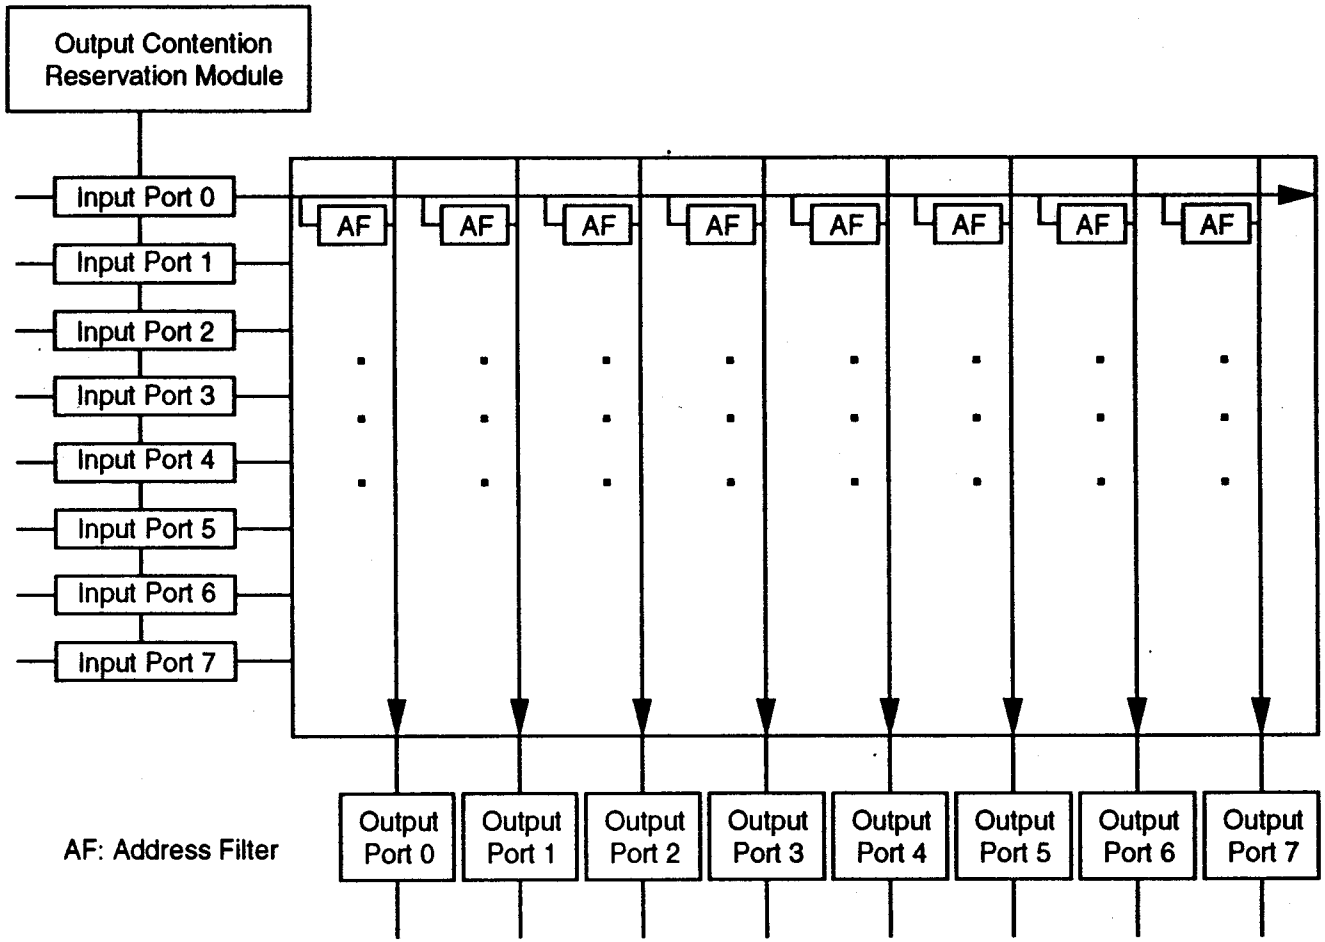

Figure 6-33. 8 x 8 Self-Routing Crossbar

	Input Buffering		Output Buffering		Input/Output Buffering	Internal Buffering
	Nonblocking Banyan-type Or Crossbar	Blocking Banyan-Type	Nonblocking Banyan-Type or Crossbar	Knockout Switch	Nonblocking Banyan-Type or Crossbar	
Larger Checking Depth (d)	applicable	applicable (with parallel copies)	no	no	applicable	applicable
Speedup Factor (S)	$1 \le S \le N$ (required output buffering)	$1 \le S \le N$ (required output buffering)	$S = N$	$S = 1$	$1 \le S \le N$	$1 \le S \le N$
Multiple Parallel Copies (p)	applicable	applicable	applicable	no	applicable	applicable (out-of-seq problem)
Larger Switching Element Size (D)	no	applicable	no	no	no	applicable
Head of Line Blocking	yes	yes	no	no	yes	yes
Internal Blocking	no	yes	no	no	no	yes
Output Contention	yes	yes	no	yes	yes	yes
Format Conversion (Speed Conversion)	no	no	yes	yes	yes	no
Packet Transfer Protocol	1 reservation with tokens 2 setup + transfer 3 transfer + retransmission	1 setup + transfer 2 transfer + retransmission	forward-and-store	forward-and-store	1 reservation with tokens 2 setup + transfer 3 transfer + retransmission	store-and-forward
Throughput	58%	< 58%	100%	100%	≥ 58% and ≤ 100%	≥ 58% (3 buffers or more)
Fault-Diagnosis	easy	easy	easy	easy	easy	hard
Hardware Complexity	low/medium	low	high	high	low/medium	high

• Throughput and packet transfer delay can be improved by adjusting the value of d, S, p, and/or D

Figure 6-34. Fast Packet Switch Comparison

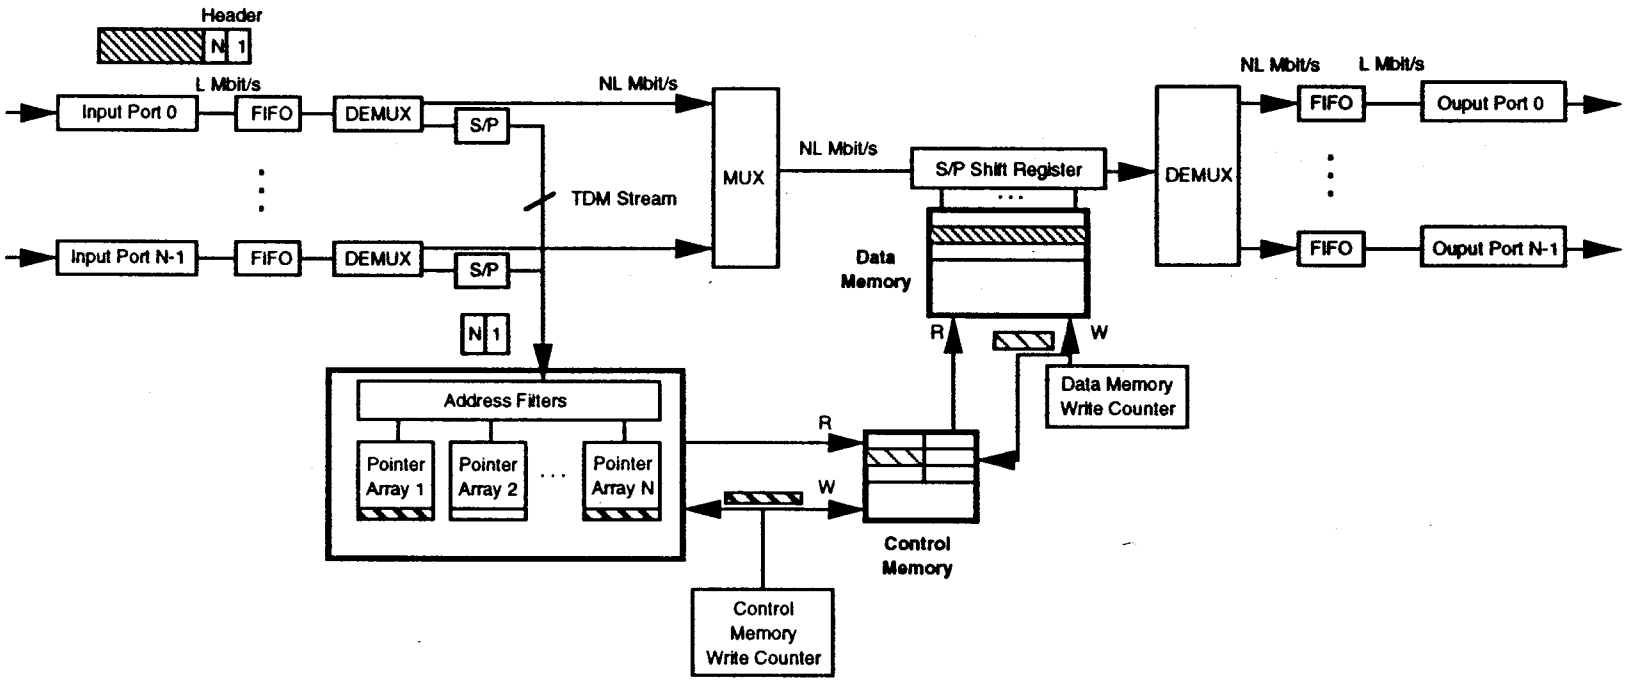

Figure 6-35. Shared-Memory Switch

obtained using the addresses of each array corresponding to each output port. And then the packets on the TDM stream are demultiplexer into different output ports.

Since the data memory and control memory are operated in random read fashion, it is not easy to keep track of the empty memory space after the packets have been read out from the memory. Link list implementation of the memory is required to efficiently use the memory space. Each time a packet is read out from the memory, the address of the packet location is recorded in a buffer pool. Each time a packet is written into the memory, an address from the pool is selected to store the packet.

Note that the shared-memory switch has no internal blocking and has no output contention.

The multicast operation is achieved using multiple writes to the pointer arraies since more than one pointer arraies will be activated at the same time for multicast connection.

The concern of this approach is the memory access time requirement for high speed application. This problem can be overcome using a wider parallel bus. The other concern is the memory size, which include the data memory and the control memory. The size of the memory is determined by the size of the switch, the utilization of each input line, and the uplink frame size.

6.2.3 Shared-Medium Switch

Two shared-medium switching architectures are considered in this subsection: bus and ring.

6.2.3.1 Bus

6.2.3.1.1 High Speed Bus with Distributed Output Memories

In this approach, all the packets from different input lines are multiplexed into a high-speed TDM bus. The speed of the TDM bus is the sum of the incoming rates. Within a physical bus, there are two separate logical buses; the first one is the packet bus and the second one is the address bus. There is an address filter attached with the address bus at each output port. This address filter is used to select the packet on the TDM bus.

Since there are chances that more than one packets will arrive at one output port, buffering is required at the output ports (see Figure 6-36). The distributed output memory approach is also suitable for the situation that format conversion at the output port is necessary. This is important in the satellite environment; for example, the format conversion between the uplink TDMA burst mode and the downlink TDM frame mode. Compared with the shared-memory approach, these approaches completely divide the memory among the output ports.

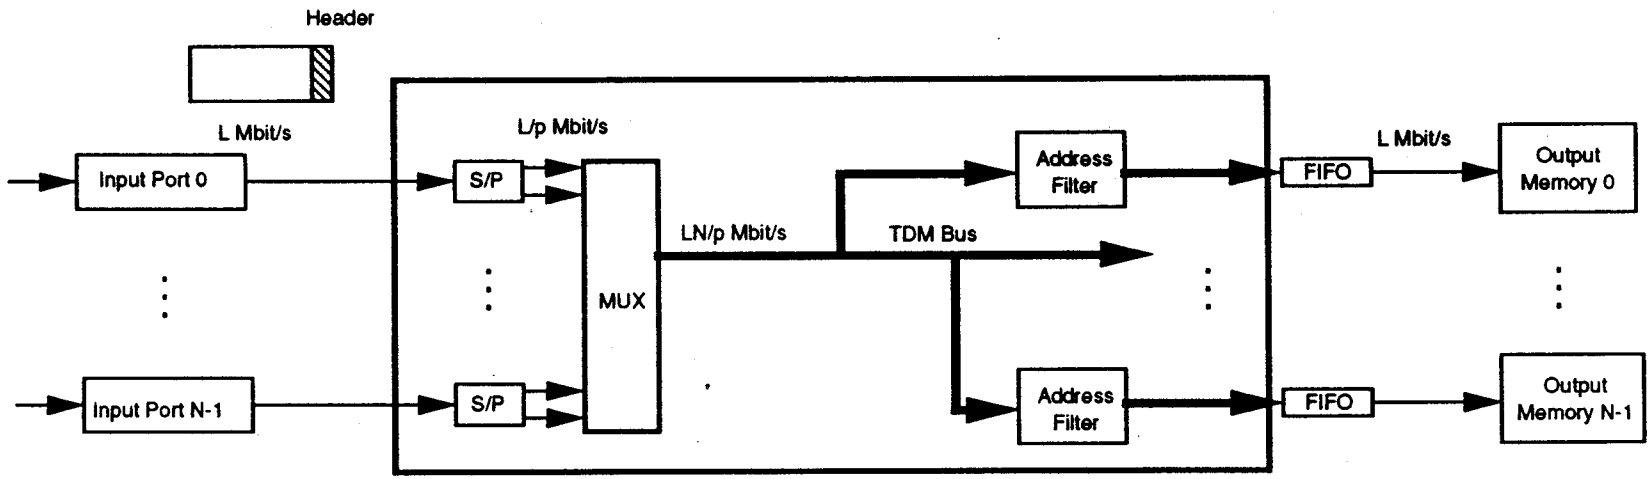

Figure 6-36. TDM Bus with Distributed Output Memories

Note that the TDM bus switch has no internal blocking and no output contention.

Since the TDM bus has an inherent broadcast capability, the multicast operation can be achieved easily.

The general concerns with the bus approach are the bus speed, the memory access time, and the memory size. The bus speed and the memory access time problem can be tackled using the bit slice approach, which will be introduced latter.

The packet filter structure depends on the addressing scheme. For an N-bit addressing scheme (where one bit is for one output port), the packet filter is a simple latch circuit. For a virtual circuit addressing scheme, the packet filter is implemented using a

comparator. For a multicast group address, the packet filter is implemented using a RAM. The content of the RAM is 1 if the output port address is in the multicast group.

6.2.3.1.2 High Speed Bus with Distributed Input Memories

The buffering of the arriving packets is performed in the input ports. As shown in Figure 6-37, the distributed input memory approach is suitable for consistent frame format between input lines and output lines. Since there is no output buffer, the output contention has to be resolved at the input port. Hence, an output port reservation device is used among the input ports.

6.2.3.1.3 Bit Slice Bus

In this approach, a S/P conversion is performed for each incoming line and each bit out of n bits of the packet data is fed into one bus switch, where n is the size of the parallel conversion (see Figure 6-38). In a sense, n bus switches are stacked so that one bus switch handles one bit out of n bits of the incoming line. Note that a complete packet routing tag is required for each switch. The result is that the bus speed and memory access time requirement are reduced by a factor of n. To further reduce the requirement, each bus switch can use a parallel bus for multiplexing instead of a serial bus.

The bit-sliced operation can be changed to block-sliced or packet sliced operation.

6.2.3.1.4 Multistage Structure Bus

The multistage approach is to use a bus switch as a basic switching element and to connect these elements into a Clos-type network or a Benes-type network (see Figure 6-39). The basic principle behind the multistage structure is to create multiple paths between any input and output pair. During the call setup phase, a path is chosen for one virtual circuit so that all the packets with the same virtual circuit follow the same path.

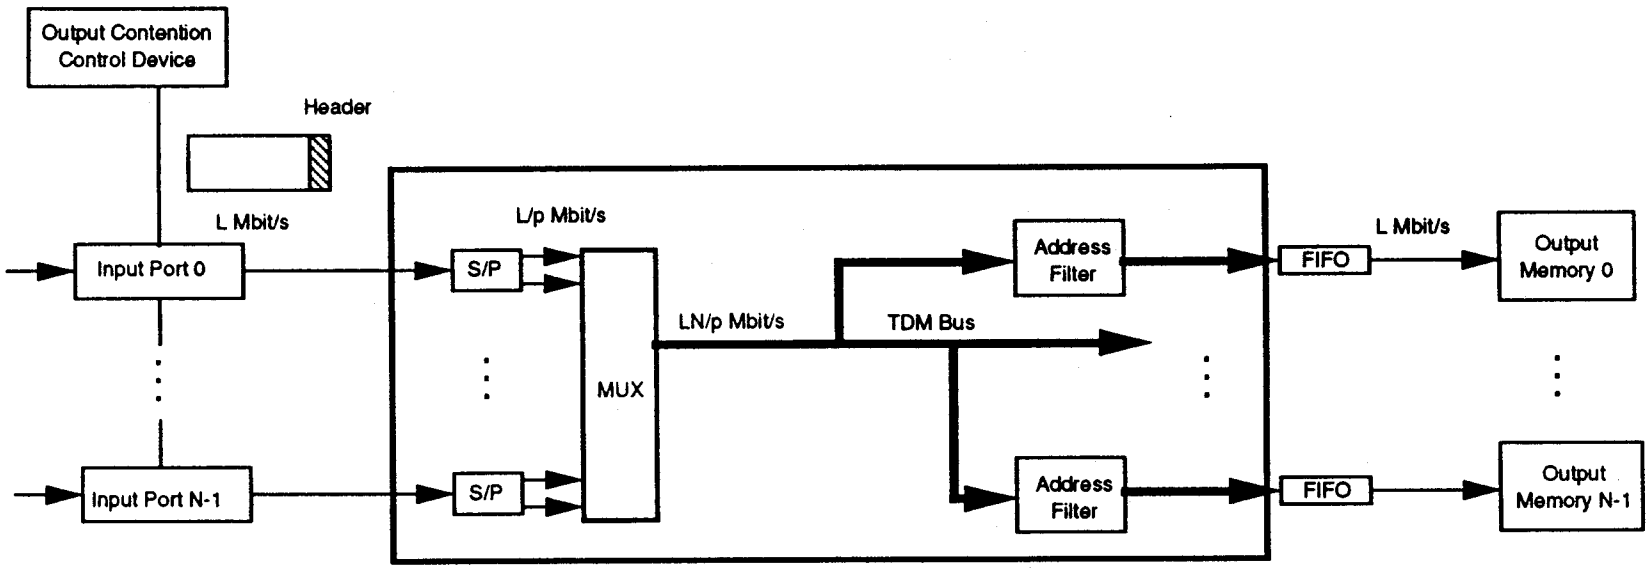

Figure 6-37. TDM Bus with Distributed Input Memories

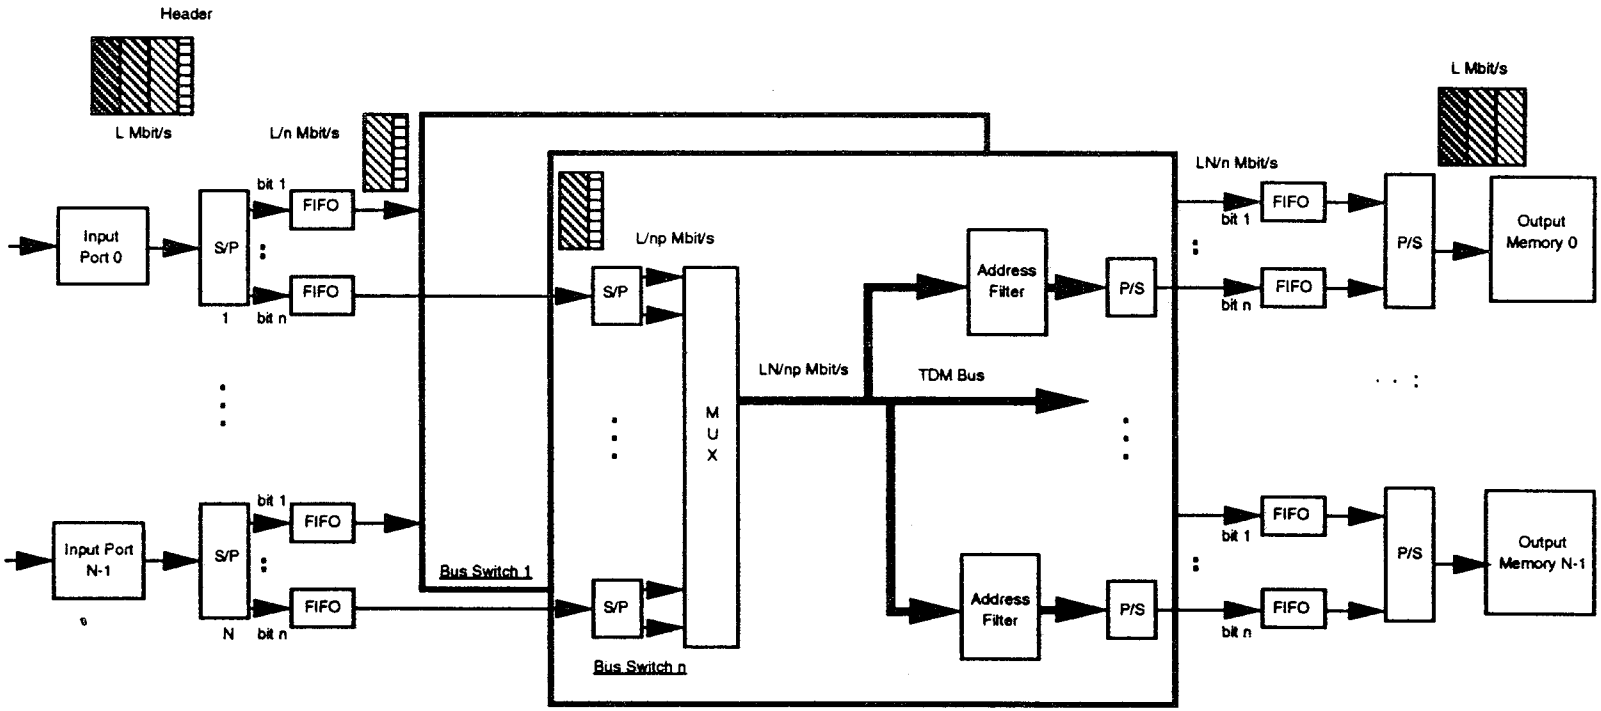

Figure 6-38. Bit-Sliced Bus with Distributed Output Memories

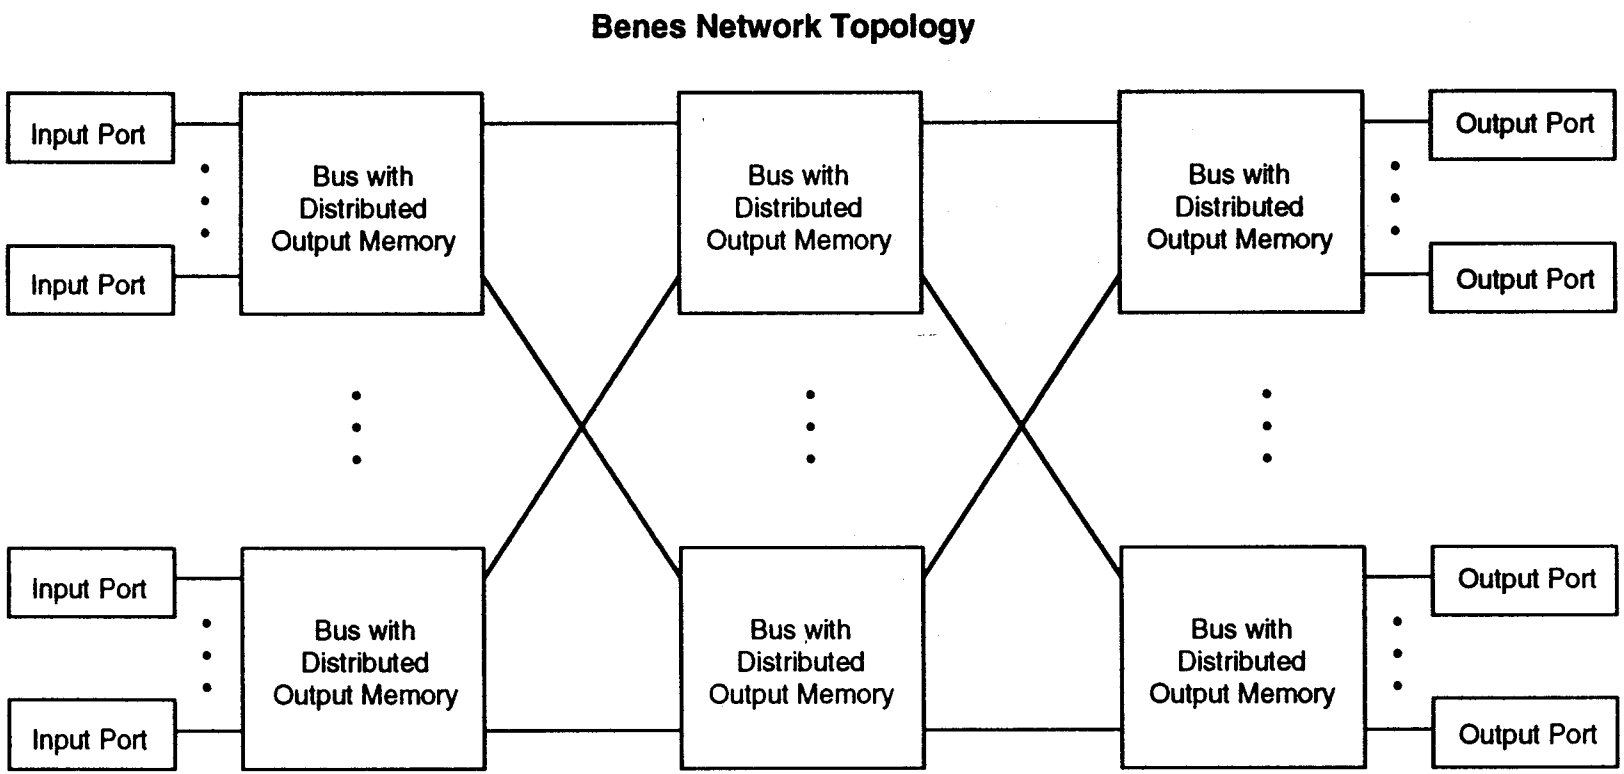

Figure 6-39. Multistage Switching Network

6.2.3.2 Ring Switch

The media access scheme of the ring switch decides the system performance. Several access schemes have been proposed: the slotted ring, the token passing, and the distributed queue dual bus (DQDB).

The throughput of the ring is limited since the transfer of packets is not performed in parallel. As shown in Figure 6-40, to increase the switch capacity, multiple rings are used to interconnect the switching modules.

In conclusion, the potential switch architectures for different capacity are listed as follows.

- low to medium capacity (< 2 Gbit/s)
  - shared-memory
  - TDM bus (including optic ring)
- high capacity (10's Gbit/s)
  - space division switch
- ultrahigh capacity (e.g. 1 Tbit/s)
  - optical switch

6.2.4 On-Board Switch Design Issues

6.2.4.1 Uplink Grouping

The design concept is that the switch size does not have to be consistent with the number of uplink beams and number of downlink beams. If the number of beams is large (e.g. 100) and the link speed is low (e.g. 1 Mbit/s), then it is not proper to use a switch whose size is the same as the number of beams on-board the satellite due to power, mass, weight, and redundancy consideration. Therefore, the low-speed uplink beams can be multiplexed into a high-speed TDM bus on-board and be fed into the switch input port (see Figure 6-41). The size of switch is reduced to N/m, where N is the total number of beams and m is the number of beams being multiplexed. The speed of the switch is increased to mS_I, where S_I is the uplink beam speed. Every packet can carry one physical address and the physical addressed is translated to logical address on-board. Or every packet can carry two addresses: one logical address and one physical address. The logical address is used to identify the output port of the switch while the physical address is used to identify the downlink beam. The switch uses the logical address to set up a path from the input port to the destined output port. There is a demultiplexer after the output port. The demultiplexer uses the physical address to route the packet to the proper downlink beam.

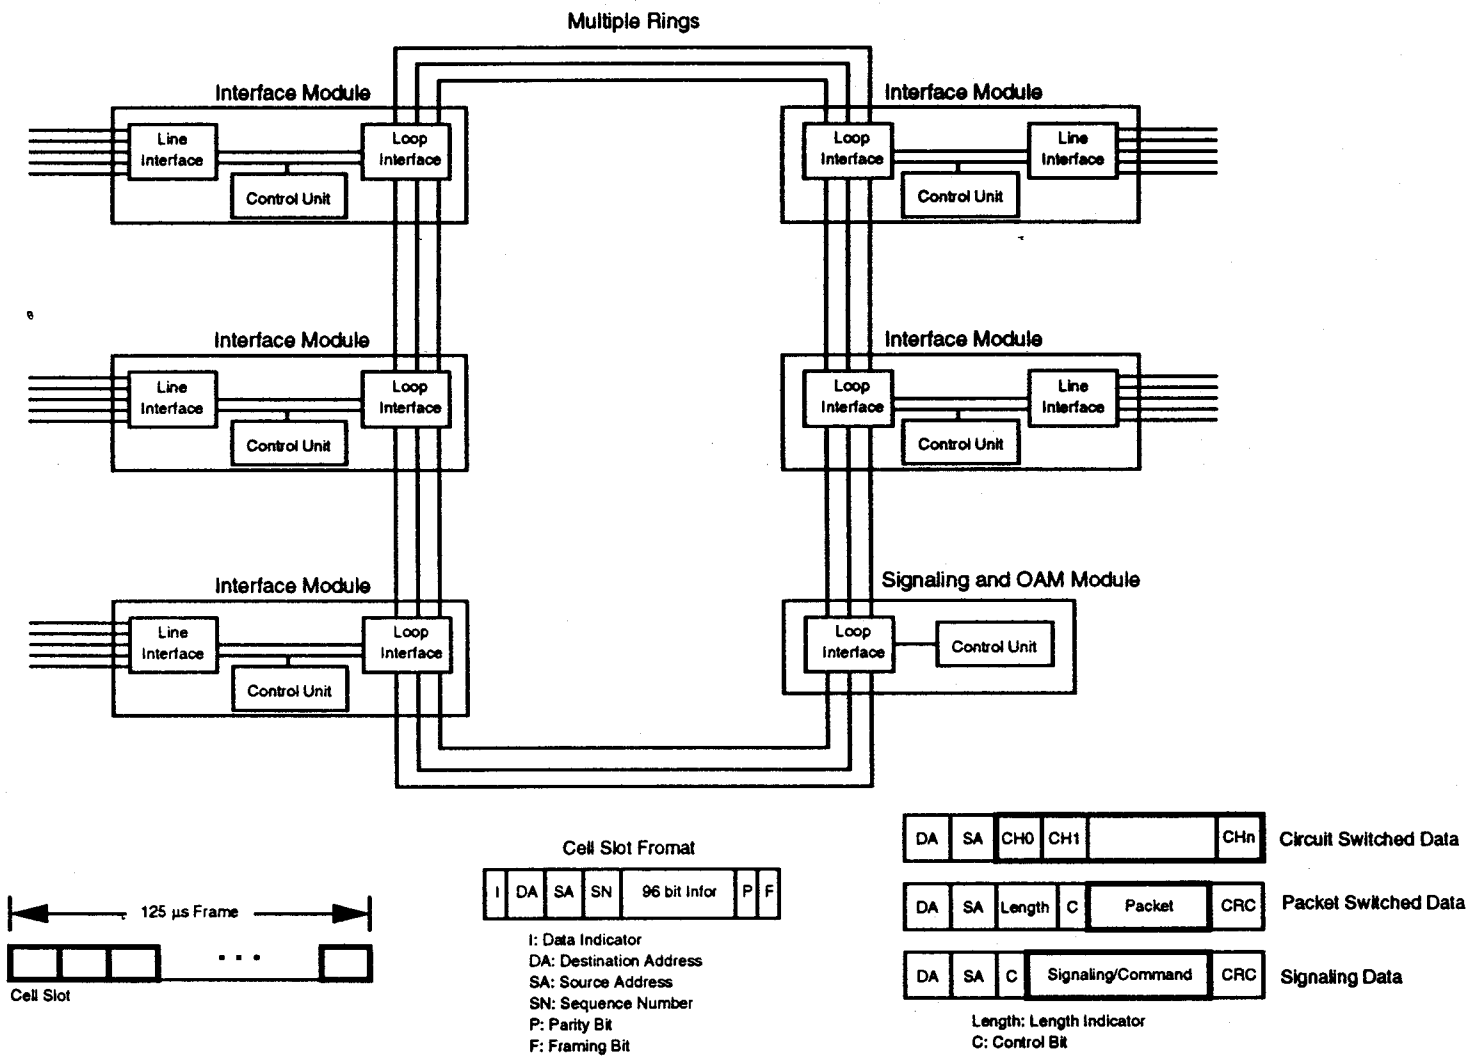

- Circuit switched data has priority over packet and signaling data in using the cell slots.
- Periodic frame transmission of circuit switched data.

Figure 6-40. Multiple Ring Switch

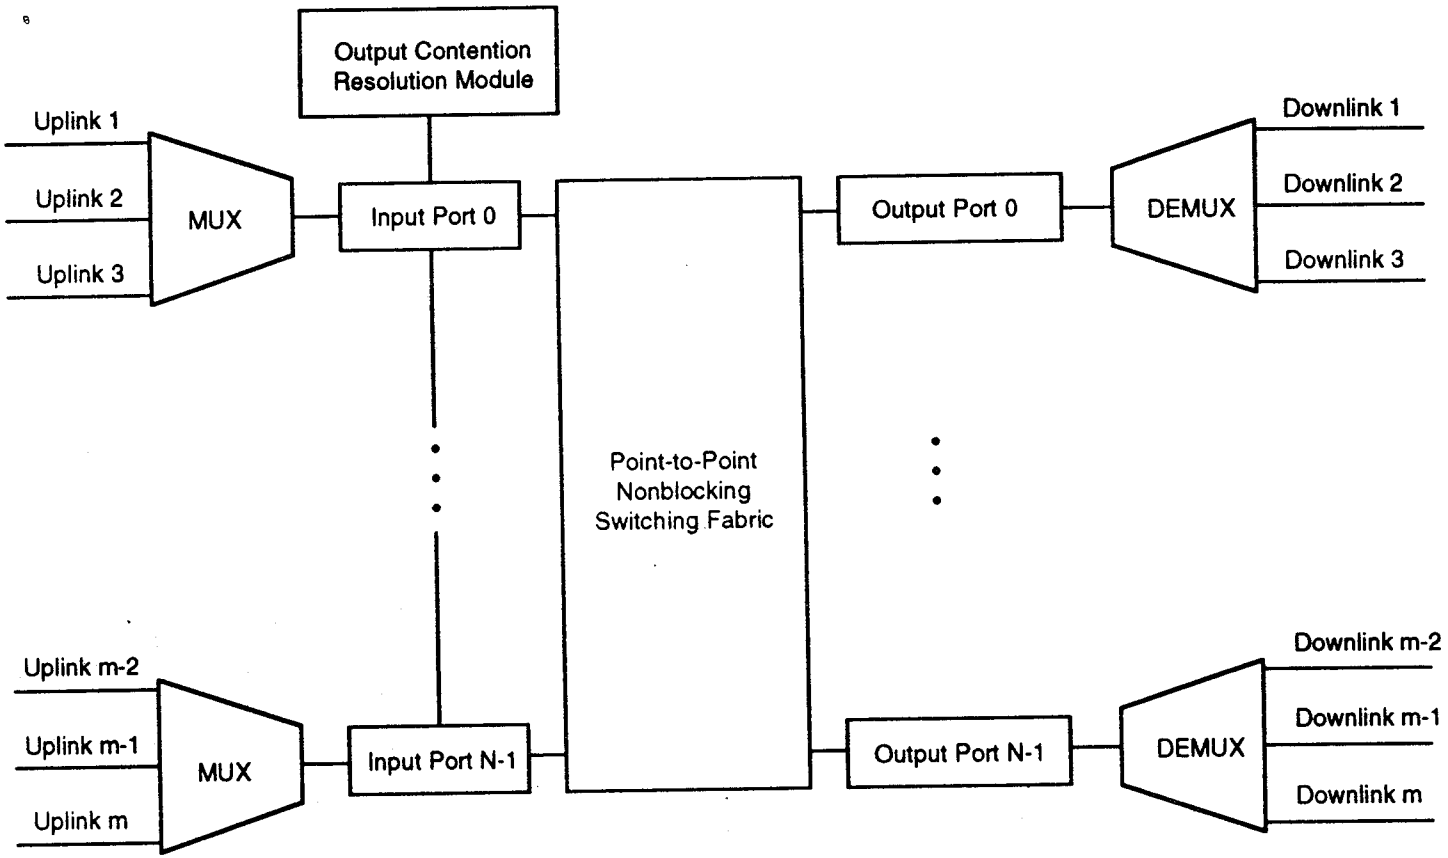

Figure 6-41. Upllink Beam Grouping

6.2.4.2 Downlink Grouping

If the number of uplink beams is small and the link speed is very high (e.g. 600 Mbit/s), then it is not easy to implement such a high-speed switch using low power device. Therefore, the high-speed uplink beam can be demultiplexed first into several lower-speed links, and these links are fed into input ports individually (see Figure 6-42). The size of switch is enlarged to N*m, where N is the total number of beams and m is the number of output links of a demultiplexer. The speed of the switch is decreased to S_l/m, where S_l is the uplink beam speed. Every packet only has to carry only one physical address (for the downlink beam). This physical address can be translated into logical address at the input port on-board. The logical address is used to set up a path within the switch. The main reason that the logical address should be translated on-board the satellite is that output contention problem of the switch is minimized using the downlink group concept. Suppose m = 3 and consider downlink 1. Whether a packet is routed to output port 0, output port 1, or output port 2, it will be multiplexed to the same downlink beam. Hence, by performing on-board translation and a careful design of the output port contention resolution scheme, the downlink beam efficiency can be largely increased by reducing the output port contention. One way of implementing this on-board translation and output port contention resolution is introduced as follows. Assume the switching fabric is point-to-point nonblocking and input buffering is employed. The token reservation scheme is used as the output port contention resolution scheme. There are N*m tokens in the token stream, where every m tokens is grouped into a super token for each downlink beam. All the packets destined to the same downlink beam can remove any token (any output port) in the super token (the downlink beam). The token removed by the input port will be the routing tag through the switch for the packet at the next time slot.

6.3 Multicast/Broadcast Operation

In general, there are six approaches to design a multicast packet switch. The first is to use a shared-medium switch. The second is to use a shared-memory approach. The third approach is to use the store-and-forward nature of the packet switch, i.e., send the multicast packet one by one from the input port. The fourth approach is to use a copy network and a routing network. The fifth approach is to use a multicast module at the output port and this multicast module is responsible for sending the multicast packet to the destinations. The sixth approach is to use a multicast banyan network or a multicast tree network.

6.3.1 Shared-Medium Approach: Inherent Multicast Capability

Due to the broadcast nature of the shared-medium, the multicast operation can be achieved without packet duplication. As previously mentioned, the general concerns of the shared-medium approach are the high bus speed. The bus speed requirement can

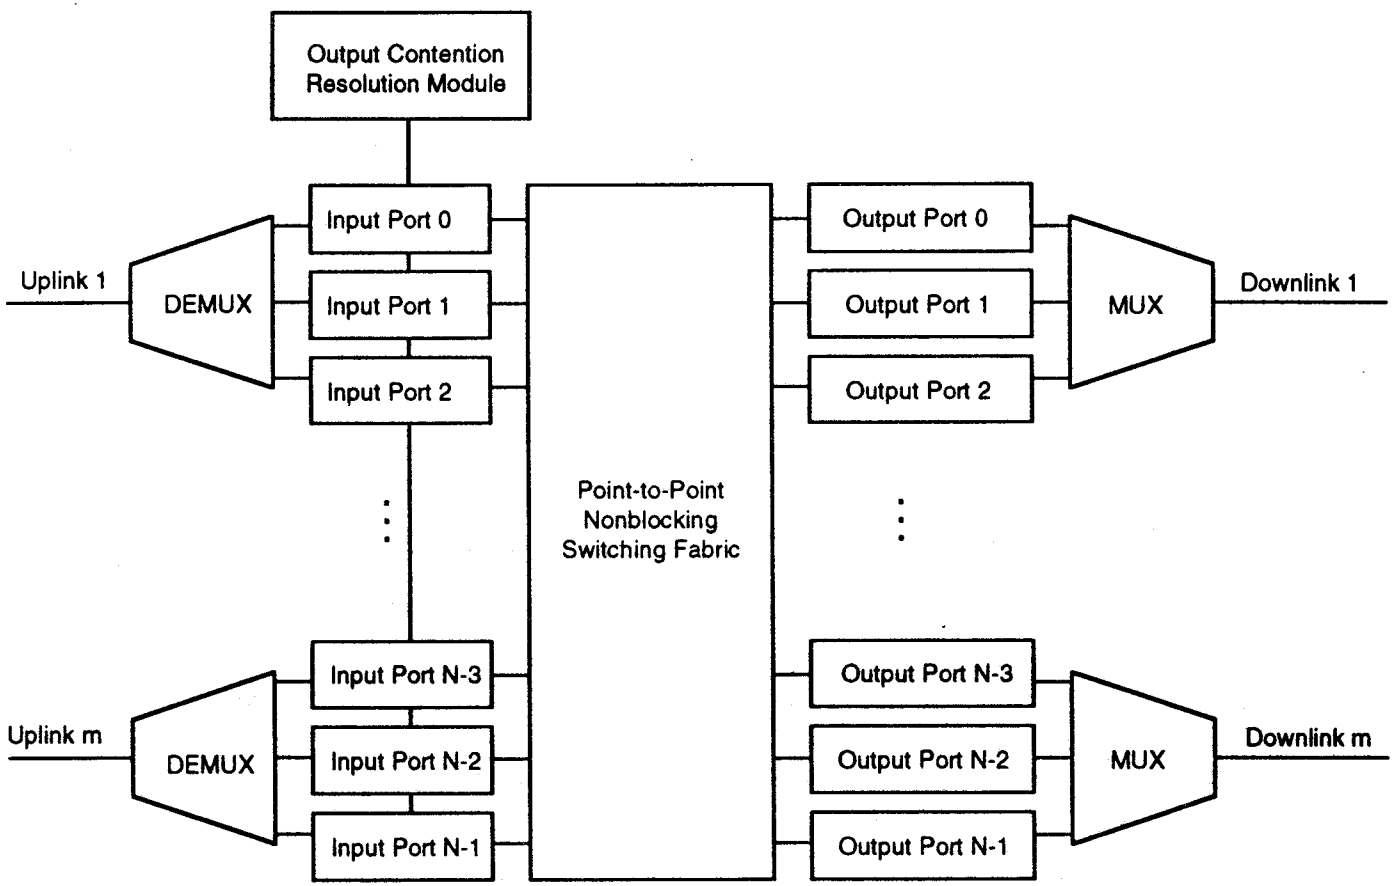

Figure 6-42. Downlink Beam Grouping

be reduced by using a wider parallel bus. Nevertheless, it is not appropriate and cost effective to use the shared-medium switch, if the capacity is above 5 Gbit/s.

6.3.2 Shared-Memory Approach

Contrast with the shared-medium approach, multicast operation in the shared-memory is not one-shot operation. Multiple read operation is necessary for the multicast operation. Since the speed of the shared-buffer approach is limited by the memory access time, it is not appropriate and cost effective to use the shared-buffered switch, if the capacity is above 5 Gbit/s.

6.3.3 Store-and-Forward at the Input Port

In this approach, the multicast operation is achieved by sending the multicast packet one by one from the input port (see Figure 6-43). The advantage of this approach is that a point-to-point switch can be used as a multicast switch; hence, the hardware cost for building a multicast switch is minimal. The disadvantages of this approach are the long delay due to the serial transfer of the multicast packet and serious congestion if the number of duplication is large.

The above approach is feasible and very cost effective if the amount of multicast traffic is small and the number of duplication of each multicast packet is small. Otherwise, serial packet duplication at the input port has to be modified so that parallel duplication is possible.

One of the methods for parallel duplication is to send the multicast packet to adjacent input ports so that packet duplication can be achieved in parallel by many input ports, and each input port only handles a portion of the multicast traffic. In a sense, a virtual copy network is implemented among the input ports using a bus structure. It can be envisioned that this procedure involves a lot of handshaking among different input ports.

If the input and output ports are combined into one module, the switching fabric can be used as a copy network. Hence, the input port can send the multicast packet to several output ports, and the output ports can relay this multicast packet to the accompanying input ports. The locations of the input ports which are used to duplicate the packets are decided at the call setup time.

6.3.4 Copy Network Plus Routing Network

In this approach, packets from different input lines are duplicated using a space-division copy network [6-13][6-14]. After the packet duplication, the point-to-point routing network routes the packets to the destinations (see Figure 6-44). The difference between the Turner's approach [6-13] and Lee's approach [6-14] is that Lee's copy network is nonblocking and Turner's copy network is blocking. Since Turner's copy network is blocking, the switching elements in the copy network are buffered.

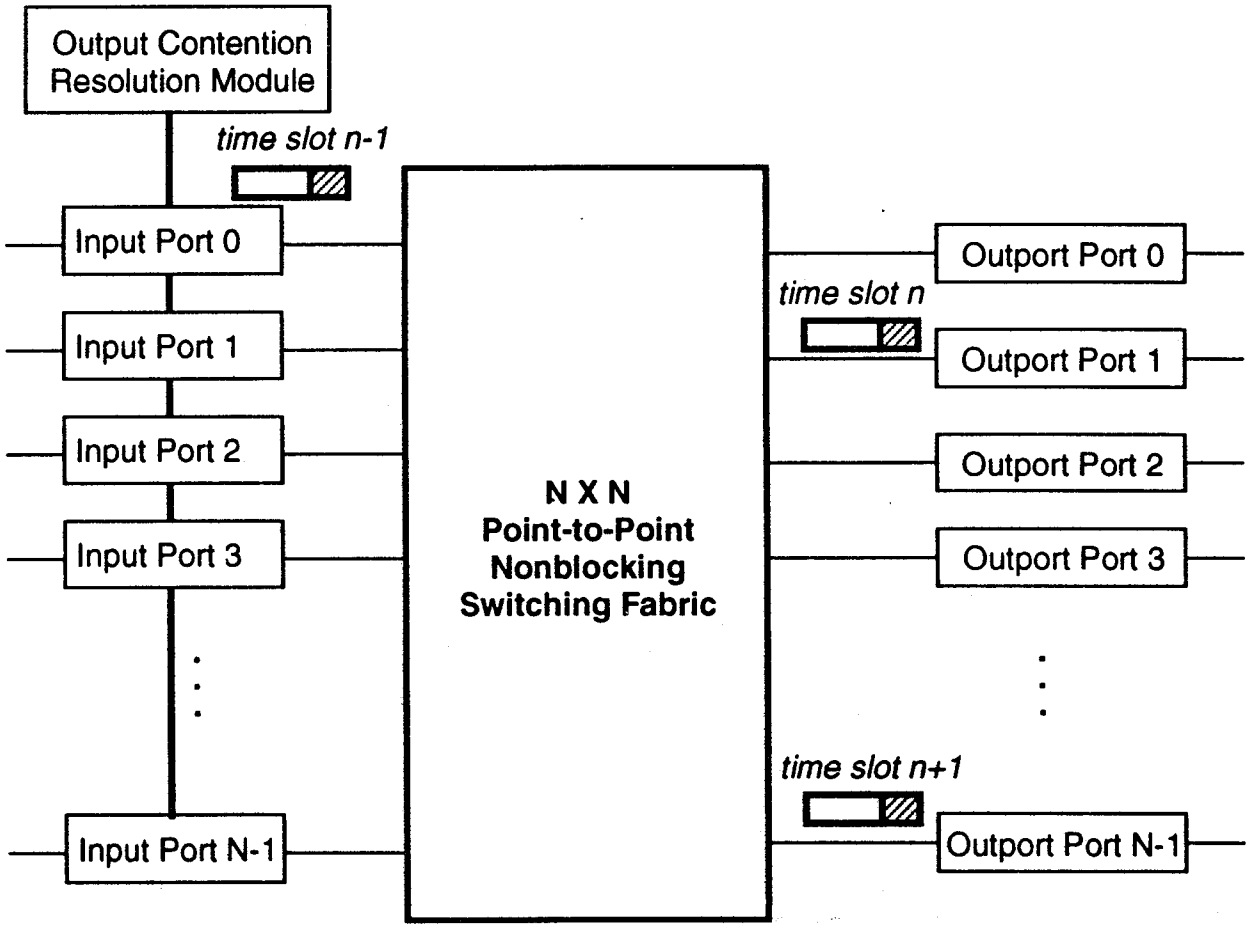

- The multicast packet is sent to different output ports one by one from the input port.

Figure 6-43. Store-and-Forward at the Input Ports

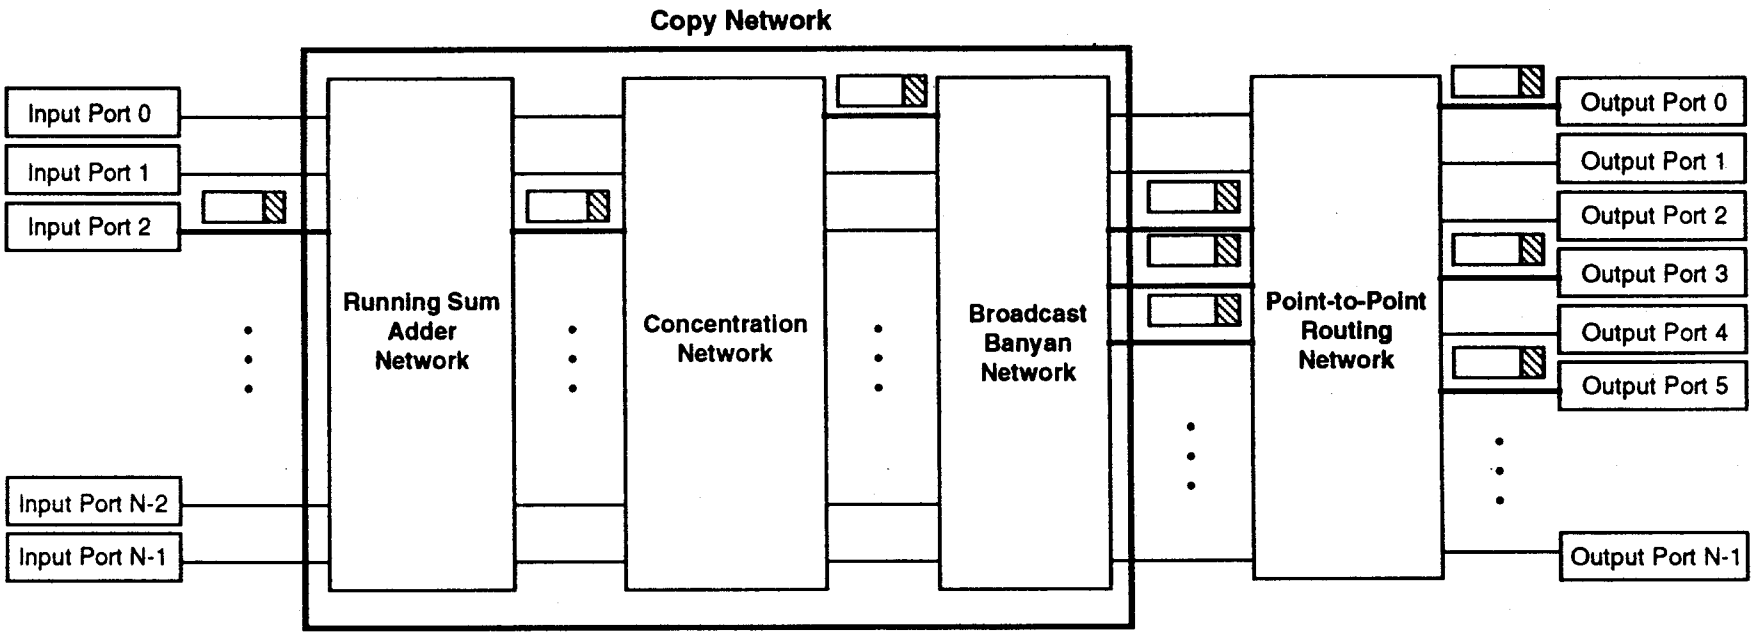

Figure 6-44. Copy Network Plus Routing Network

Evidently, Lee's copy network is superior than Turner's copy network in several aspects: non-buffered-banyan network, nonblocking property, and constant latency time. For these reasons, the Lee's copy network is used as the representative of this approach.

The first operation for the copy network is to duplicate the exact number of copies for each multicast packet. This requires the incoming packets to carry a copy index in the header. For the copy network to be nonblocking, the copy index of each multicast packet has to be translated into an address interval; this operation is performed on the incoming packets sequentially from the top to the bottom. This procedure can be implemented using a running adder network and an address interval encoder. After the multicast packets with proper address intervals have been generated, a concentration network is necessary to concentrate these packets to satisfy the nonblocking condition of the copy network. Since these packets' address intervals are monotonically increasing and they are concentrated, a banyan network can duplicate these packets without any blocking. After the copies have been generated, a table is necessary to translate the header of each copy to the destination address so that the routing network can route the packet to the destined output port.

The are several disadvantages using this approach. The first is the delay incurred for every packet (including unicast and multicast packets) passing through the copy network and the routing network, and the hardware complexity incurred by the copy network. Since the duplicated packets might use different input ports of the point-to-point switching network for routing at different time, these packets might be transmitted out-of-sequence due to different levels of input congestion at different time. It is noted that if the routing network is nonblocking, there is no out-of-sequence problem. For the same reason, if the positions of the duplicated packets at the input ports of the point-to-point routing network are changed very often due to connection addition or deletion, the translation table has to be updated very often also. This puts a large burden on the signaling processor. This situation becomes very severe for video distribution services since viewers are most likely change the channels very often. Also if the number of duplication each time exceeds the size of the switch, the copy network cannot handle this situation, and the overflowed packets are simply dropped. Based on the above discussion, the copy network is cost-effective if the amount of the multicast traffic is large.

6.3.5 Multicast Module at the Output Port

In this approach, there are multiple multicast modules at the output ports. All the multicast packets are relayed to these multicast modules first. And then the multicast modules send the multicast packet to the destined output ports through a proper multicast interconnection network (see Figure 6-45). The number of multicast modules required depends on the amount of multicast traffic. The multicast knockout switch uses a similar approach [6-12] as shown in Figure 6-46. The knockout switch uses the bus approach to interconnect the inputs and outputs. There are N broadcast buses in the switch for the point-to-point applications. If there are M multicast modules, then

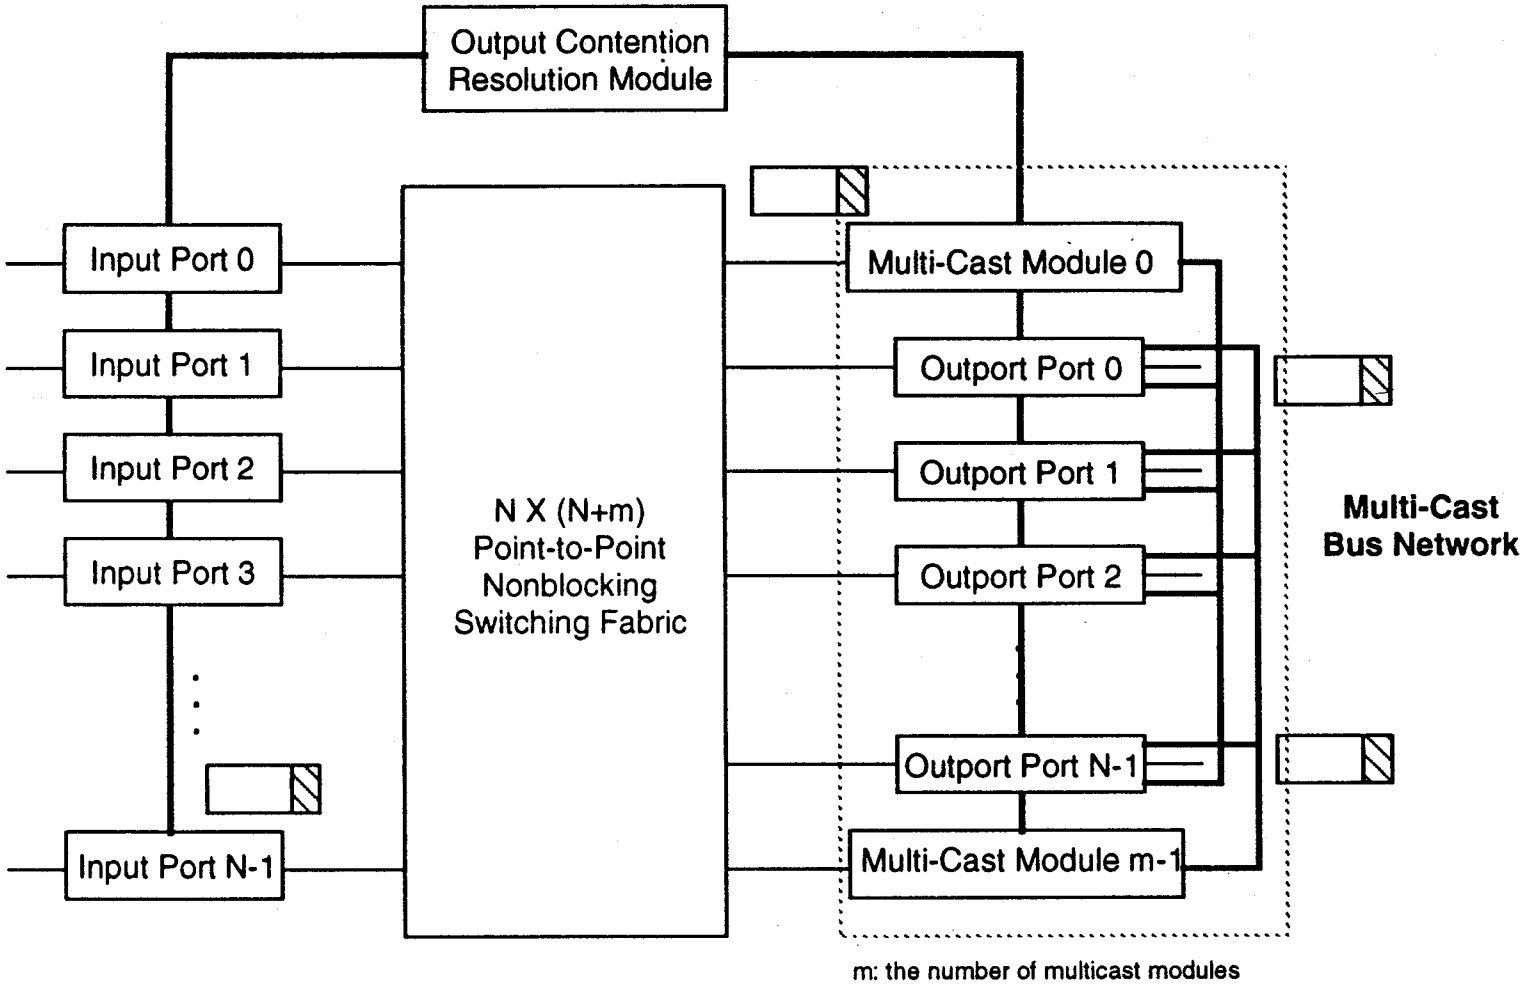

- Packet transfer sequence for point-to-point packets is from one input port to one output port.
- Packet transfer sequence for a multicast packet is from one input port to one multicast module. And then the multicast module sends this multicast packet to different output ports.

Figure 6-45. Multicast Modules at the Ouput Ports

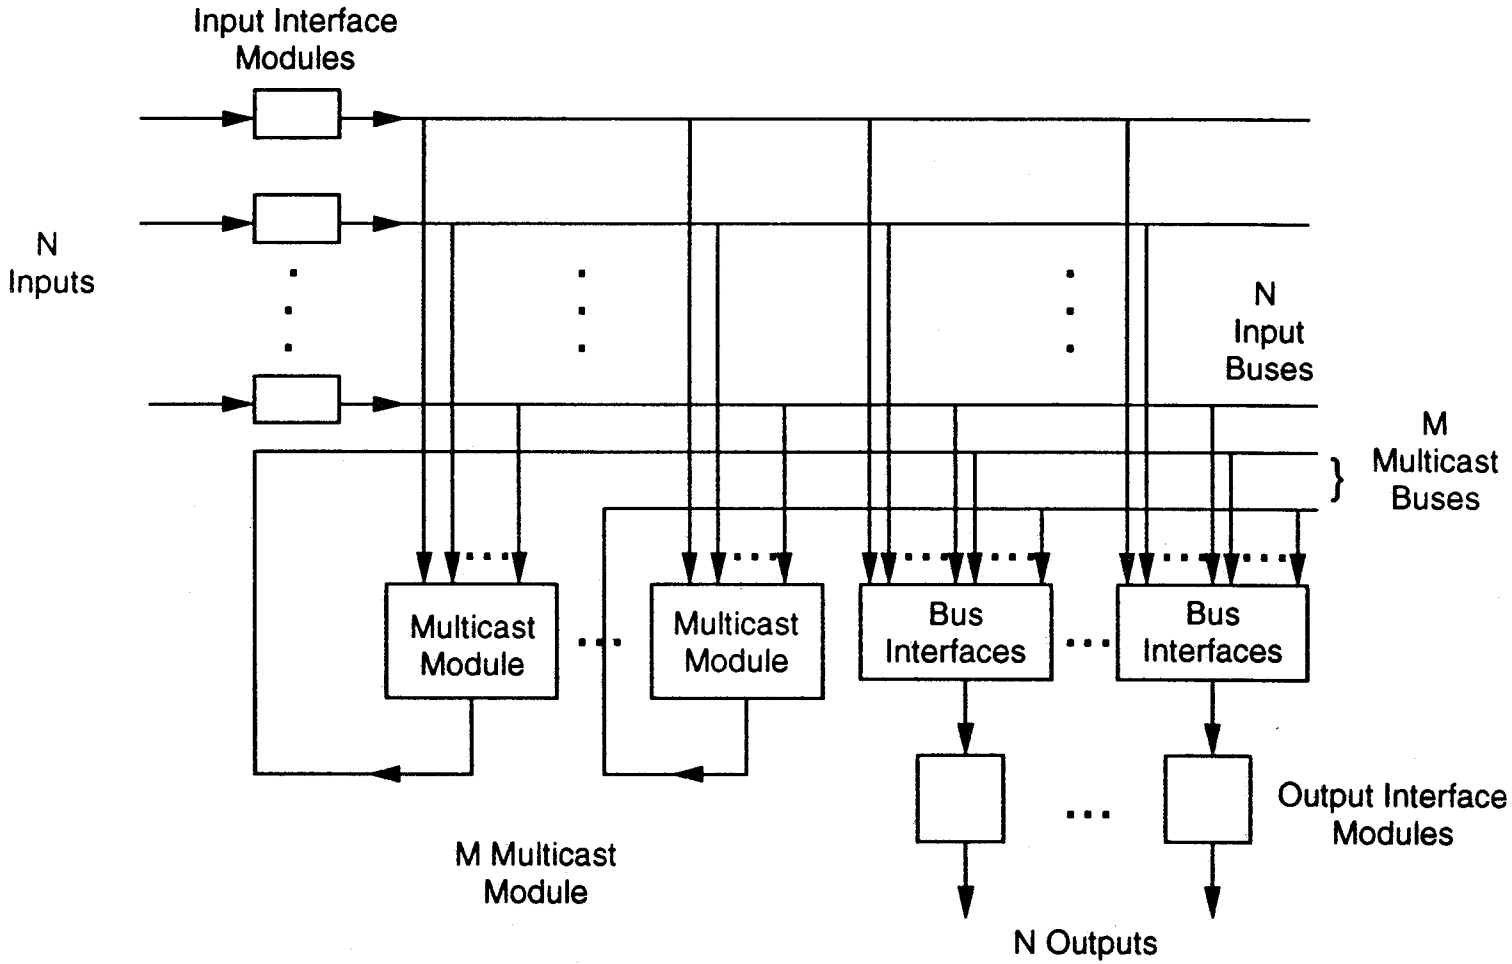

Figure 6-46. Multicast Knockout Switch

the total number of buses is N + M and the size of the switch becomes N x (N+M). There are (N+M) filters at each bus interface of the output port, where each filter is for one input; hence, the total number of filters for the switch are N^2 + NM. It can be seen that the complexity of the bus interface is very high. The desired point-to-point switching fabric is the banyan-type network, which is assumed to be the switching fabric in the discussion below. If the banyan-type network is used as the switching fabric, then the number of filters necessary for the bus interface at each output port is only M, where M is the number of multicast modules.

The disadvantage of this approach is that some of the output ports are allocated for the multicast modules. The result is the switch may become a nonsymmetric switch if a symmetric switch is chosen as the point-to-point switching fabric. However, this might not be a problem for an on-board switch since, usually, the number of uplink beams is larger than the number of downlink beams. If symmetry is a requirement of the switch, the solution is to use a larger switch so that the number of inputs and outputs can be made equal. The other disadvantage is that the output port has to provide interfaces for two packet transfer protocols: one for the point-to-point network and the other for the multicast network.

6.3.6 Multicast Banyan or Tree Networks

The multicast banyan network (see Figure 6-47) or tree network can duplicate and route the multicast packet simultaneously, which is very attractive. In a multicast banyan network [6-15], each output port is associated with one bit in the routing tag. Hence the size of the routing tag of a multicast packet is N. There are two registers, one for each output, holding control bits at the switching element. The operation of each switching element is to AND the routing tag of the packet and the control bits at each register. If the result of the AND operation is 1, a copy of the packet is send to the output. Although the multicast banyan is cost-effective, it has the disadvantages of a long routing tag, large registers in the switching element, and serious internal blocking problem.

It is found that the multicast banyan network can become a nonblocking multicast switching network by using a sorting network at every stage of the multicast banyan network (see Figure 6-48). The invention entitled "Nonblocking Multicast Fast Packet/Circuit Switching Networks" are in the process of applying for a patent [6-16]. These switching networks route the multicast packets with proper multicast routing tags to the destinations in a distributive and nonblocking fashion.

Although the tree network is internal nonblocking, it has the disadvantages of a long routing tag, large registers in the switching element, and high cost if the size of the switch is large.

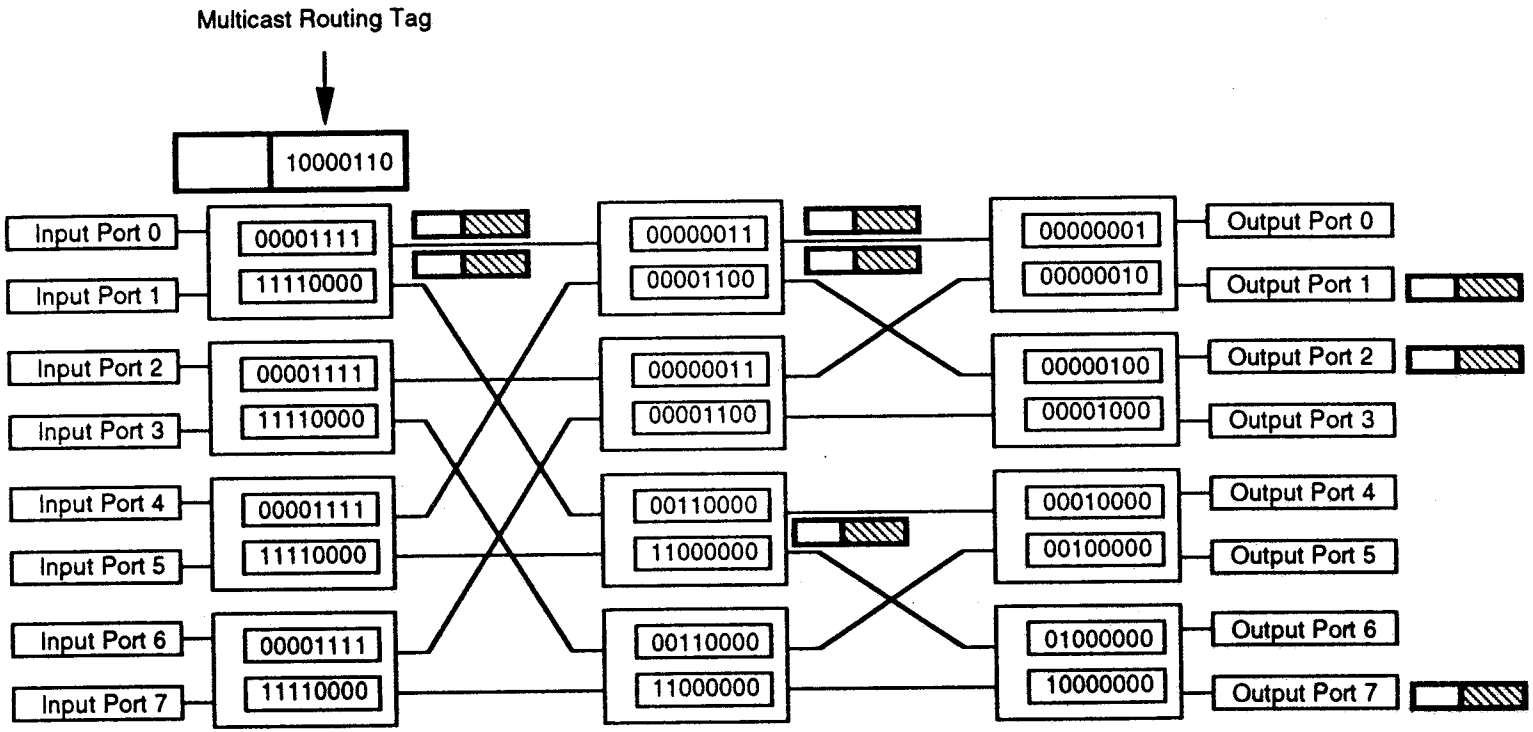

Figure 6-47. Multicast Banyan Switch

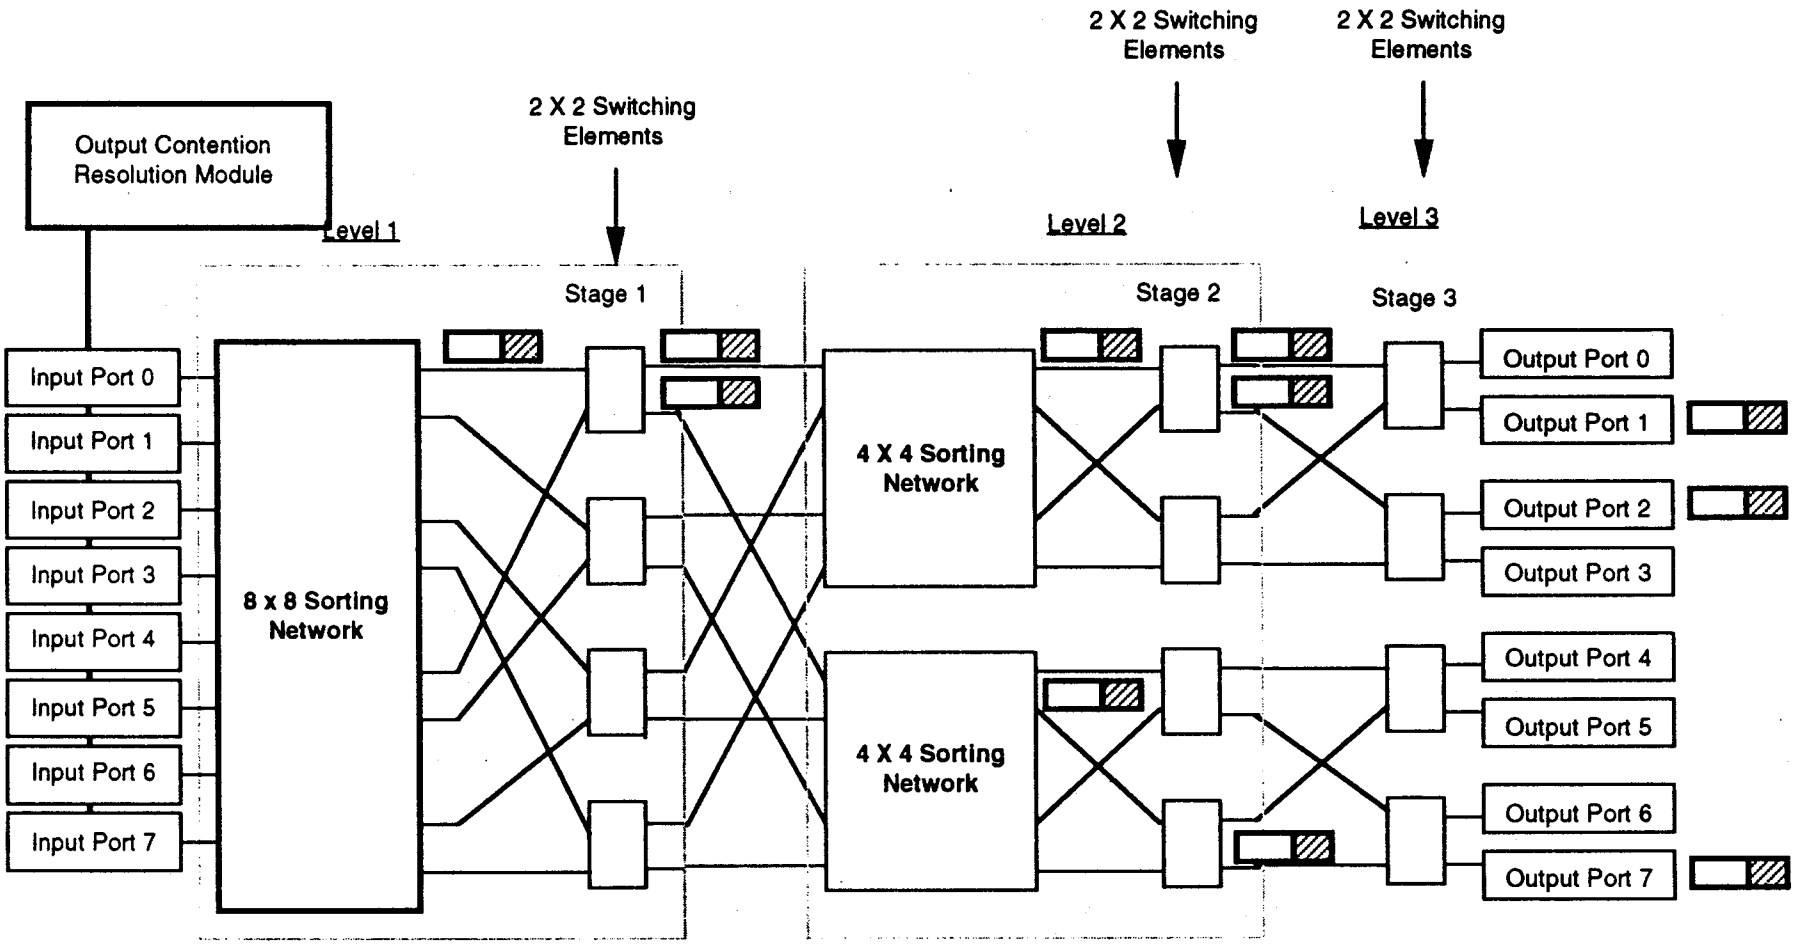

Figure 6-48. Point-to-Multipoint Banyan Network (Point-to-Multipoint Nonblocking Network)

The multicast switching architecture comparison is provided in Table 6-1.

Table 6-1. Multicast Switch Trade-Off

	SWITCHING ARCHITECTURE					
FUNCTION	A	B	C	D	E	F
• high speed	XXX	XXX	XXX	X	XXX	X
• delay	XX	XXX	X	XX	XX	XX
• packet transfer protocol complexity	XX	XXX	XXX	XX	X	XX
• hardware complexity	XX	XX	XXX	XX	X	XX
• self-routing	y	y	y	y	y	y
• nonblocking switching fabric	y	y	y	y	y	y
• growth capability	y	y	y	y	y	y
• topology symmetry	n	y	y	y	y	y
• packet sequencing	y	y	y	y	n	y
• fault-tolerance	y (if m > 1)	n	n	n	n	n

A: Multicast module at the output port
B: Modified multicast banyan network
C: Store-and-forward at the input port
D: Shared-medium
E: Copy network plus routing network
F: Shared-memory

XXX: Good
XX: Fair
X: Poor
y: Yes
n: No
m: the number of multicast modules

6.4 Priority Control

To provide different levels of QOS for different classes of services, priority control of the switch is necessary. For ATM cells, priority control is performed using the CLP bit in the ATM header. For the satellite virtual packet (SVP), a QOS field is provided so that the traffic can be segregated and different control can be applied to different classes. There are two types of priority control: priority control for scheduling the packet transfer and priority control for congestion control. In addition to scheduling the the packet transfer using priority control, two other aspects have to be considered. The first is the sequence the packets entering a queue, i.e., the multiplexing scheme. The second is the sharing policies of buffer among different classes of packets.

6.4.1 Scheduling Method (Queueing Discipline)

In general, there are two scheduling methods: a static priority scheme and a dynamic priority scheme.

6.4.1.1 Static Priority Scheme

In this scheme, the high priority packets are always scheduled (served) first. This scheme has been analyzed and the results are presented in Appendix C. Although the performance of the high priority packets is good, the performance of the low priority cell is significantly degraded if the amount of high priority packets is large.

6.4.1.2 Dynamic Priority Scheme

In this scheme, the scheduling method is dynamic according to the status of the queue. During the normal operation, the high priority packets are always scheduled first. When the number of low-priority packets in the queue exceeds a certain threshold, then the low priority packets are scheduled first until the number is decreased below the threshold.

6.4.2 Sorting Implementation for Multiplexing Operation

In this approach, the packets with different priorities will be sorted first before entering the switch. The high priority packets will be put into the queues first before the low priority packets. If the buffer is full, then the low priority packets are dropped first.

6.4.3 Buffer Management

A variety of buffer allocation schemes with respect to packet priorities have been proposed [6-17]. They include complete partitioning, complete sharing, sharing with maximum queues, sharing with a minimum allocation, sharing with a maximum queue and minimum allocation.

These allocation schemes are introduced below. Assume there are B buffers in the storage space and the number of priorities is M. Let the number of buffers dedicated to packets with priority i is b_i and the maximum number of buffers that packets with priority i can occupy is B_i.

6.4.3.1 Complete Partitioning

The storage space B is partitioned into M disjoint areas; a packet may enter the buffer only if the area associated to its priority is not filled. In this case,

$$b_i > 0,\ b_i = B_i, \text{ and } \sum_{i=1}^{M} b_i = B.$$

6.4.3.2 Complete Sharing

The storage is shared by all the packets; in a sense, there is no priority control. In this case $b_i = 0$ and $B_i = B$.

6.4.3.3 Sharing with Maximum Queues

In this case $b_i = 0$ and $B_i < B$. The summation of all the B_i must be greater than total space B if some sharing is provided.

6.4.3.4 Sharing with a Minimum Allocation

In this case, $b_i > 0$ and $B_i = B$. A minimum number of buffers b_i is always reserved for one priority, and the following common pool of buffers are to be shared among all packets:

$$B - \sum_{i=1}^{M} b_i$$

6.4.3.5 Sharing with a Maximum Queue and Minimum Allocations

The buffer allocation scheme combines both sharing with maximum queue and sharing with a minimum allocation,i.e., $b_i > 0$ and $B_i < B$.

6.4.3.6 Link List Implementation

In this implementation, at the input queue, the packets are structured into a number of link lists, and the number is equal to the number of priorities. Note that the assumption here is that the priority of an ATM connection should not be changed. Otherwise, the packet might transmit out of order. This link list approach is also suitable for output queueing and shared buffer approaches.

6.5 References

[6-1] L. R. Goke and G. J. Lipovski, "Banyan Networks for Partitioning Multiprocessing Systems," First Annual Computer Architecture, pp. 21-28, 1973.

[6-2] K. E. Batcher, "Sorting Networks and Their Applications," AFIPS , vol.32, pp. 307-314, 1968.

[6-3] B. Bingham and H. Bussey, "Reservation-Based Contention Resolution Mechanism for Batcher-Banyan Packet Switches," Electronic Letters, vol. 24, no. 13, pp. 772-773, June 1988.

[6-4] F. A. Tobagi, "Fast Packet Switch Architectures for Broad-band Integrated Services Digital Networks," Proceedings of the IEEE, vol. 78, no. 1, pp. 133-167, Jan, 1990.

[6-5] H. Ahmadi and W. E. Denzel, "A Survey of Modern High-Performance Switching Techniques," IEEE JSAC, vol. 7, no. 5, pp. 1091-1103, Sep. 1989.

[6-6] K. W. Sarkies, "The Bypass Queue in Fast Packet Switching," IEEE Trans. on Communications, vol. 39, no. 5, May 1991.

[6-7] N. Arakawa, Akira Noiri and H. Inoue, "ATM Switch for Multi-Media Switching System," ISS, vol. 5, pp. 9-14, 1990.

[6-8] A. Cisneros, "Large Packet Switch and Contention Resolution Device," ISS, vol. 3, pp. 77-83, 1990.

[6-9] A. Hunag and S. Knauer, "Starlite: A Wideband Digital Switch", IEEE GLOBECOM, pp. 121-125, 1984.

[6-10] J. Y. Hui and E. Arthurs, "Broadband Packet Switch for Integrated Transport", IEEE JSAC, vol. 5, no. 8, pp. 1264-1273., Oct. 1987.

[6-11] J. J. Hickey, W. S. Marcus, " The Implementation of A High Speed ATM Packet Switch Using CMOS VLSI," ISS, 1990, pp.75-84.

[6-12] K. Y. Eng, M. G. Hluchyj and Y. S. Yeh, "Multicast and Broadcast Services in a Knockout Packet Switch," IEEE INFOCOM, pp. 29-34, 1988.

[6-13] J. S. Turner, "Design of a Broadcast Packet Switching Network," IEEE INFOCOM, pp. 667-675, 1986.

[6-14] T. T. Lee, "Nonblocking Copy Network for Multicast Packet Switching," IEEE JSAC, vol. 6, no. 9, pp. 1455-1467, Dec. 1988.

[6-15] G. Nathan, P. Holdaway, and G. Anido, "A Multipath Multicast Switch Architecture," 1988.

[6-16] D.-J. Shyy, "Nonblocking Multicast Fast Packet/Circuit Switching Networks," COMSAT Invention Disclosure No. 31-E-10, June 1991.

[6-17] P. Kermani and L. Kleinrock, "Analysis of Buffer Allocation Schemes in a Multiplexing Node," IEEE ICC, pp. 269-275, 1977.

Section 7

Fast Packet Switch Designs

This section presents detailed hardware implementation of several switch network architectures. For point-to-point switches, the Sorted Banyan Network and the Self-Routing Crossbar-Based Network are described. For multicast switches, the Self-Routing Multicast Banyan Network, the Self-Routing Multicast Crossbar Network, and the Point-To-Point Switch Network with Multicast Output Ports are discussed.

The requirements for the switch networks are listed as follows:

- Maximum Network Size: 30 x 30
- Port Data Rate: Serial Data at 155 Mbps
- ATM Packets (53 bytes/packet)
- Switch Fabric Routing Tag Information set by Signal Processor

Each approach will be evaluated by the following criteria:

- Power Consumption
- Mass/Size
- ASIC Design Complexity
- Fault Tolerance

7.1 Point-To-Point Switch Networks

7.1.1 Sorted Banyan Network

A block diagram of the Sorted Banyan Network is shown in Figure 7-1 below.

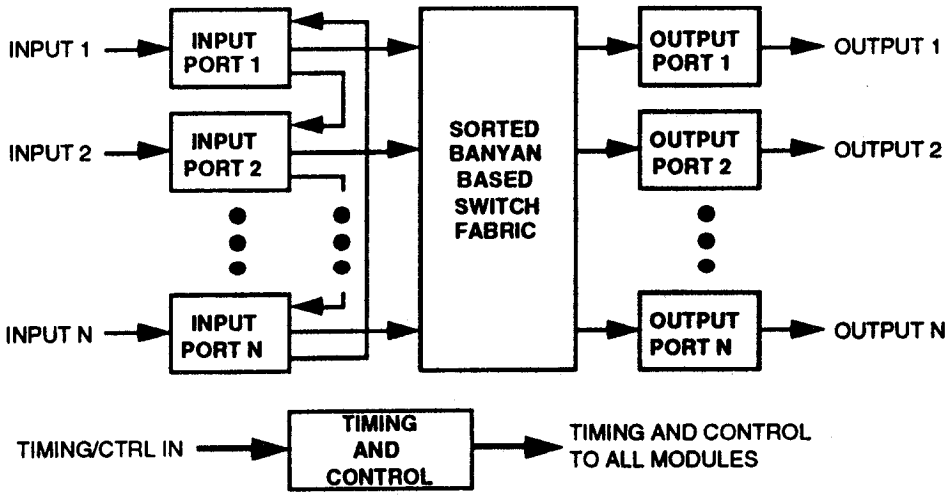

Figure 7-1. Sorted Banyan Network Block Diagram

Serial data from the demodulators is routed through the switch to the modulators in the following steps:

- The serial packet data enters the respective Input Port where it is stored in a 128 packet FIFO type buffer.
- As each packet is stored in the packet buffer, the routing tag associated with the packet is obtain from the routing map memory then stored with the packet in the buffer.
- The Input Port performs Output Port congestion control to ensure that there are no blocked packets at the Output Port. The Input Ports convey their desired destination to each other by means of a serial daisy chained message every packet slot time. When an Input Port reserves an Output Port for the next packet transmission, no other Input Port further down the daisy chain may select the reserved Output Port. The four oldest in the packet buffer will be examined for transmission through the Switch Fabric to minimize head-of-line blocking.
- The selected packet, if any, from each Input Port is sent to the Switch Fabric. The Switch Fabric will route packets from all Input Ports to the desired Output Ports as specified by the routing tag. The Switch Fabric will need to operate at a higher rate than the input serial data to compensate for the added bits for the routing tag and for the lost packet transfers through the Switch Fabric due to head-of-line blocking.
- The Output Port will accept serial data from the Switch Fabric then store it in a 128 packet FIFO type buffer. The serial packet data will be output to the modulator for downlink transmission.

7.1.1.1 Formats

7.1.1.1.1 Routing Tag Format

The format of the routing tag appended to each packet is shown Figure 7-2.

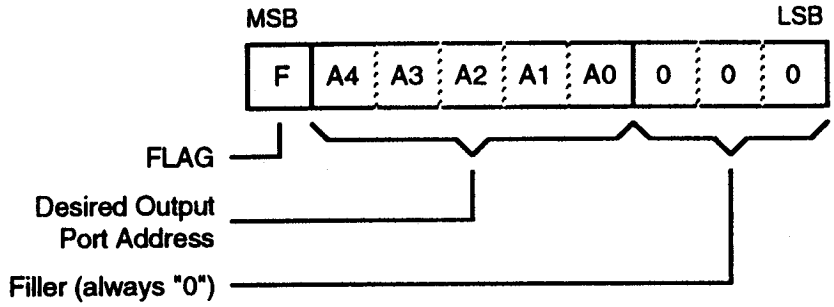

Figure 7-2. Sorted Banyan Routing Tag Format

The FLAG bit, "F", is an indication of a valid packet. If this bit is false, "0", then the packet is invalid and should not be routed through the Switch Fabric. A4 to A0 represent the binary address of the destination Output Port number. The three least significant bits are always "0".

The FLAG bit is the most significant bit and is transmitted first.

7.1.1.1.2 Output Port Congestion Control Message Format

The format of the daisy chain message for the Output Port congestion control is given in Figure 7-3 shown below. This message provides the means by which the Switch Network ensures that no two inputs to the Switch Fabric have the same destination Output Port. The message is sent through each Input Port via a daisy chain interconnect. The Input Port will set a token associated with the desired Output Port provided that the token has not been previously set by another port. The first Input Port to set the token for a given Output Port has reserved that Output Port for the next packet transmission slot period.

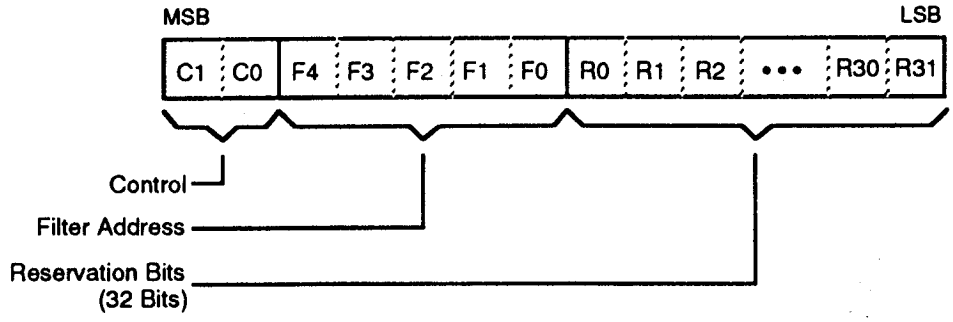

Figure 7-3. Sorted Banyan Congestion Control Message Format

The control message is sent completely around the daisy chain ring five time for every packet slot. The Control bits specify the type of the message as follows:

Table 7-1. Control Bits Definition

CODE (C1C0)	TYPE
00	NOP
01	NORMAL
10	RESET
11	NOP

The first message of a packet slot time will have the RESET code to notify the Input Ports that a new packet sequence is starting. The following four messages will have the NORMAL code to indicate to the Input Ports that normal processing is required. The remaining message(s) in the packet slot will contain the NOP code to prevent the Input Ports from taking any action on the message.

The five bit filter address is used to provide a rotating priority scheme for the Input Ports so that the selection of Output Ports will not favor any particular Input Port. The process for this priority scheme is as follows:

- The filter address is zero ("00000") for all messages except for the first message of a packet slot.
- The filter address in the first message of the previous packet slot is incremented (modulo 32) to provide the filter address in the current packet slot.
- Each Input Port is assigned a unique port number from 0 to 31 in increasing order per location in the daisy chain scheme. The first port in the daisy chain will be assigned port number 0, the second port number 1, third port number 2, etc.
- Each Input Port will examine the filter address in the first message of the packet slot period. Only those Input Ports with port numbers greater than or equal to the filter address will respond to the first message. If the port passes the address filter, the packet at the head of the line will be check for desired destination Output Port. The Input Port with the port number equal to the filter address will be the first port to be able to act on the message; it will therefore have the highest priority for a given packet slot period.
- Each Input Port will be able to act on the second message of the packet slot period since the filter address is "00000". The ports which did not respond in the first message will check the first packet in their buffer; the ports which did respond will check the second packet in their buffer since the first packet was checked in the previous message.

- Each Input Port will be able to act on the third message of the packet slot period since the filter address is "00000". The ports which did not respond in the first message will check the second packet in their buffer; the ports which did respond will check the third packet in their buffer.

- Each Input Port will be able to act on the fourth message of the packet slot period since the filter address is "00000". The ports which did not respond in the first message will check the third packet in their buffer; the ports which did respond will check the fourth packet in their buffer.

- Each Input Port will be able to act on the fifth message of the packet slot period since the filter address is "00000". The ports which did not respond in the first message will check the fourth packet in their buffer; the ports which did respond do not act on this message since they have already checked all four messages at the head of the buffer.

- A NOP message will follow the fifth message and will be valid till the first message of the next packet slot period. No Input Ports will act on a NOP message.

The thirty-two reservations bits represent each of the thirty-two Output Ports. Reservation bit *Ri* corresponds to Output Port *i*. All bits are reset to "0" in the first message of the packet slot period. As this message is passed through the Input Ports, reservation bit *i* will be set if the following conditions are met:

- The Control mode is not NOP and the Input Port address is greater than or equal to the filter address in the message.

- Reservation bit *i* is "0" when it entered the Input Port.

- The packet being checked desires to send its data to the selected destination Output Port *i*.

7.1.1.2 Input Port

A block diagram of the Input Port is shown in Figure 7-4. The Input Port performs the following major tasks:

- Serial data enters the Input Port via the DATA IN signal. This data is aligned to the common control signals CLK IN and NEW PACKET. Transitions of the serial data occur on the rising edge of CLK IN. NEW PACKET determines the start of the received packet.

  The differentially received serial data and control lines are converted to single-ended ECL by the I/O modules.

- The Input Port ASIC will convert the serial data to parallel using the control signals.

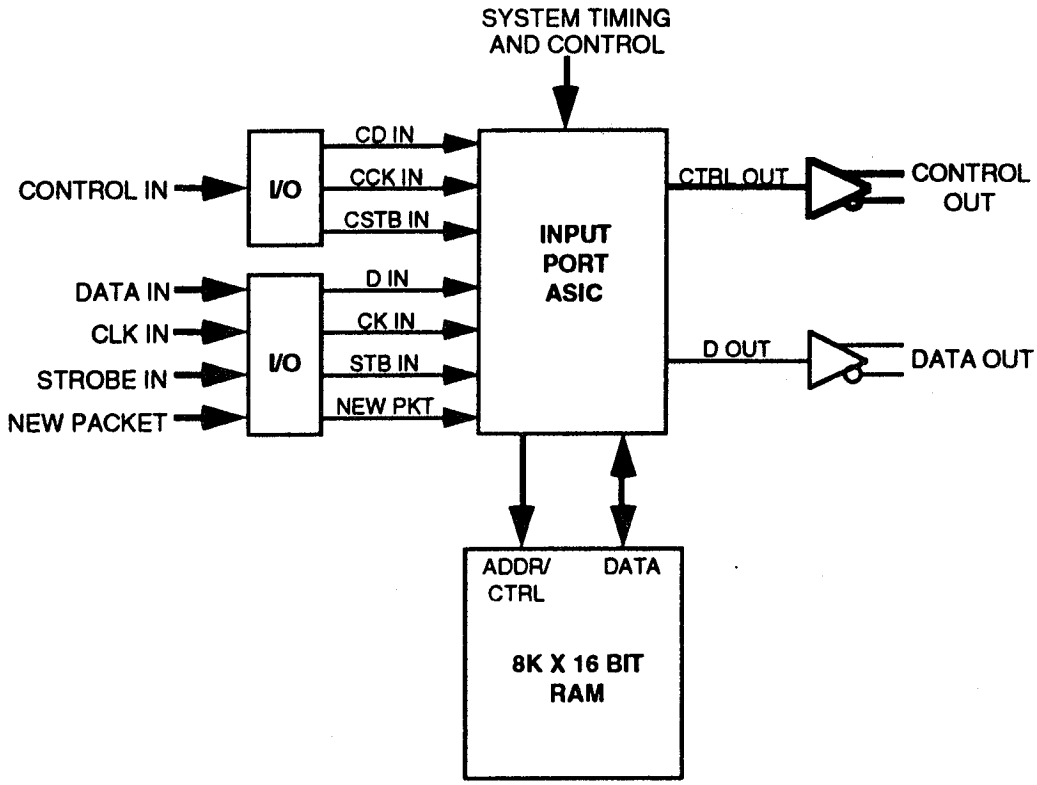

Figure 7-4. Sorted Banyan Input Port Block Diagram

- The parallel input data is then stored in the first half of the RAM which is configured as a 4K by 16 bit FIFO type buffer. The buffer is configured as 128 pages of 32 words (64 bytes). The 53 bytes of packet data is stored sequentially starting at byte 1 (byte 0 is left open for the routing tag) of the next available page - the ten bytes at the end of the page are not used.

- As an input packet is being stored in the buffer, its routing tag is being obtained. The Input Port ASIC will examine up to sixteen bits of the destination address in the packet header. These sixteen bits will be used to obtain the routing tag by the mapping shown in Figure 7-5 below.

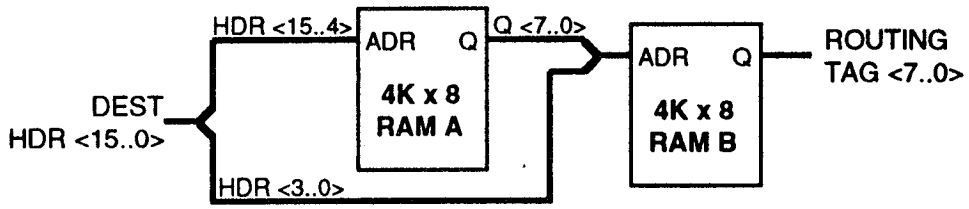

Figure 7-5. Routing Tag Memory Mapping

RAM A and RAM B are actually implemented as two 2K x 16 RAMs instead of two 4K by 8 RAMs. The least significant bit of the address is used to

determine if the high or low byte of the memory access is used. These two 2K x 16 RAMs occupy the lower 4K of the 8K x 16 RAM.

Both RAMs have been intialized by the signal processor. The most significant twelve bits of the destination address will be condensed by the mapping in RAM A into eight bits. These eight bits and the least significant four bits of the destination address will be used to fetch the routing tag from RAM B. This method provides a mapping of 64K addresses (2^{16}) to a minimum of 256 routing tags with only 16K bytes of memory.

- The eight bit routing tag is stored in the byte 0 location of the buffer page associated with the packet.

- Output Port congestion control is performed. The congestion control information is contained on the CONTROL IN signals. These differential signals are converted to single-ended ECL by the I/O module. The message is received serially on the CD IN signal with the associated clock, CCK IN. The CSTB IN provides a strobe timing signal to delineate the fields in the message stream.

  The congestion control circuitry is initialized when a RESET mode is detected in the message. If the Input Port port number is greater than or equal to the filter address and the mode is not NOP, the congestion control circuity is enabled. The congestion control circuitry will check the token in the message associated with the desired Output Port. If the token is not set, the control circuitry will set it to reserve the selected Output Port then prepare the associated packet for transmission to the Switch Fabric during the next packet slot period. If the token is set, the desired Output Port has been reserved by another Input Port so no action is taken except to prepare the subsequent packet in line for checking with the next control message.

  The congestion control message will be sent five times during each packet slot. Every Input Port will have an opportunity to check all four packets at the head-of-line in the packet buffer. Only one packet is checked in an Input Port for each message. In this manner, the oldest packets in the buffer of all Input Ports are checked before subsequent packets in any buffer are checked.

- Once the packet to be transmitted to the Switch Fabric is determined, the packet, including the routing tag, must be fetched from the packet buffer. This fetch starts as the current packet transmission is ending. The STROBE IN signal identifies the start of the packet slots within the Switch Fabric. The packets will be transmitted at a higher rate than the port data rate to compensate for the added bits in the routing tag (about 2%) and for the blocked packets due to head-of-line blocking (about 20%). Therefore, the minimum Switch Fabric rate must be 190 MHz (155 MHz * 102% * 120%).

- The fetched 16 bit words from the packet buffer are converted to serial data to be sent to the Switch Fabric.

The Input Port ASIC is responsible for performing all the functions above. The four main functions of the ASIC are:

- Port Side Processing. This includes the input serial data conversion to parallel, generating the RAM packet buffer address to store the parallel word, and generating the RAM address to obtain the routing tag.
- Output Port Congestion Control. This includes all functions associated with the Output Port Congestion Control.
- Switch Fabric Side Processing. This includes storing of the start locations of the four packets at the head-of-line, generating the routing tag address for use in the Output Port Congestion Control, generating the address for the packet data to be transmitted to the Switch Fabric, and converting the parallel data to serial format.
- RAM Processing. This includes all address/data buffering and multiplexing and the generation of the control signals to the RAM memory.

7.1.1.2.1 Port Side Processing

A block diagram of the Port Side processing functions is shown in Figure 7-6

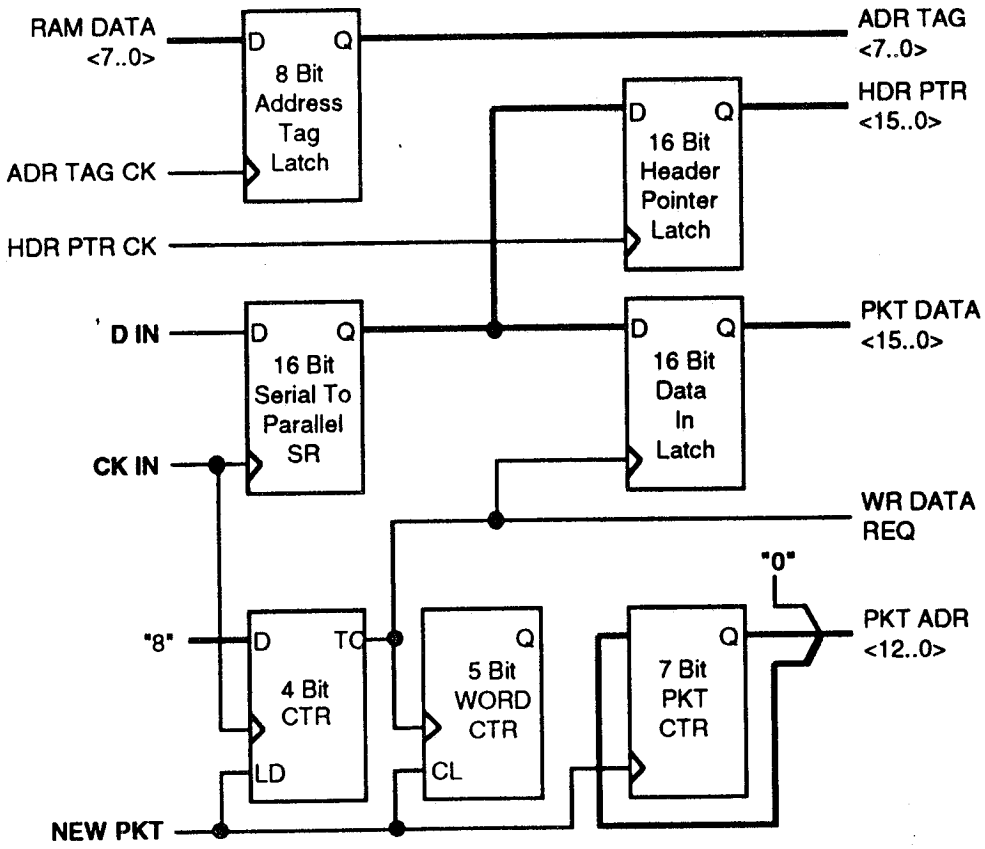

Figure 7-6. Port Side Processing Block Diagram

Serial to Parallel Conversion. The "16 Bit Serial To Parallel SR" (S/P SR), the "4 Bit CTR" (BIT CTR), and the "16 Bit Data In Latch" (DATA LATCH) perform the serial to parallel conversion function.

- BIT CTR represents the bit count within a sixteen bit word boundary of the incoming serial data stream. BIT CTR is initialized to "8" on NEW PKT and is clocked on the data clock signal, CK IN. The terminal count output of the BIT CTR, TC, will occur eight bits after the start of the packet and every sixteen bits thereafter.

- Serial data, D IN, is clocked into the S/P SR by the data clock, CK IN to convert it to parallel data for storage in the RAM memory.

- The output of the S/P SR is latched into the DATA LATCH by the terminal count output of the BIT CTR at which time a request is made to the RAM Processing module to store the output of the DATA LATCH in the packet buffer. This request must be serviced before the next request is made sixteen clock periods later.

  The first buffer word will contain only eight bits (byte 0 of the packet) in the low byte of location 0 of the packet page and the subsequent 52 bytes of the packet will be stored in the next 26 locations of the buffer.

Generate Buffer Address for Packet Data. The buffer address for the packet data is generated by the "4 Bit CTR" (BIT CTR), the the "5 Bit WORD CTR" (WORD CTR), and the "7 Bit PKT CTR" (PKT CTR). The BIT CTR and WORD CTR are initialized at the start of every packet by NEW PKT. BIT CTR represents the bit count of the incoming serial data. The terminal count output, TC, of the BIT CTR goes active when a word of the packet has been latched in DATA LATCH. The WORD CTR therefore represents the word count of the incoming packet. The PKT CTR is incremented for every new packet by NEW PKT. The PKT CTR represents 1 of 128 packet pages within the packet buffer. The address for the parallel data is obtained by appending the 7 bit output of the PKT CTR to the 5 bit output of the WORD CTR, resulting in 12 bits (4K data buffer address space). The most significant bit of the RAM address must be forced to "0" to access the data buffer portion of the RAM.

Generate Addresses for Obtaining Routing Tag. The RAM address for obtaining the routing tag is accomplished by the "8 Bit Address Tag Latch" (TAG LATCH) and the "16 Bit Header Pointer Latch" (HDR LATCH). The circuitry for obtaining the routing tag is able to map 16 bits (64K unique address) to a minimum of 256 routing tags. These sixteen bits, which should be two consecutive bytes within the packet, are latched into HDR LATCH.

The most significant twelve bits of HDR LATCH are used for the first access. The most significant eleven of the twelve bits are used for the least significant bits of the address; the two most significant bits of the address are set to "10" to select the third 2K block within the RAM memory. The least significant bit of the HDR LATCH determines if the

low or high byte is to be used. The selected byte of the first fetch is latched into the TAG LATCH.

The second address is formed from the contents of the TAG LATCH and the least significant four bits of the HDR LATCH. The two most significant bits of the RAM address are "11" to select the last 2K block of the RAM memory. The next eight bits will use the contents of the TAG LATCH. The least significant three bits of the address will be formed by the most significant three of the four bits from the HDR LATCH. The least significant bit of the HDR LATCH will determine if the low or high byte of the fetched word will be used.

The above process requires that the contents of the two specified maps be determined by the signal processor prior to the reception of the packet. The signal processor will be able to change one entry in the RAM at a time, i.e., the map is not a ping-pong type memory where the entire map can be changed offline then flipped online.

7.1.1.2.2 Output Port Congestion Control

A block diagram of the Output Port Congestion Control is shown in Figure 7-7. Figure 7-8 shows the the Mode Circuitry and Figure 7-9 shows the Address Filter Circuitry. A timing diagram for the Mode Circuitry and the Address Filter Circuitry is shown in Figure 7-10. A timing diagram for the setting of a token bit is shown in Figure 7-11.

The control message will experience a one clock delay through the circuitry. The serial data, CD IN, and strobe signal, CSTB IN, are clocked into flip-flop FF1 and FF2 on the negative edge of the input clock, CCK IN, to obtain the DIN 1 and STB 1 signals. Flip-flop FF3 reclocks STB 1 on the positive edge of CCK IN to compensate for the delay to the data through flip-flop FF4. Under most conditions, the or-gate, G1, will pass DIN 1 to flip-flop FF4 for transmission to the next Input Port. The only exception is when this Input Port wishes to set one of the tokens to reserve the associated Output Port.

The following conditions must be met for this Input Port to be allowed to set one of the Output Port token bits.

- This Input Port does not already have an Output Port reserved for the next packet slot period. This information is present in flip-flop FF5. When set, it is an indication that a token has already been set in a previous message of the packet slot period. See final criteria below for a more detailed description of the operation of this flip-flop.

- The Mode Control bits at the beginning of the message must be a non-NOP code. The Mode Control block diagram is shown in Figure 7-8. The first two bits of the message are latched into flip-flops FF6 and FF7. These two bits are decoded to generate the signals for NOP, NORMAL, and RESET. The RESET signal is further waveshaped to provide the RST signal as shown in the timing diagram. See the timing diagram in Figure 7-10.

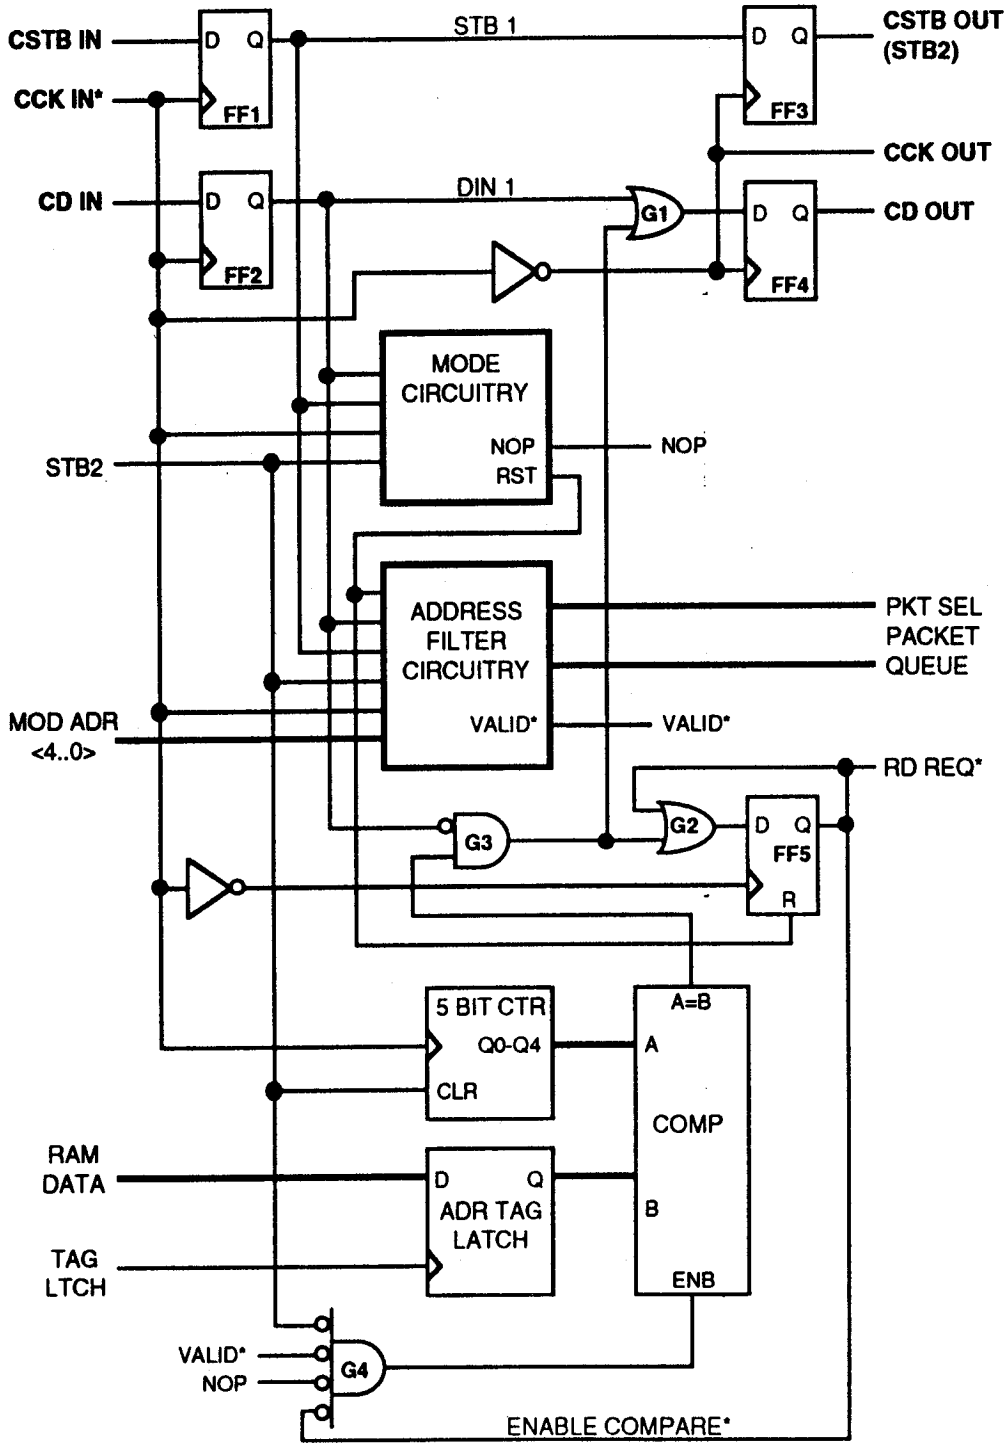

Figure 7-7. Output Port Congestion Control Block Diagram

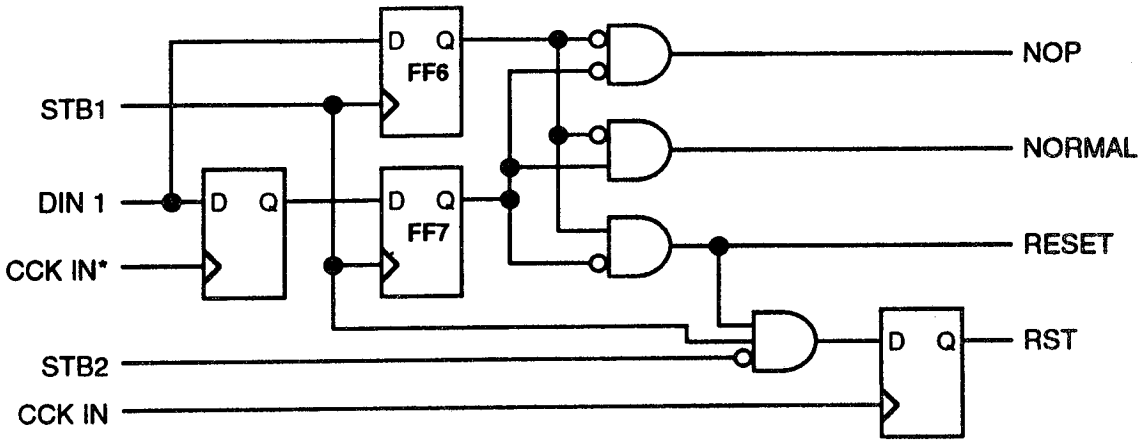

Figure 7-8. Mode Circuitry Block Diagram

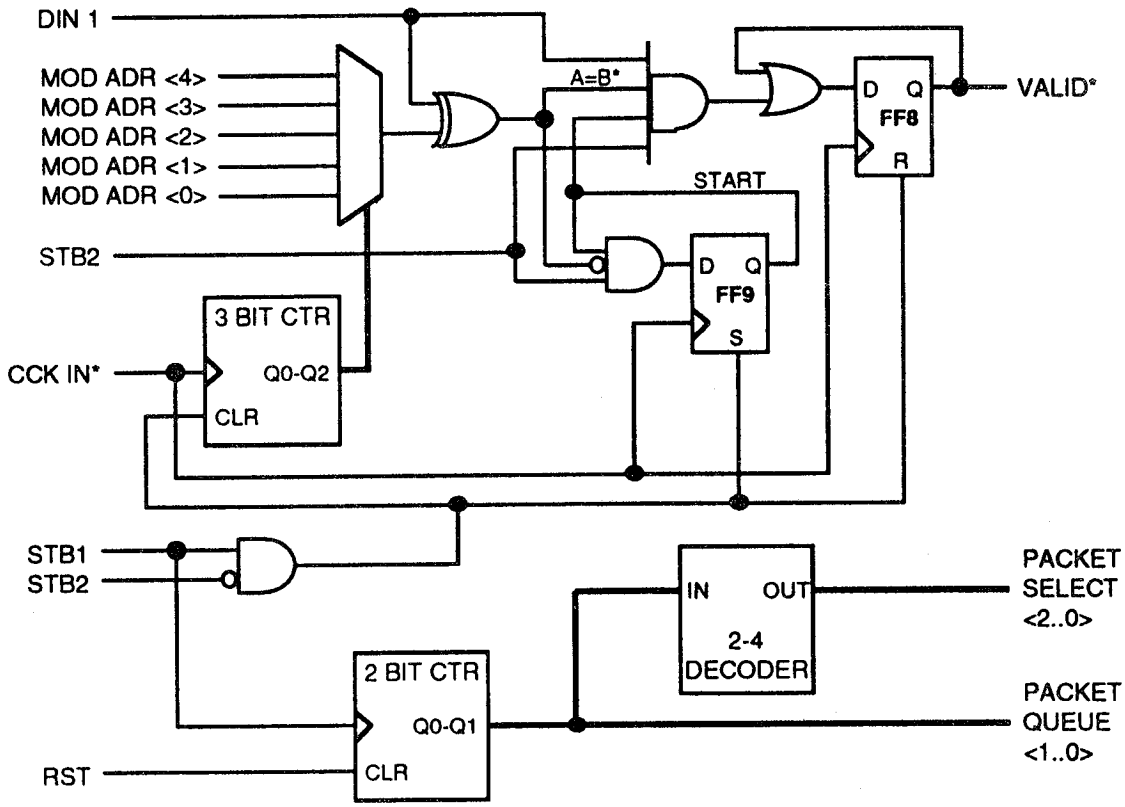

Figure 7-9. Address Filter Block Diagram

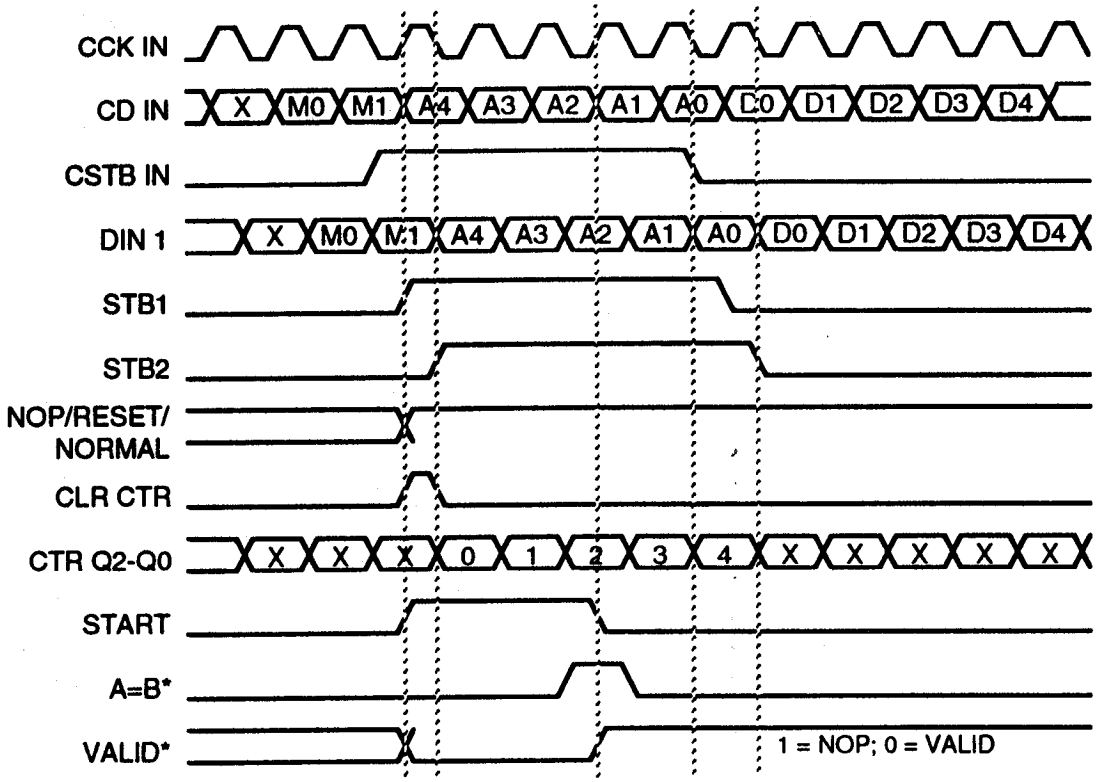

Figure 7-10. Mode Control And Address Filter Timing Diagram

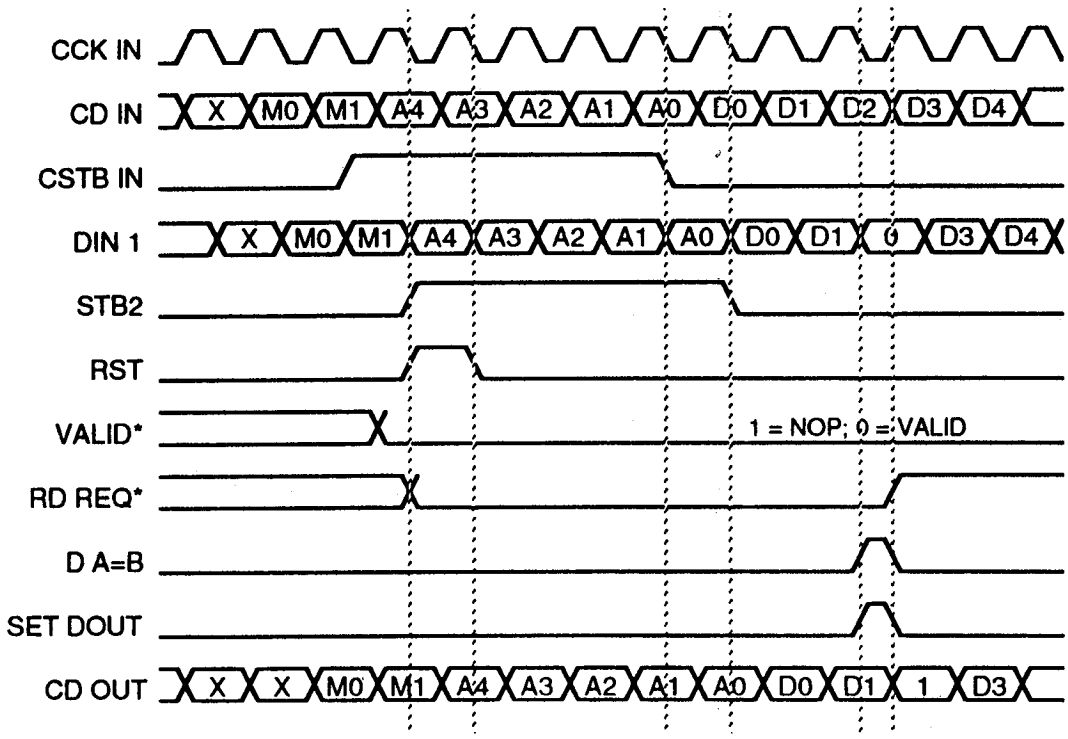

Figure 7-11. Timing Diagram

- The port number for this Input Port must be greater than or equal to the Filter Address in the control message. The Address Filter Circuitry block diagram is shown in Figure 7-9. The output of flip-flop FF8 indicates if the Input Port port number is greater than or equal to the Filter Address - a low indicates that the Input Port passed the address filter. The output of flip-flop FF9 is used to enable the 4-input and gate which feeds flip-flop FF8. The circuitry performs as follows:

    - When STB 1 goes active, the counter is cleared to zero, flip-flop FF8 is reset and flip-flop FF9 is set.

    - When STB 2 goes high, the check for the first bit is performed. The multiplexor will present port number bit A4 to the XOR gate. If this bit compares to the DIN 1 bit (A4 of the Filter Address), then no action is taken. If this bit differs from DIN 1, then the following tasks are implemented:

        1. FF9 is reset to disable all further checks

        2. If DIN 1 is high, it indicates that the Filter Address is larger than the port number so flip-flop FF8 is set and remains set until the next RST pulse.

        3. If DIN 1 is low, it indicates that the Filter Address is smaller than the port number so flip-flop FF8 remains low until the next STB 1 pulse.

    - If no action was taken for the first bit check, then the second bit, A3, is checked on the next clock cycle. Again, if the two A3 bits compare, no action is taken; if they differ, the three actions in the above step are implemented.

    - The above step is repeated until the the XOR detects a difference or the STB 2 signal goes low, indicating the end of the Filter Address field. If the STB 2 signal goes low with no differences detected, the port address is equal to the Filter Address and flip-flop FF8 should remain reset. Figure 7-10 shows an example where the port number was detected as smaller than the filter address on bit A2.

- The final criteria for the Input Port to be allowed to set a reservation token bit in the message is that the token has not been previously set by another Input Port. When set, flip-flop FF5 indicates that this Input Port has a packet to send to the Switch Fabric in the next packet slot period. This bit is reset on RST and if all conditions are satisfied, it will be set by the end of the current packet slot period. The "5 Bit CTR" (PORT CTR) is reset on STB 2 and clocked by the input clock, CCK IN. Its output represents the Output Port associated with the current token received on DIN 1. When PORT CTR is "00000", the token bit present on DIN 1 is for Output Port 0, "00001" indicates port 1, "00010" indicates port 2, etc. The output of PORT CTR is compared

with the data in "ADR TAG LATCH" which contains the address of the desired destination Output Port. When these two values match, DIN 1 is checked. If the token is set, no action is taken. If the token is not set, the output of or-gate G1 will be high, setting the token in the control message for the subsequent Input Ports. Also, flip-flop FF5 will be set until the next RST pulse. The setting of flip-flop FF5 will prevent the setting of tokens for subsequent messages within the same packet slot period.

- The Address Filter Circuitry is responsible for determining the packet to be checked by the Congestion Control Circuitry. The "2 Bit CTR" (PKT CK CTR) specifies, via PACKET QUEUE, which packet latch in the Switch Fabric Side Processing will be used to fetch the next routing tag. These two bits are also decoded to provide the PKT SEL signals to update the packet latches every packet slot period. The PKT CK CTR is cleared on RST and clocked every STB 1.

7.1.1.2.3 Switch Fabric Side Processing

A block diagram of the Switch Fabric side processing is shown in Figure 7-12.

Storing Start Packet Locations for the Four Head-Of-Line Packets. The "7 Bit PKT CTR" (PKT CTR) contains the start location for the next packet to be serviced from the data buffer. The four packet latches, "7 Bit Pkt *i* Ltch" (PKT*i* LTCH), contain the current start locations for the four packets at the head-of-line. PKT0 LTCH contains the start location for the packet at the head of the line followed by PKT1 LTCH, PKT2 LTCH, and PKT3 LTCH. A packet which is transmitted to the Switch Fabric can be associated with any one of the four packet latches. Once the packet is transmitted the contents of all packet latches down the line should move up by one and the new packet header should be latched into PKT3 LTCH. For example, if the packet that was transmitted was represented by PKT1 LTCH, then the contents of PKT2 LTCH should be moved to PKT1 LTCH, the contents of PKT3 LTCH should be moved to PKT2 LTCH, and the PKT CTR contents should be latched into PKT3 LTCH. PKT0 LTCH remains unaffected. The four packet latches now contain the four packets at the head of the line in the same time sequence in which they entered the data buffer.

The PKT CLK line increments the PKT CTR and latches the contents of the PKT CTR into PKT3 LTCH. If PKT2 LTCH is to be updated, PKT SEL 2 will be high to enable the transfer of PKT3 LTCH contents to PKT2 LTCH. Similarly, if PKT1 LTCH is to be updated, PKT SEL 1 will be high to enable the PKT2 LTCH and PKT1 LTCH for update. If PKT0 LTCH is to be updated, PKT SEL 0 will be high enabling all packet latches to be updated.

Generation of Routing Tag RAM Address. The routing tag is stored in the first location, high byte, of the 32 word page in the data buffer. This page represents one packet and is specified by the seven bits held in the packet latches, PKT*i* LTCH. The RAM address to fetch the routing tag for the packet associated with PKT*i* LTCH is formed as follows:

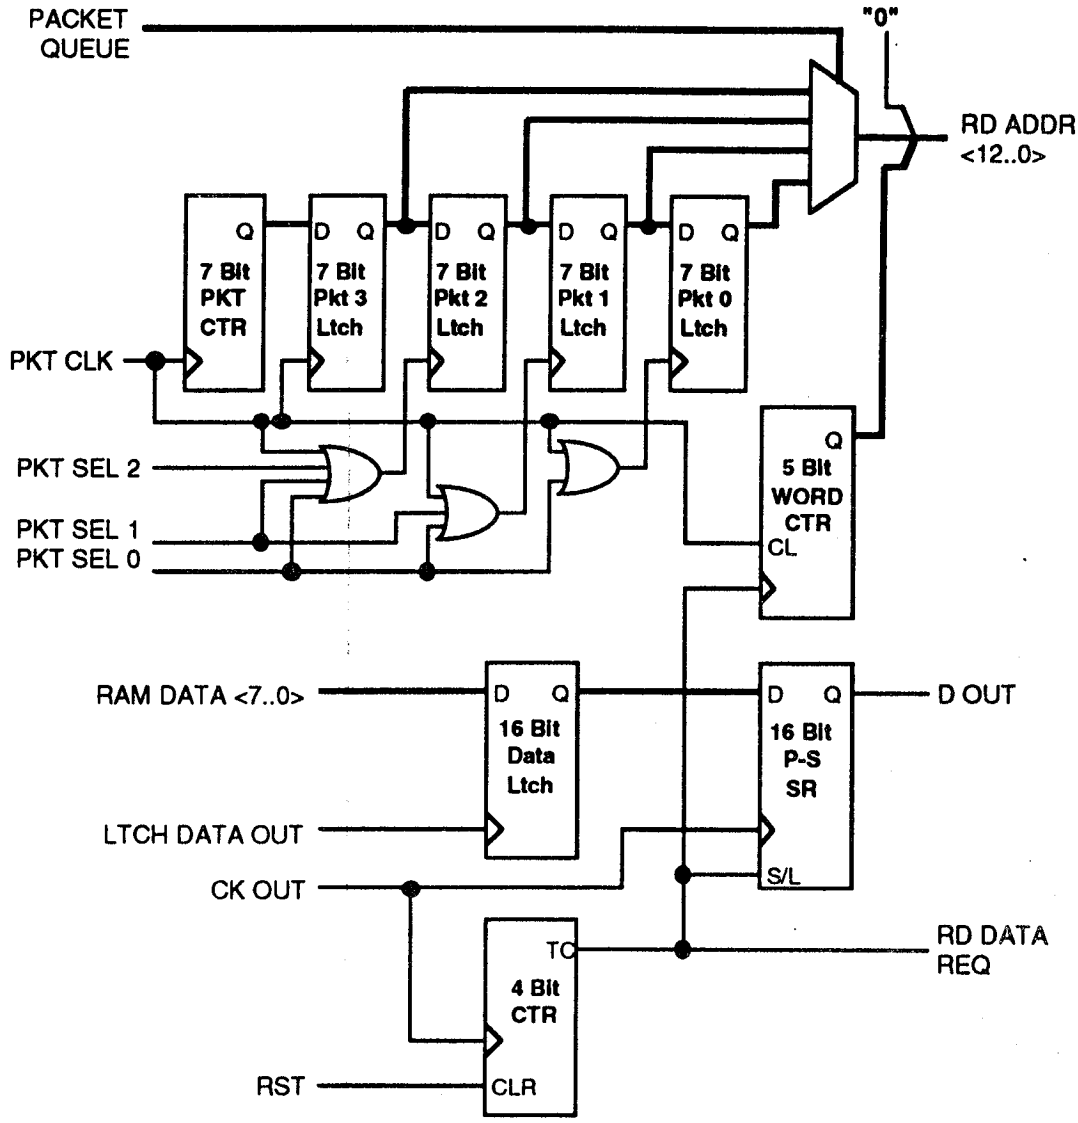

Figure 7-12. Switch Fabric Side Processing Block Diagram

- The most significant bit, A12 is forced to "0" to select the data buffer portion of the RAM memory.
- The next seven bits, A11 to A5, are the contents of the PKT*i* LTCH.
- The least significant five bits, A4 to A0, represent the word count and are forced to "00000" to select the first word of the packet. The routing tag is in the high byte of the selected word.

Generation of Packet Data RAM Address. The RAM address to fetch the packet data associated with the selected packet latch, PKT*i* LTCH, is formed as follows:

- The most significant bit, A12 is forced to "0" to select the data buffer portion of the RAM memory.

- The next seven bits, A11 to A5, are the contents of the PKT*i* LTCH.
- The least significant five bits, A4 to A0, represent the word count and are obtained from the word counter, "5 Bit WORD CTR" (WORD CTR). The WORD CTR is cleared at the beginning of every packet transmission. It is clocked from the terminal count output, TC, of the bit counter, "4 Bit CTR" (BIT CTR).

Parallel to Serial Conversion. Data fetched for output to the Switch Fabric is latched into "16 Bit Data Ltch" (DATA LATCH). The BIT CTR is reset at the beginning of the packet transmission by RST. When the terminal count of the BIT CTR is active every sixteen bit periods, the "16 Bit P-S SR" (P/S SR) will be parallel loaded. For the next fifteen clocks of the data clock, CK OUT, the data will be serially shifted out on D OUT of the P/S SR.

7.1.1.2.4 RAM Processing

A block diagram of the RAM Processing circuitry is shown in Figure 7-13

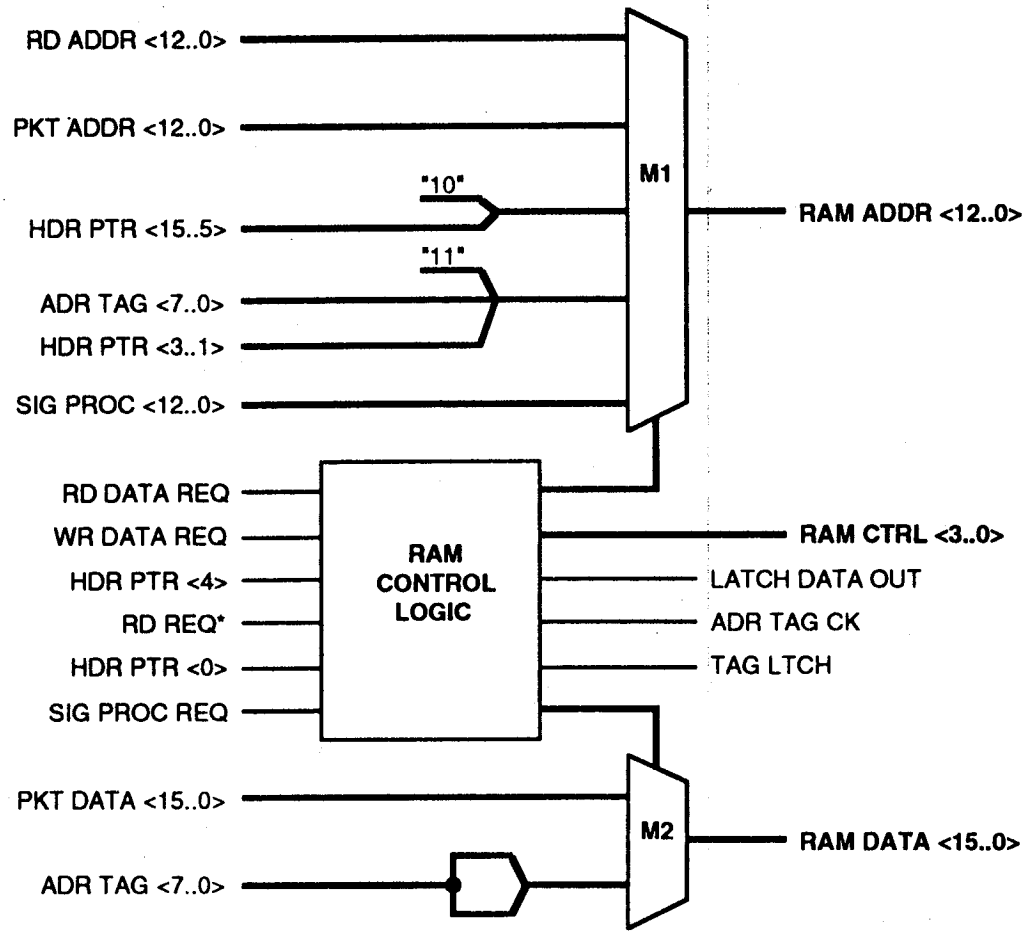

Figure 7-13. RAM Processing Block Diagram

Address Multiplexing. The multiplexor M1 selects the desired RAM address for the current access as determined by the RAM Control Logic. Sources of RAM addresses are:

- RD ADDR <12..0>. Address from the Switch Fabric Side Processing module. This address is used to fetch the routing tag for the Output Port Congestion Control and the packet data for transmission to the Switch Fabric.
- PKT ADDR <12..0>. Address from the Port Side Processing module. This address is used to store the incoming packet data and associated routing tag.
- "10" | HDR PTR <15..5>. Address from the Port Side Processing module. This address is the indirect address for fetching the routing tag. The most significant two bits are "10" to select the routing tag indirect address map portion of the RAM memory.
- "11" | ADR TAG <7..0> | HDR PTR <3..1>. Address from the Port Side Processing module. This address will fetch the routing tag from the RAM memory. The most significant two bits are "11" to select the routing tag portion of the RAM memory.
- SIG PROC <12..0>. Address from the Signal Processor. This address is used by the Signal Processor for read and write accesses to the RAM memory.

Data Multiplexing. The multiplexor M2 selects the desired RAM write data for the current access as determined by the RAM Control Logic. Sources of RAM address are:

- PKT DATA <15..0>. Write data from the Port Side Processing module. This is the parallel input packet data to be stored in the data buffer.
- ADR TAG <7..0>. Address tag from the Port Side Processing module. This is the eight bit address tag used for routing the packet through the Switch Fabric.

RAM Control. The RAM Control Logic module will accept requests from all sources for accesses to the RAM memory. It will arbitrate the requests, then generate the proper memory and register control signals to execute the desired access. Priority of accesses will be arbitrated as follows:

- RD DATA REQ. Requests for a read access to fetch the packet data to be transmitted to the Switch Fabric. This has the highest priority since it has the least amount of time between requests. The rate of requests is approximately 12.5 MHz (200 MHz / 16 Bits).
- WR DATA REQ. Requests for a write access to store the packet data to the data buffer. This has the second highest priority with a rate of request at approximately 9.7 MHz (155 MHz / 16 Bits).
- RD REQ*. Request from the Output Port Congestion Control for a read access to fetch the next address tag. This has the third highest priority with a rate of request at approximately 2.6 MHz (100 MHz / 39 Bits).

- Indirect Address Tag Fetch. Request from internal RAM Control Logic for a read access to the routing tag indirect address map. This has the fourth highest priority with a rate of request at approximately 183 KHz (9.7 MHz / 53 Wds/Pkt).
- Routing Tag Fetch. Request from internal RAM Control Logic for a read access to the routing tag memory. This has the fifth highest priority with a rate of request the same as the indirect address tag fetch of approximately 183 KHz.
- Routing Tag Write. Request from internal RAM Control Logic for a write access to the data buffer. This has the sixth highest priority with a rate of request the same as the indirect address tag fetch of approximately 183 KHz.
- SIG PROC REQ. Request from the Signal Processor for a read or write access to the RAM memory. This has the lowest priority since there is no real time requirement for these accesses.

7.1.1.2.5 Input Port Characteristics

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-2.

Table 7-2. Input Port ASIC Estimated Gate Count

FUNCTION	GATES/FNCTN	QUANTITY	TOTAL GATES
D FLIP-FLOP	7	200	1400
LOGIC GATE	1.75	575	1006
TOTAL + 20%			2900

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-3.

Table 7-3. Input Port ASIC Estimated Power Consumption

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	7	4	0.580 W
OUTPUT I/O	0	18	0	4	0.530 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	125	30	25	20	0.095 W
LOGIC GATE	300	240	35	0	0.050 W
TOTAL + 20%					1.8 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Input Port Power Characteristics. The Input Port will require one ASIC chip, two 8K by 8 bit memories, differential receivers and drivers, and minor support logic requiring an estimated ten square inches of board space. The total power consumption of the Input Port is summarized in Table 7-4.

Table 7-4. Input Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.8 W
RAM	1.5 W
I/O & SUPPORT	1.5 W
TOTAL	4.8 W

7.1.1.3 Output Port

A block diagram of the Output Port is shown in Figure 7-14.

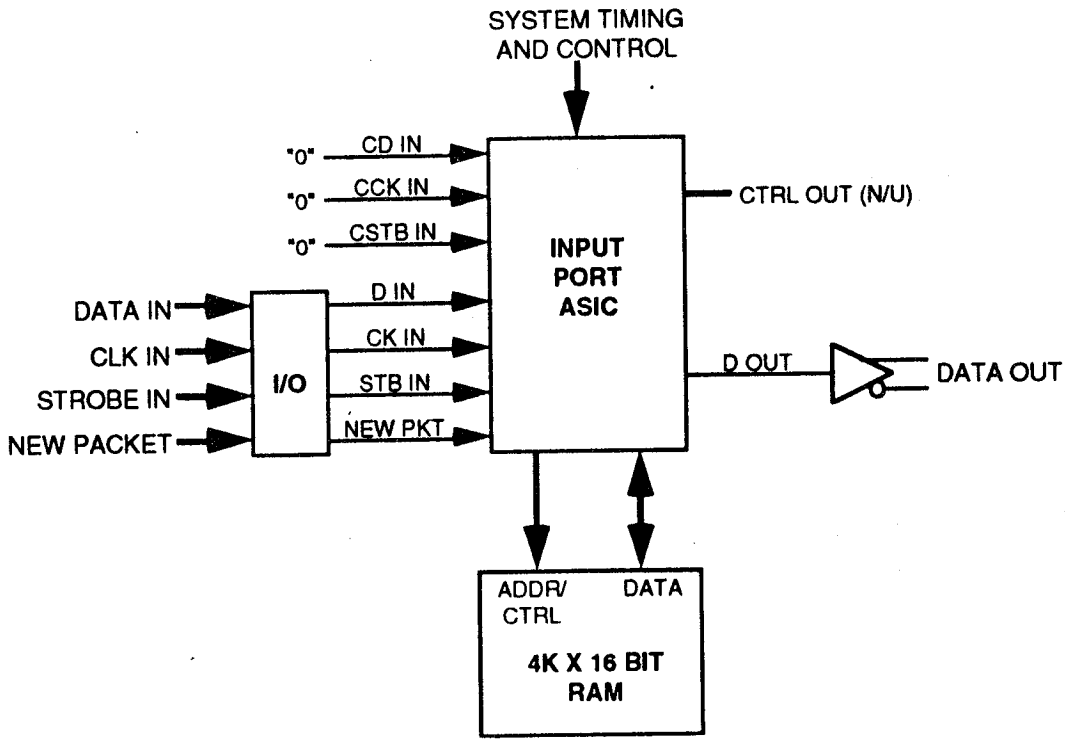

Figure 7-14. Output Port Block Diagram

The Output Port uses the same ASIC as the Input Port. Only a subset of the Input Port ASIC is required reducing the chip power consumption. The functions required for the Output Port are as follows:

- Serial data enters the Output Port from the Switch Fabric via the DATA IN signal. This data is aligned to the common control signals CLK IN and NEW PACKET and is received at the Switch Fabric rate (> 190 Mbps). Transitions of the serial data occur on the rising edge of CLK IN. NEW PACKET determines the start of the received packet from the Switch Fabric.

  The differentially received serial data and control lines are converted to single-ended ECL by the I/O modules.

- The Input Port ASIC will convert the serial data to parallel using the control signals.

- The parallel input data is then stored in the RAM which is configured as a 4K by 16 bit FIFO type buffer. The buffer is configured as 128 pages of 32 words (64 bytes). The 54 bytes of packet data are stored sequentially starting at byte 0 of the next available page - the ten bytes at the end of the page are not used.

- When available in the data buffer, packets are continually sent to the modulator. The fetch for the next packet starts as the current packet transmission is ending.

- The fetched 16 bit words from the packet buffer are converted to serial data to be sent to the Switch Fabric. The output serial data is transmitted at the port rate of 155 Mbps.

The Input Port ASIC is responsible for performing all the functions above. The subset of functions performed for the Output Port are:

- Switch Side Processing. This includes the input serial data conversion to parallel and generating the RAM packet buffer address to store the parallel word.

- Port Fabric Side Processing. This includes generating the address for the packet data to be transmitted to the modulator and converting the parallel data to serial format.

- RAM Processing. This includes all address/data buffering and multiplexing and the generation of the control signals to the RAM memory.

7.1.1.3.1 ASIC Description

See the discussion for the Input Port ASIC functions for a description of these modules.

7.1.1.3.2 Output Port Characteristics

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-2. This is the same ASIC as used on the Input Port and will require approximately 2900 gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-5. The values have been altered from the Input Port design to reflect the parts of the ASIC not used for the Output Port.

Table 7-5. Input Port ASIC Estimated Power Consumption For Output Port

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	4	4	0.440 W
OUTPUT I/O	0	18	0	1	0.345 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	105	30	25	20	0.095 W
LOGIC GATE	300	240	35	0	0.050 W
TOTAL + 20%					1.4 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Output Port Power Characteristics. The Output Port will require one ASIC chip, one 4K by 16 bit memory, differential receivers and drivers, and minor support logic requiring an estimated seven square inches of board space. The total power consumption of the Output Port is summarized in Table 7-6.

Table 7-6. Output Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.4 W
RAM	1.0 W
I/O & SUPPORT	1.0 W
TOTAL	3.4 W

7.1.1.4 Switch Fabric

7.1.1.4.1 Switch Fabric Configurations

The Sorted Banyan Switch Fabric requires both a Banyan switching matrix and a Batcher sorting matrix. The Banyan switching matrix will be presented, followed by the Batcher sorting matrix, and concluding with the combined network to provide the Sorted Banyan Switch Fabric.

7.1.1.4.1.1 Banyan Switch Matrix

The Sorted Banyan Switch Fabric will require 8 x 8, 16 x 16, and 32 x 32 Banyan switch matrices. The standard configuration for these matrices are shown in Figures 7-15 to 7-17.

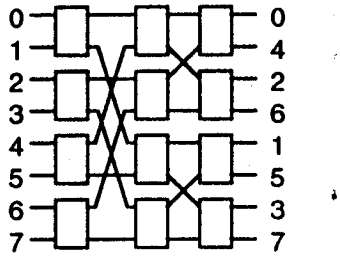

Figure 7-15. Standard 8 x 8 Banyan Matrix Configurations

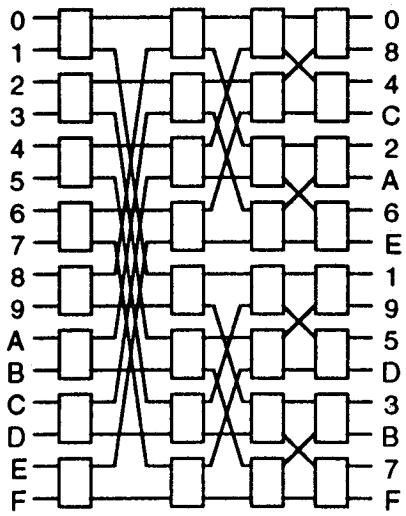

Figure 7-16. Standard 16 x 16 Banyan Matrix Configurations

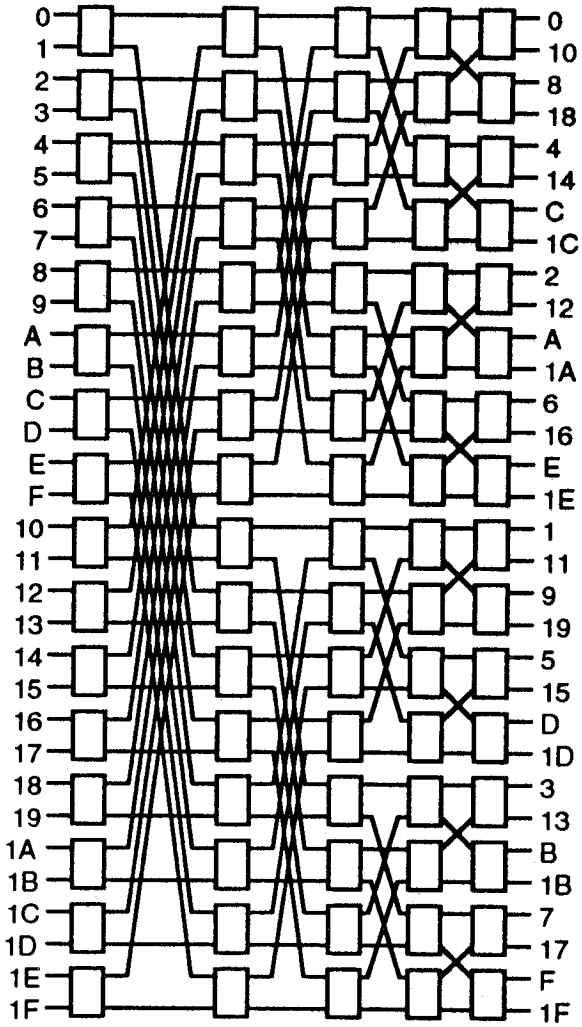

Figure 7-17. Standard 32 x 32 Banyan Matrix Configurations

The ASIC design for the Banyan matrices will incorporate multiple 2 x 2 cells to reduce the physical size and power consumption overhead of the Switch Fabric. The ASIC will be configurable in the two modes shown in Figure 7-18.

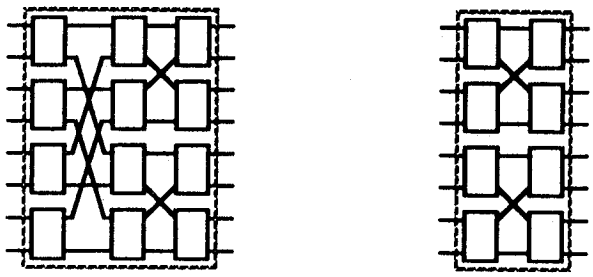

Figure 7-18. Banyan ASIC Configurations

The three Banyan matrix configurations can be implemented using the two modes of the Banyan ASIC shown above. The 8 x 8 Banyan matrix is simply the MODE A ASIC. The topology of the Standard 16 x 16 Banyan matrix can be rearranged as shown in Figure 7-19 to permit the use of four MODE B ASICs. The arrows show the translation of the highlighted cells from their original position in the same column of the standard configuration. The 32 x 32 Banyan matrix can be rearranged as shown in Figure 7-20 to permit the use of four MODE A and four MODE B ASICs.

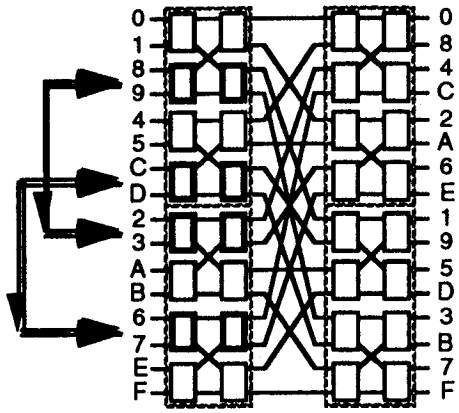

Figure 7-19. 16 x 16 Banyan Matrix Implementation

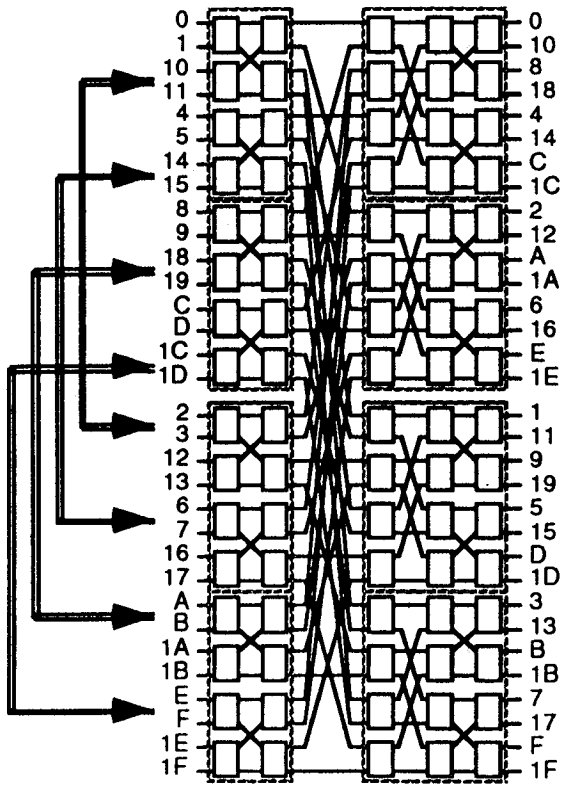

Figure 7-20. 32 x 32 Banyan Matrix Implementation

7.1.1.4.1.2 Batcher Sorting Matrix

The Sorted Banyan Switch Fabric will require 8 x 8, 16 x 16, and 32 x 32 Batcher sorting matricies. The standard configuration for these matrices are shown in Figures 7-21 to 7-23.

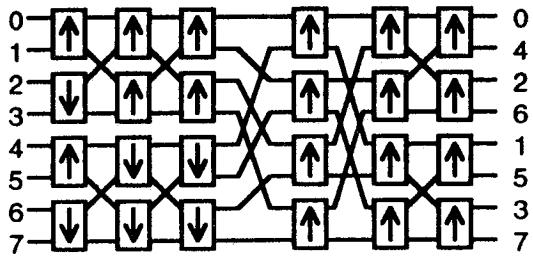

Figure 7-21. Standard 8 x 8 Batcher Sorting Matrix Configuration

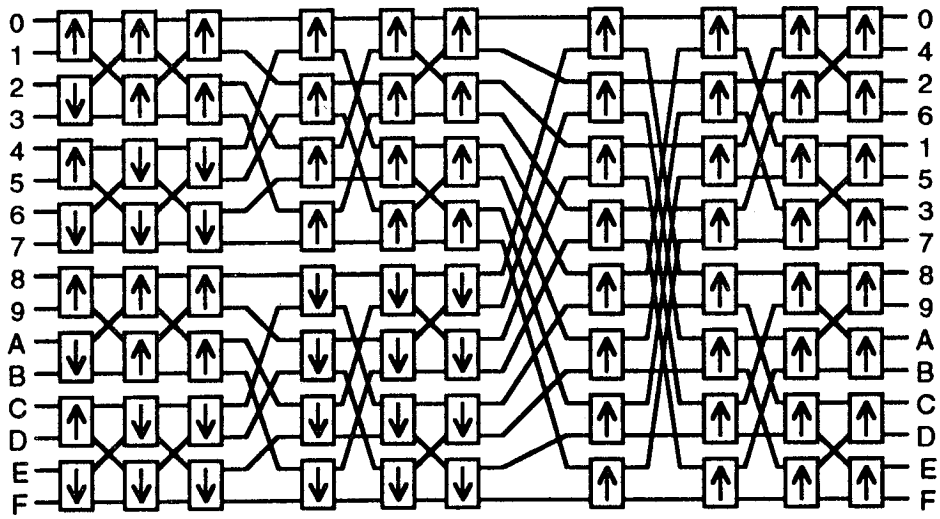

Figure 7-22. Standard 16 x 16 Batcher Sorting Matrix Configuration

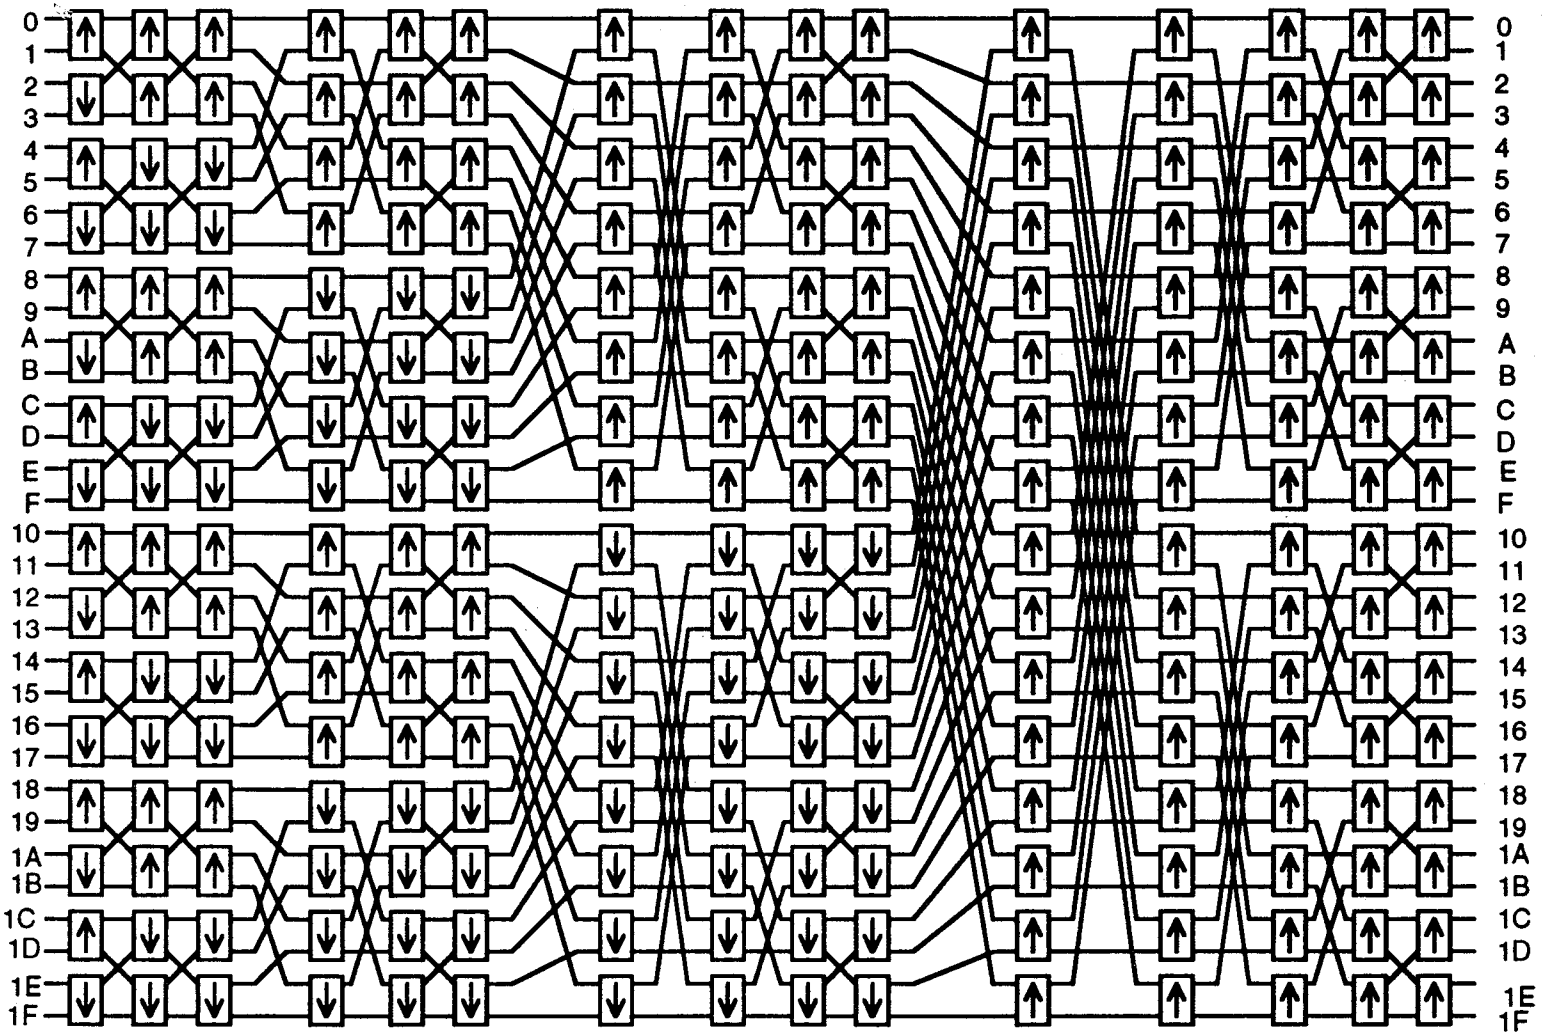

Figure 7-23. Standard 32 x 32 Batcher Sorting Matrix Configuration

The ASIC design for the Batcher matrices will incorporate multiple 2 x 2 cells to reduce the physical size and power consumption overhead of the Switch Fabric. The ASIC will be configurable in the four modes shown in Figure 7-24. The differences between each of the two versions for the MODE B and MODE D configurations are merely pin assignment changes for the ASIC.

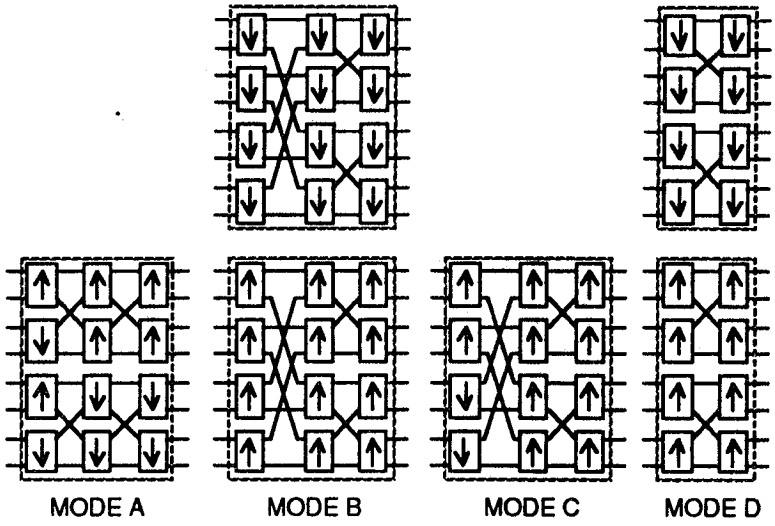

Figure 7-24. Batcher ASIC Configurations

The three Batcher matrix configurations can be implemented using the four modes of the Batcher ASIC shown above. The 8 x 8 Batcher matrix is the MODE A and MODE B ASIC connected as shown in Figure 7-25. The topology of the Standard 16 x 16 Batcher matrix can be rearranged as shown in Figure 7-26 to permit the use of two MODE A, two MODE B and four MODE D ASICs. The arrows show the translation of the highlighted cells from their original position in the same column of the standard configuration. The 32 x 32 Batcher matrix can be rearranged as shown in Figure 7-27 to permit the use of four MODE A, twelve MODE B, and four MODE C ASICs. Figure 7-28 and Figure 7-29 show the translation of the cells to provide the implemented topology. Figure 7-28 is the standard Batcher matrix where the cells which move have been numbered. Figure 7-29 show the new locations of these cells. Note that cells do not move from column to column, only from row to row within the column, so that the numbering of cells within a column will always restart at 0.

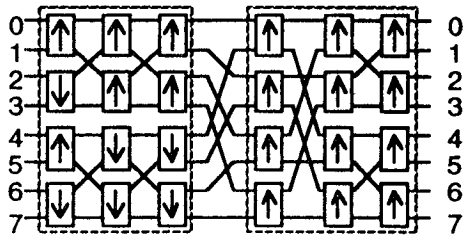

Figure 7-25. 8 x 8 Batcher Matrix Implementation

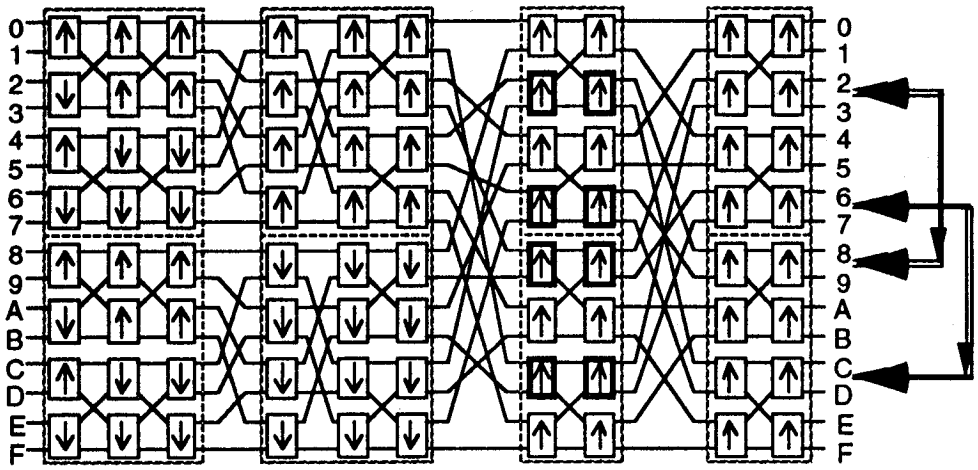

Figure 7-26. 16 x 16 Batcher Matrix Implementation

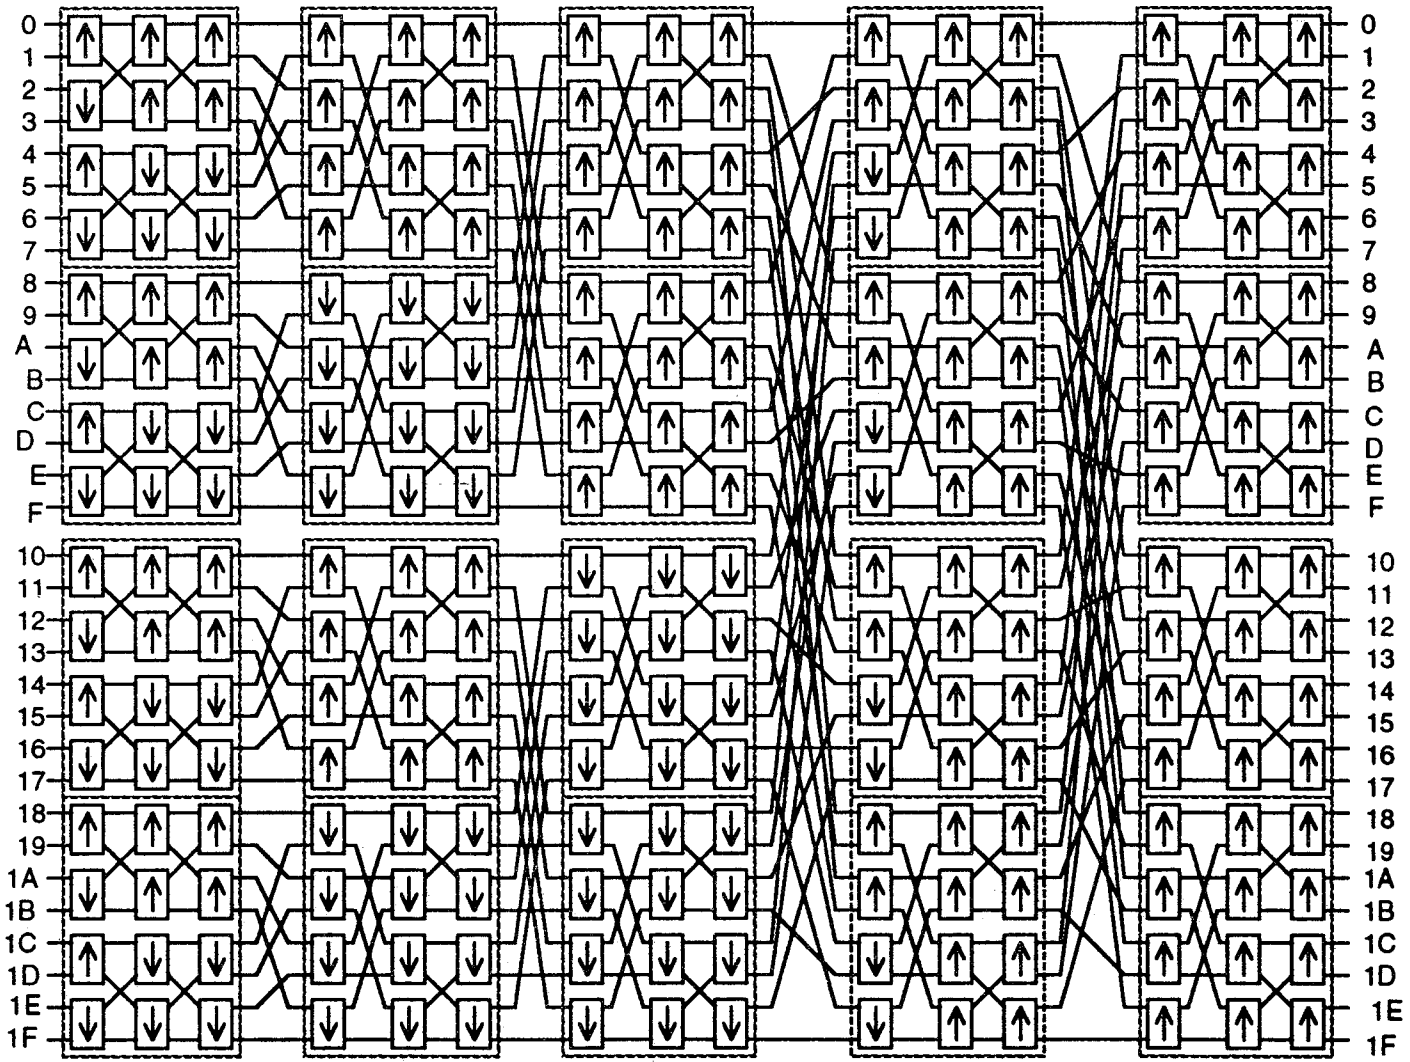

Figure 7-27. 32 x 32 Batcher Sorting Matrix Implementation

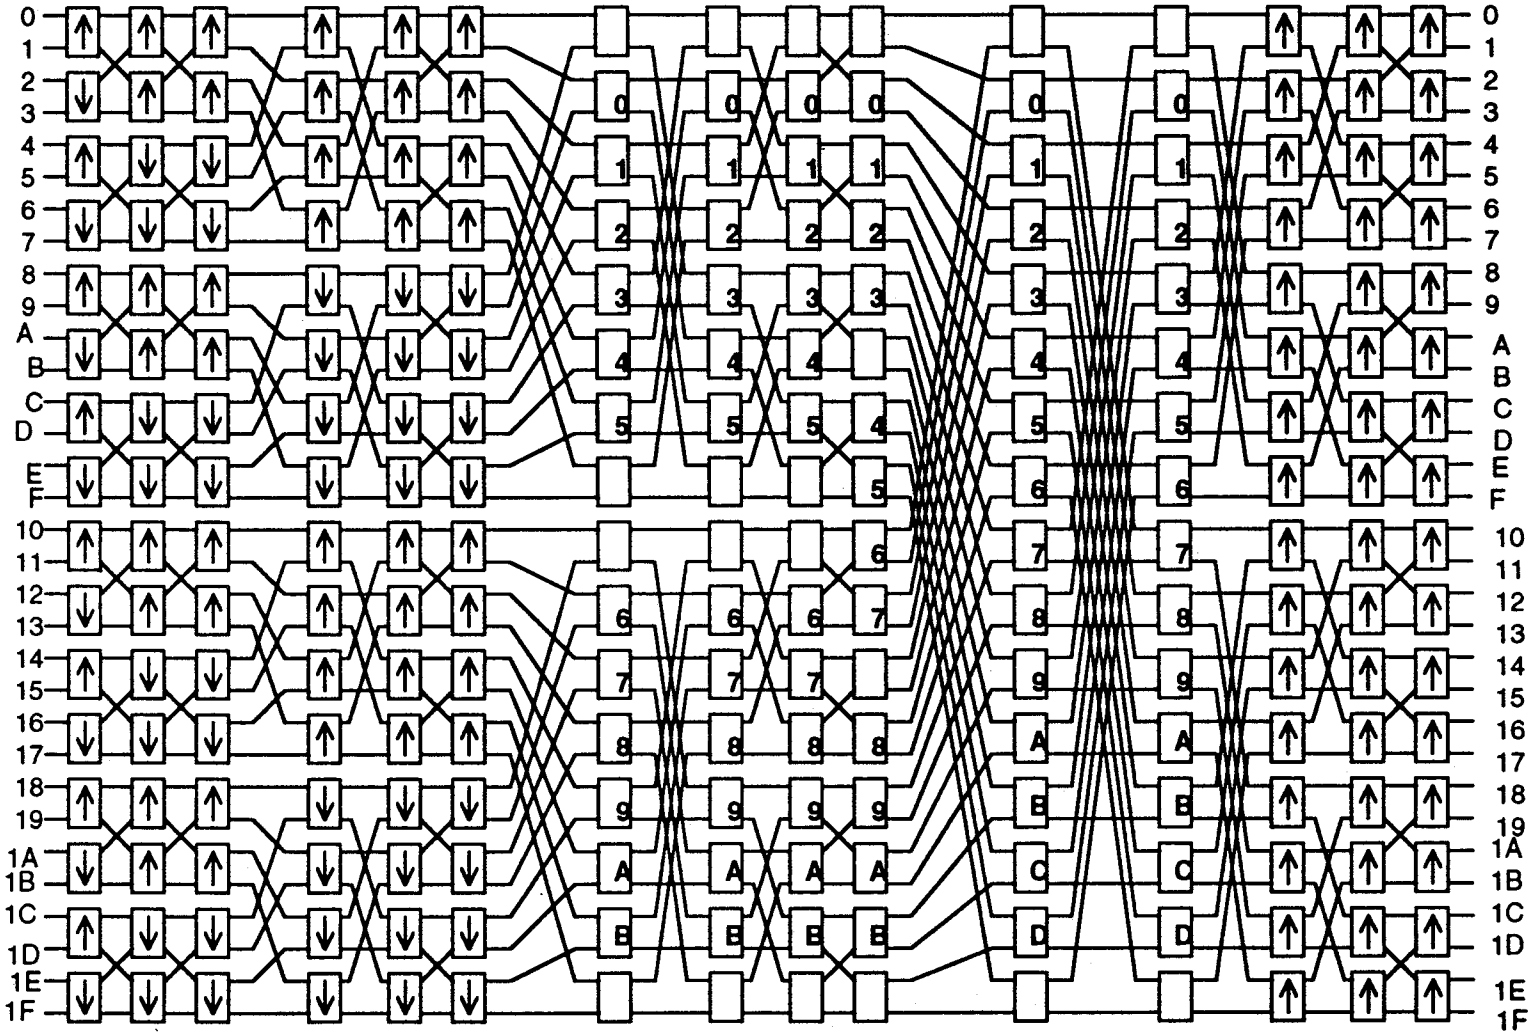

Figure 7-28. Standard 32 x 32 Batcher Sorting Matrix Configuration With Numbered Cells

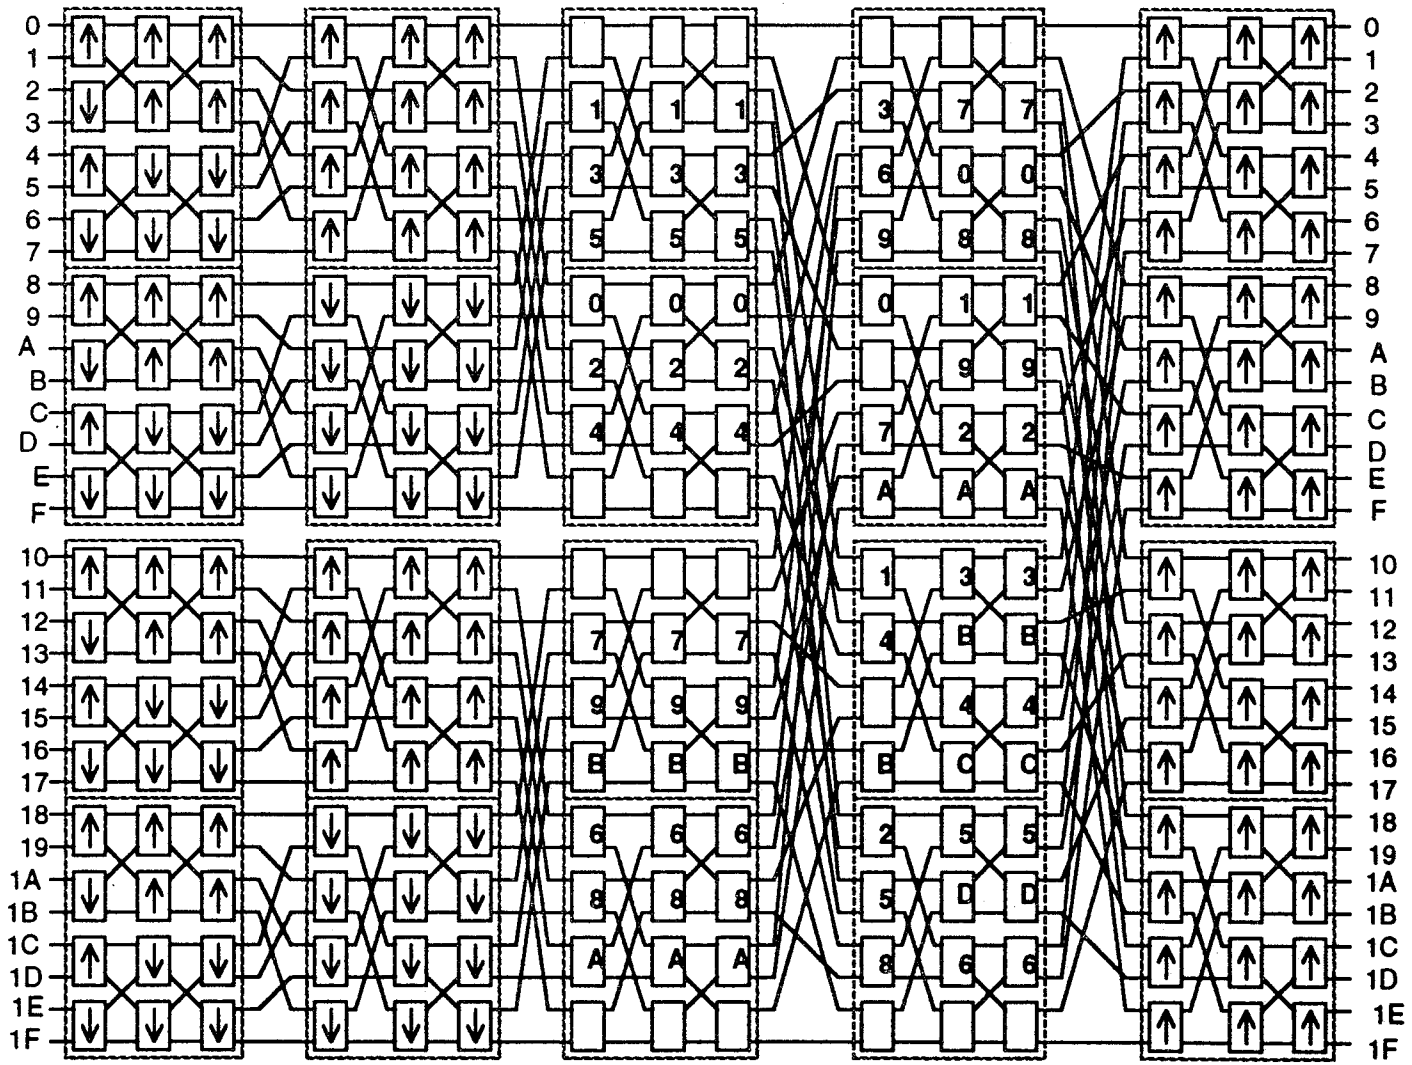

Figure 7-29. 32 x 32 Batcher Sorting Matrix Implementation With Numbered Cells

7.1.1.4.1.3 Sorted Banyan Switch Fabric

The Sorted Banyan Switch Fabric is formed by combining the Banyan and Batcher matrices. Configuration for an 8 x 8, 16 x 16, and 32 x 32 is shown in Figures 7-30 to 7-32.

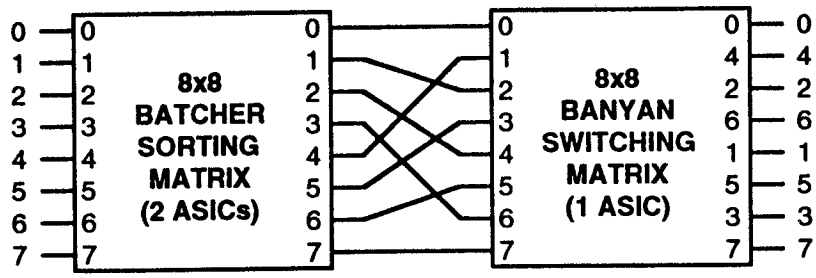

Figure 7-30. 8 x 8 Sorted Banyan Switch Fabric

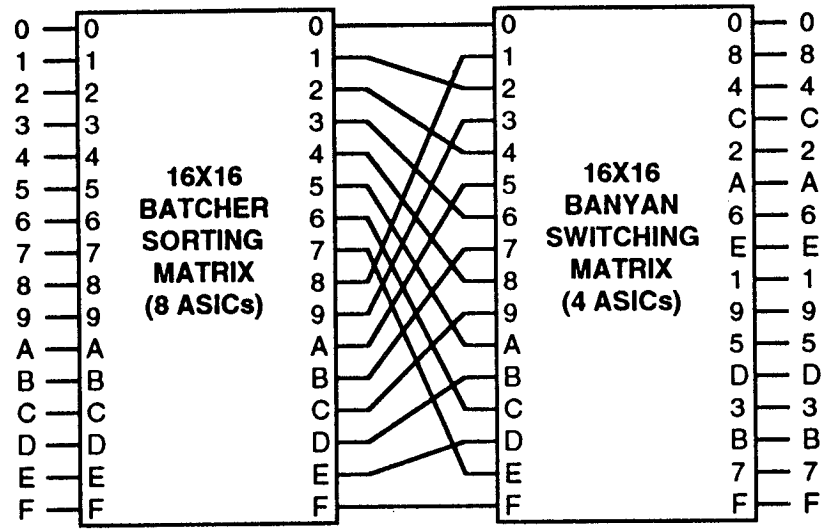

Figure 7-31. 16 x 16 Sorted Banyan Switch Fabric

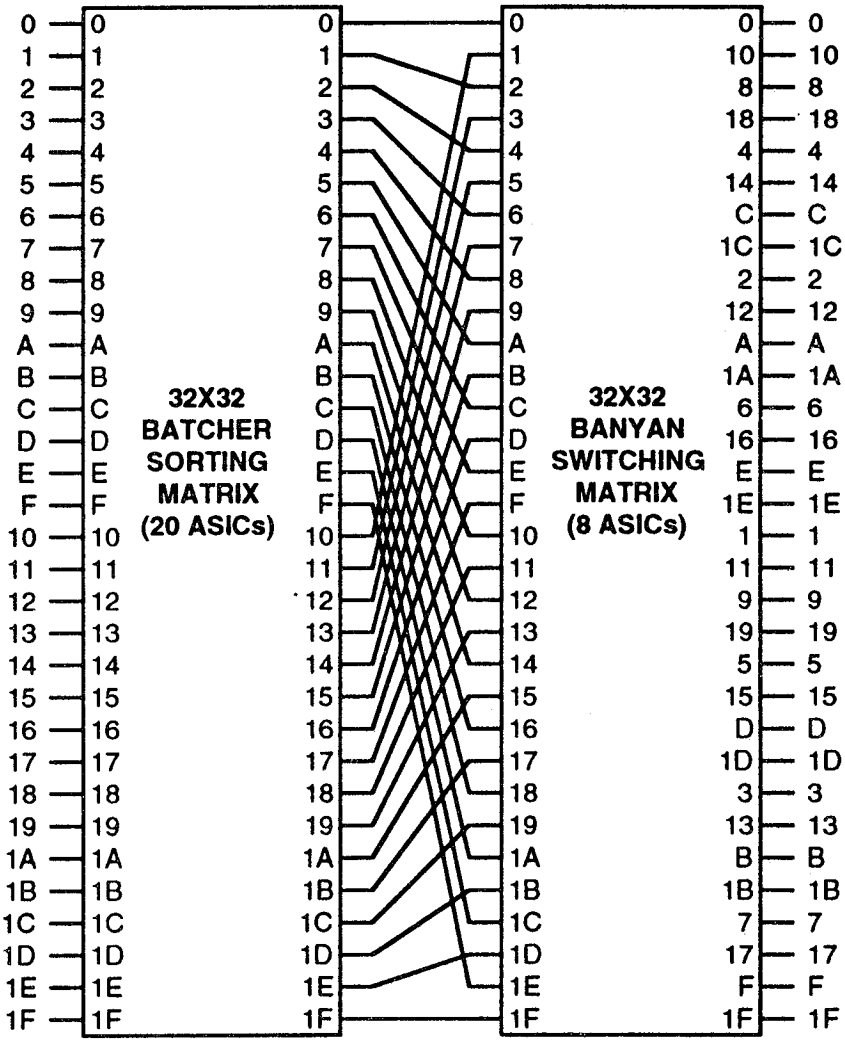

Figure 7-32. 32 x 32 Sorted Banyan Switch Fabric

7.1.1.4.2 Banyan ASIC Design

A block diagram of the Banyan ASIC is shown in Figure 7-33.

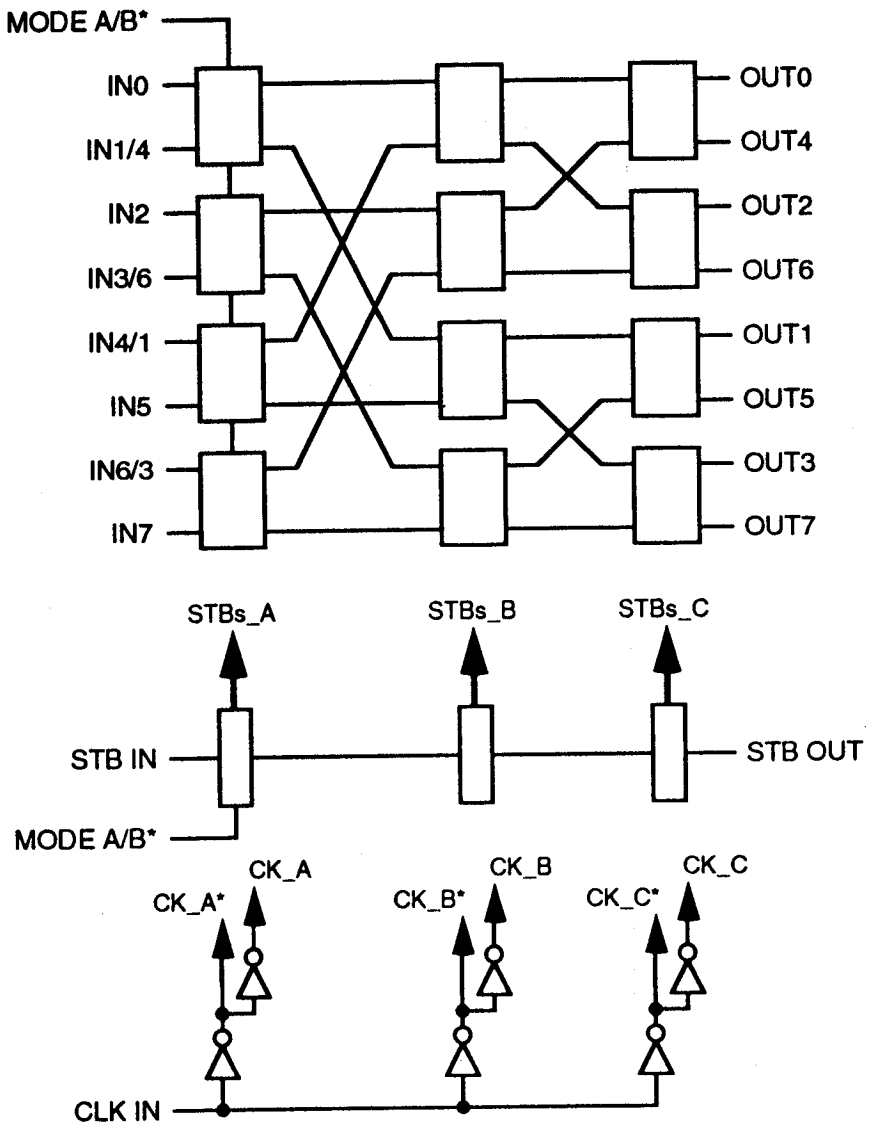

Figure 7-33. Banyan ASIC Block Diagram

Mode Control, MODE A/B*. The two modes of the Banyan ASIC are controlled by MODE A/B*. When high, the ASIC is configured as an 8 x 8 Banyan switching element. When in this mode, the inputs to the switch shown in the block diagram are defined by the first of the two numbers separated by a slash. When MODE A/B* is low, the leftmost column of switching cells are forced to pass data straight through (IN_A to OUT_A; IN_B to OUT_B). In this mode, the inputs to the switch are defined by the second number as shown in the block diagram. See Figures 7-34 and 7-35 for the configuration for these two modes. The MODE B ASIC will still have the delay of three cells since the leftmost cell is not bypassed but simply forced to a fixed configuration.

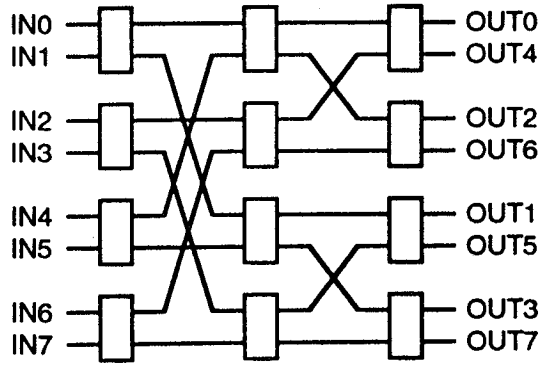

Figure 7-34. MODE A Banyan Routing

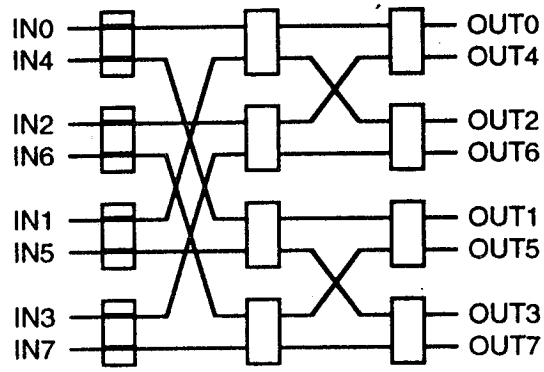

Figure 7-35. MODE B Banyan Routing

Switching Cell. A block diagram of the Banyan switching cell is shown in Figure 7-36. A detailed block diagram and timing diagram of the implementation is shown in Figures 7-37 and 7-38.

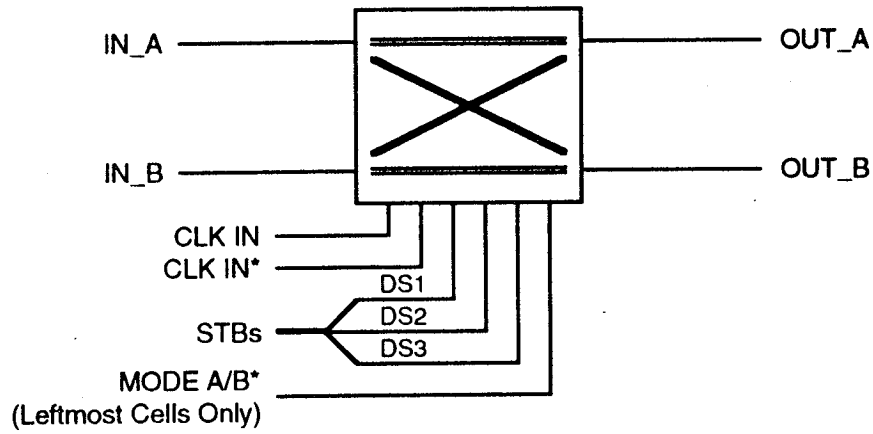

Figure 7-36. Banyan Switching Cell Block Diagram

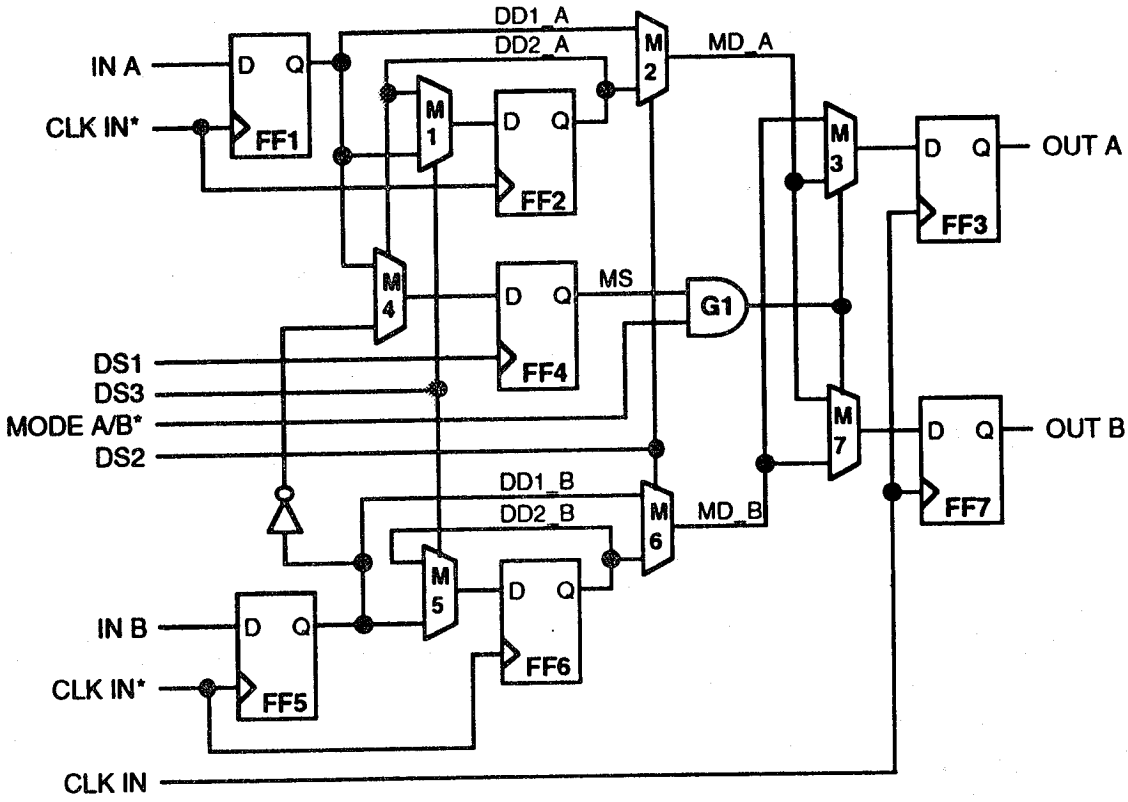

Figure 7-37. Banyan Switching Cell Detailed Block Diagram

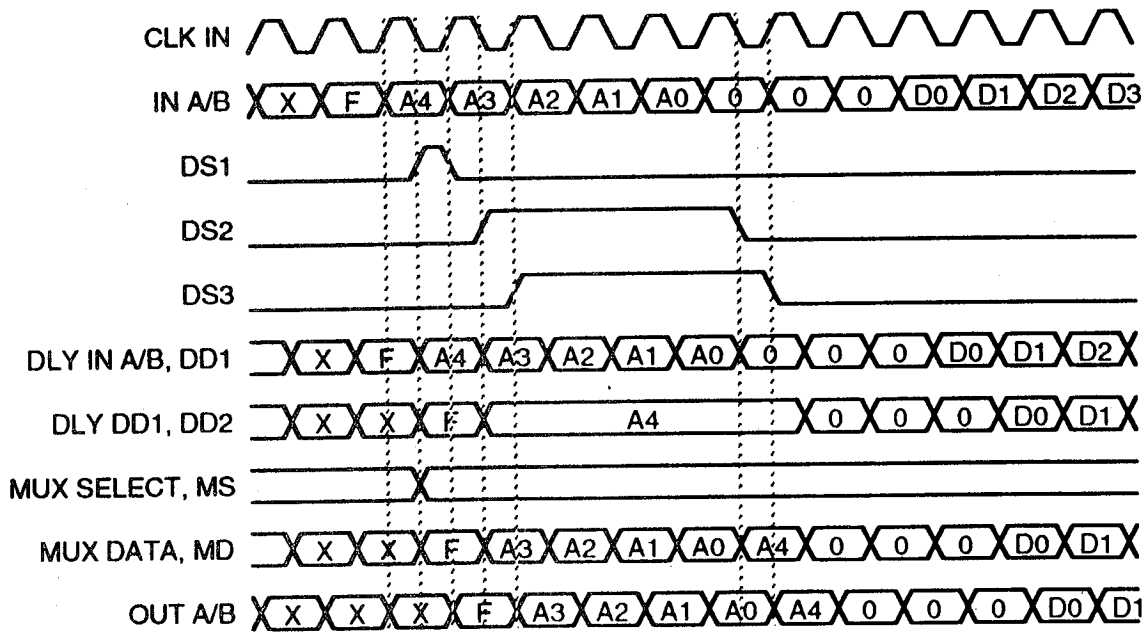

Figure 7-38. Banyan Switching Cell Timing Diagram

Flip-flop FF1 mid-bit samples the data input, IN A, to provide 1/2 clock delay data, DD1. Multiplexor M1 and flip-flop FF2 provide the storage of the first address bit and the delay of DD1 for the non-address bits. Multiplexor M2 will output on MD_A/B the desired address rotation as shown in the timing diagram. Flip-flops FF5 and FF6 and multiplexors M5 and M6 provide the same function for the B side data, IN B.

Transparent latch FF4 stores the switch selection for the duration of the packet transfer. The selection available is either straight through (IN A to OUT A; IN B to OUT B) or cross (IN A to OUT B; IN B to OUT A). The selection is determined by the following algorithm:

- If the flag bit, F, on DD2_A is set, the first address bit of the routing tag received on DD1_A will determine the switch configuration. Straight is selected if the first address bit is "0" and cross is selected if the first address bit is "1".
- If the flag bit, F, on DD2_A is reset, the first address bit of the routing tag received on DD1_B will determine the switch configuration. Straight is selected if the first address bit is "1" and cross is selected if the first address bit is "0".

Latch FF4 is transparent when the DS1 strobe is high. At this time, the select input to multiplexor M4 is the flag bit, F, from DD2_A. If it is high, DD1_A is used to determine the state of FF4; if low, DD1_B is used to determine the state of FF4. The output of FF4 will latch when DS1 goes low then remain in this state until the next DS1 strobe. Note that the A side data has priority over the B side data. If there is any conflict, i.e., both inputs wish to go to the same output, the A side will always win. This should not pose a problem since collisions should never occur due to the congestion control scheme implemented in the Input Port modules.

The output of flip-flop FF4, MS, controls the two output multiplexors, M3 and M7. When MS is low, the configuration of the switch cell will be straight; when high, it will be cross. The MODE A/B* signal influences the switch cell at this point. When low, it will force the configuration of the cell to be straight regardless of the state of FF4. This gate is only present on the leftmost column of cells; in all other cells the output of FF4 controls the output multiplexors directly.

The flip-flops FF3 and FF7 reclock the data from the two output multiplexors to retime them to switch on the rising edge of the input clock, CLK IN.

Common Control/Timing. The input clock, CLK IN, is buffered and provided to each cell in true and inverted form. Figure 7-39 is a block diagram of the circuitry required to generate the strobe signals, STBs, for a given column of switching cells. The timing diagram is shown in Figure 7-40.

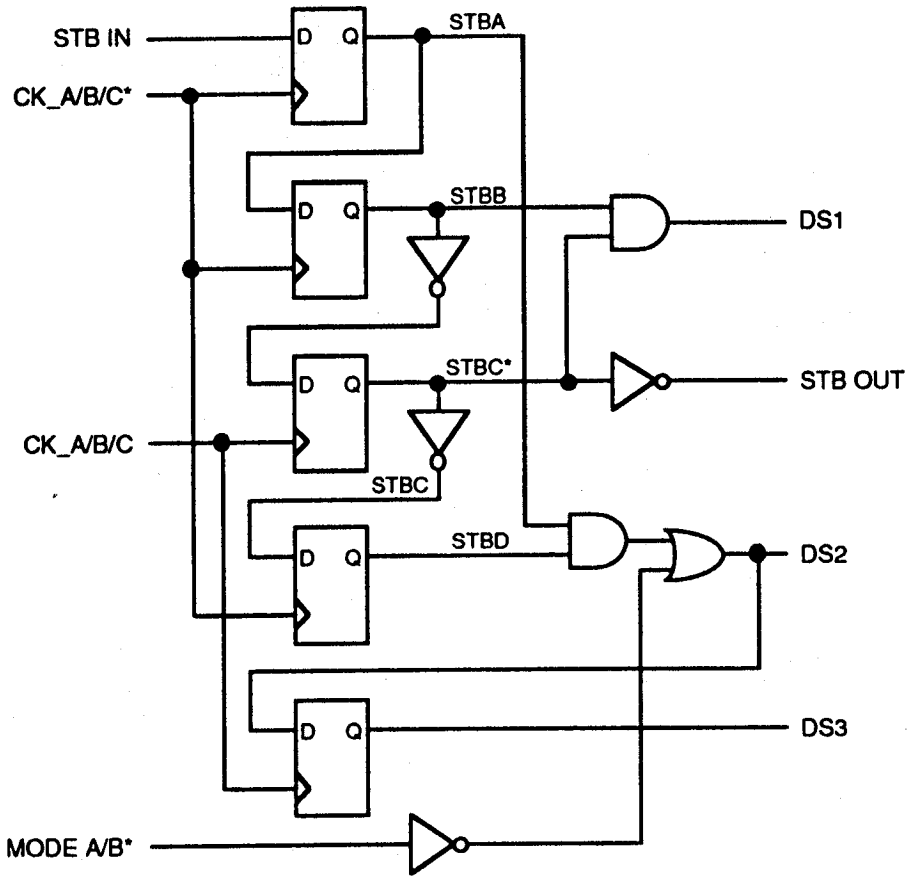

Figure 7-39. Strobe Generation Block Diagram

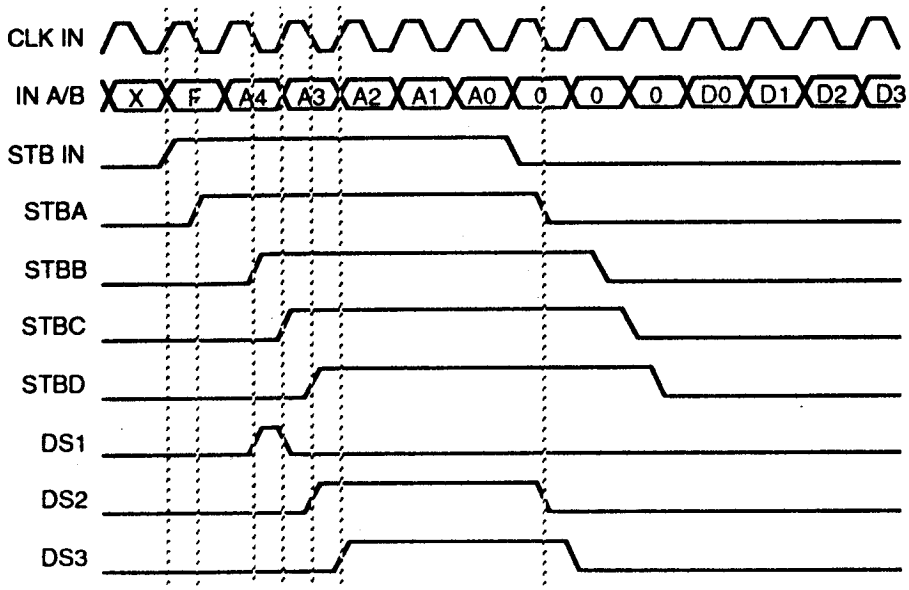

Figure 7-40. Strobe Generation Timing Diagram

C-4

The strobe signals are generated from a string of D-type flip-flops to provide the desired signals. The input to the third stage is inverted to allow the use of inverters, rather than buffers, at the output of the flip-flop. Inverters tend to have much less delay than buffers. The MODE A/B* signal (and associated or-gate) is used for the leftmost column of cells only. When low, the MODE A/B* signal will force DS2 high causing M2/6 in the switching cell to continually pass DD1_A/B to MD_A/B.

7.1.1.4.3 Banyan ASIC Characteristics

Banyan ASIC Gate Count. The estimated gate count for the Banyan ASIC is shown in Table 7-7.

Table 7-7. Banyan ASIC Estimated Gate Count

FUNCTION	GATES/FNCTN	QUANTITY	TOTAL GATES
D FLIP-FLOP	12	100	1200
LOGIC GATE	2	400	800
TOTAL + 20%			2400

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Banyan ASIC Power Consumption. The estimated power consumption for the Banyan ASIC is shown in Table 7-8.

Table 7-8. Banyan ASIC Estimated Power Consumption

FUNCTION	POWRE/DEVICE	QUANTITY	TOTAL POWER
INPUT I/O	6.4	10	0.064 W
OUTPUT I/O	7.81	10	0.078 W
D FLIP-FLOP	2.77	100	0.277 W
LOGIC GATE	.4	400	0.160 W
TOTAL + 20%			0.7 W

7.1.1.4.4 Batcher ASIC Design

The four Batcher configurations required to implement all Batcher Sorting Matrices are shown in Figure 7-41.

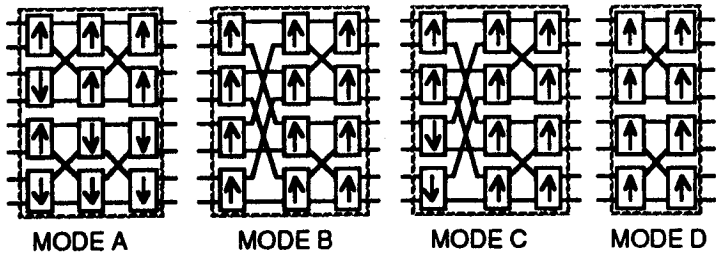

Figure 7-41. Batcher ASIC Modes

The MODE A version can be reconfigured to more easily match the other modes as shown in Figure 7-42. Figure *a* shows the ASIC as given above with inputs and outputs numbered. Figure *b* inverts the lower half of the ASIC. Figure *c* transposes the sorting direction for the second and third cells in the leftmost column. The inputs to a sorting cell can be transposed without affecting the operation of the cell. Figure *d* transposes the inputs to the second and third cells of the middle column.

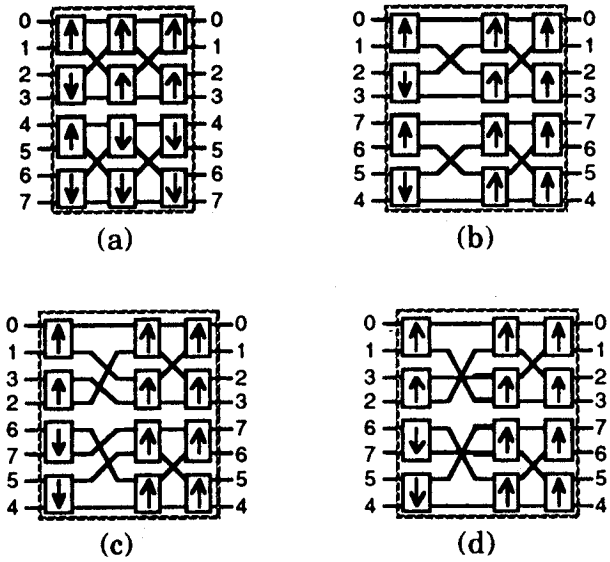

Figure 7-42. MODE A ASIC Reconfiguration Steps

The four modes to be implemented in the Batcher ASIC are shown in Figure 7-43.

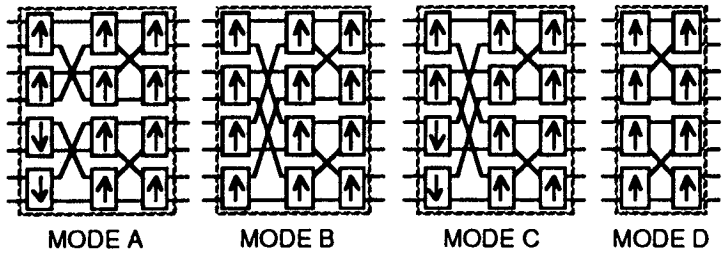

Figure 7-43. Batcher Configurations To Be Implemented In ASIC

A block diagram of the Batcher ASIC is shown in Figure 7-44.

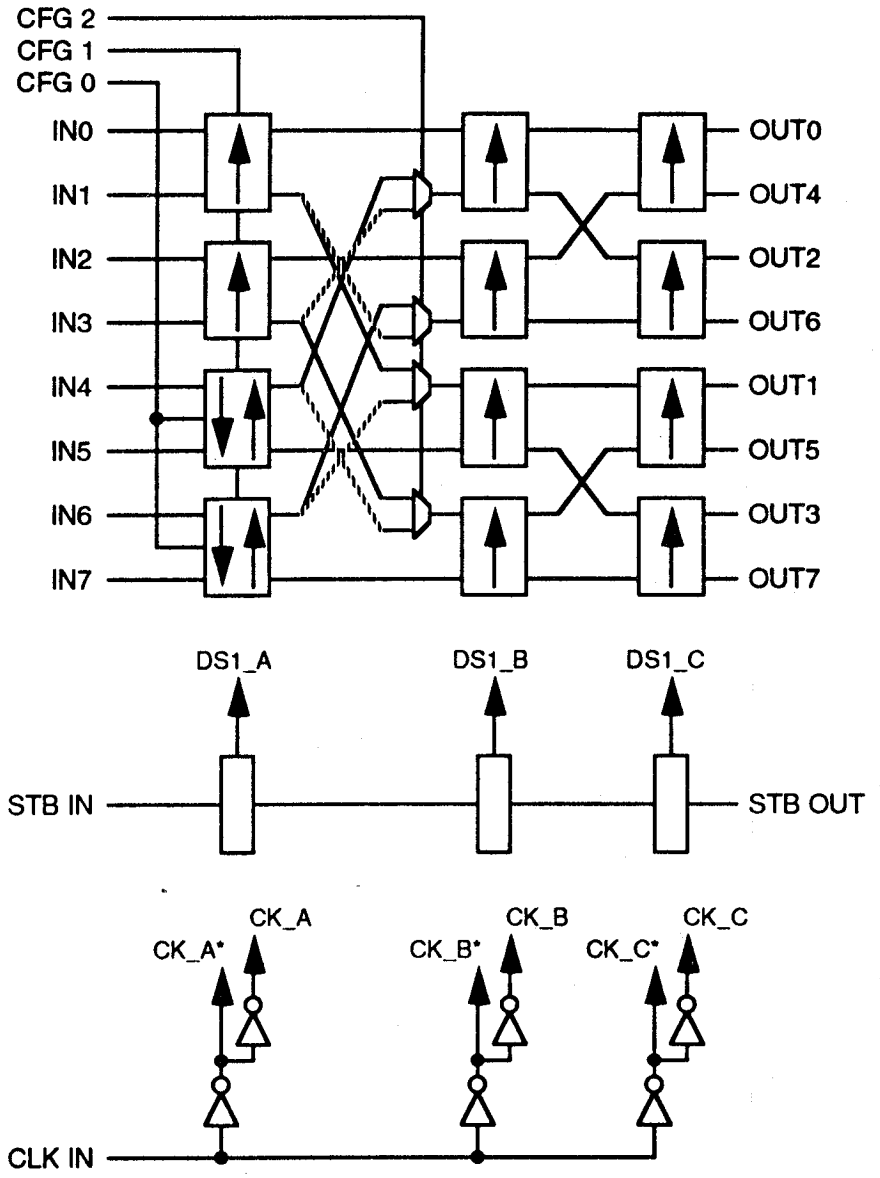

Figure 7-44. Batcher ASIC Block Diagram

Mode Control: CFG 2,1,0. The three mode control signals, CFG 2, CFG 1, and CFG 0, determine the configuration of the Batcher ASIC. CFG 2 controls the interconnect between the leftmost and middle column of cells. When high, the ASIC is configured for MODE A or MODE D; when low, it is configured for the other two modes. CFG 1 disables the sorting operation of the leftmost cells. When low, these cells are configured for MODE D as always straight through (top input to top output; bottom input to bottom output); when high, these cells operate in their normal sorting mode. CFG 0 determines the type of sorting performed by the two lowest cells in the leftmost column. When low, these cells operate where the larger input is routed to the top output; when high, the

larger input is routed to the bottom output. Figures 7-45 to 7-48 show the effective configurations for the valid combinations of the mode control bits.

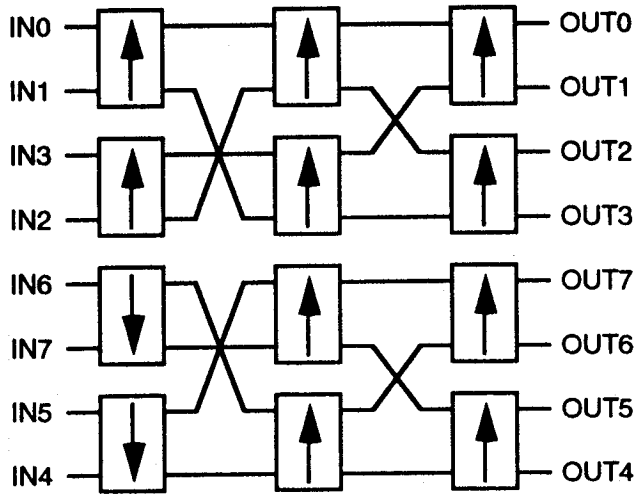

Figure 7-45. MODE A ASIC Configuration (CFG 2,1,0 = "111")

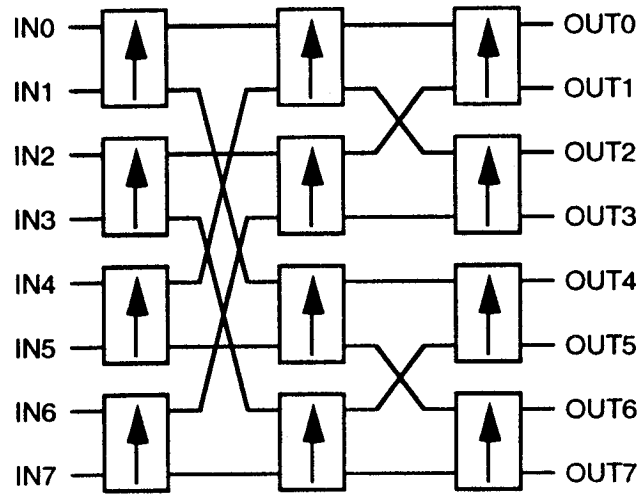

Figure 7-46. MODE B ASIC Configuration (CFG 2,1,0 = "010")

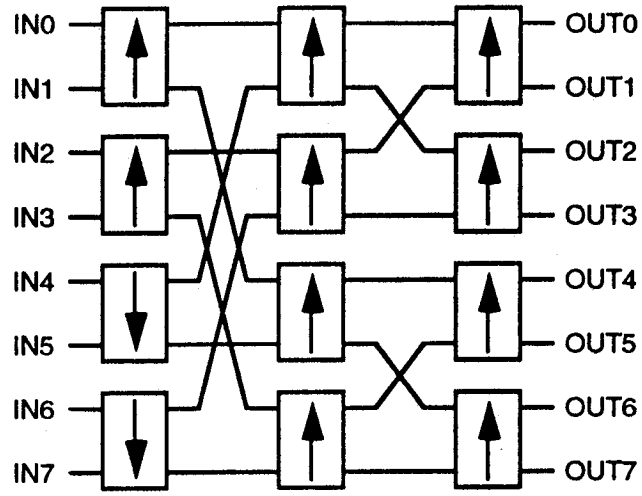

Figure 7-47. MODE C ASIC Configuration (CFG 2,1,0 = "011")

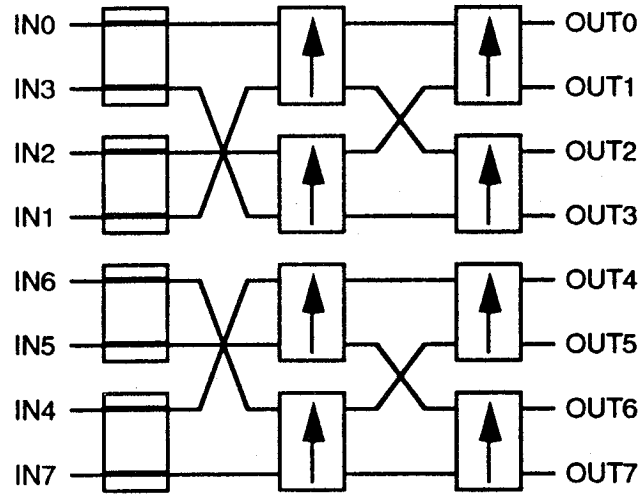

Figure 7-48. MODE D ASIC Configuration (CFG 2,1,0 = "110")

Sorting Cell. A block diagram of the Batcher sorting cell is shown in Figure 7-49. A detailed block diagram and timing diagram of the implementation is shown in Figures 7-50 and 7-51.

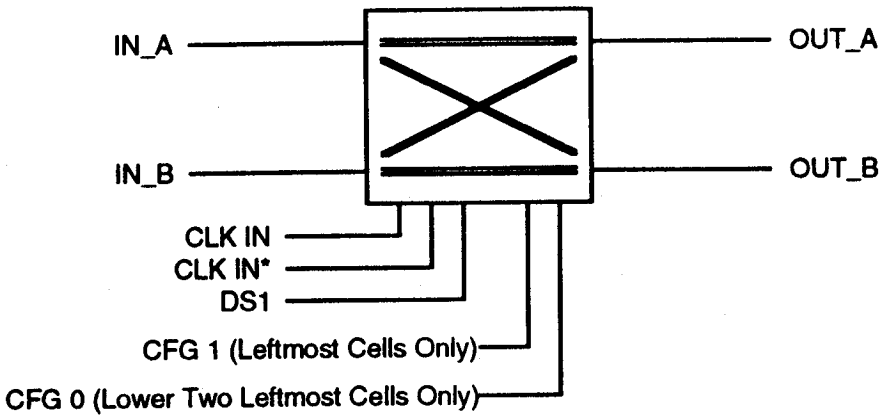

Figure 7-49. Batcher Sorting Cell Block Diagram

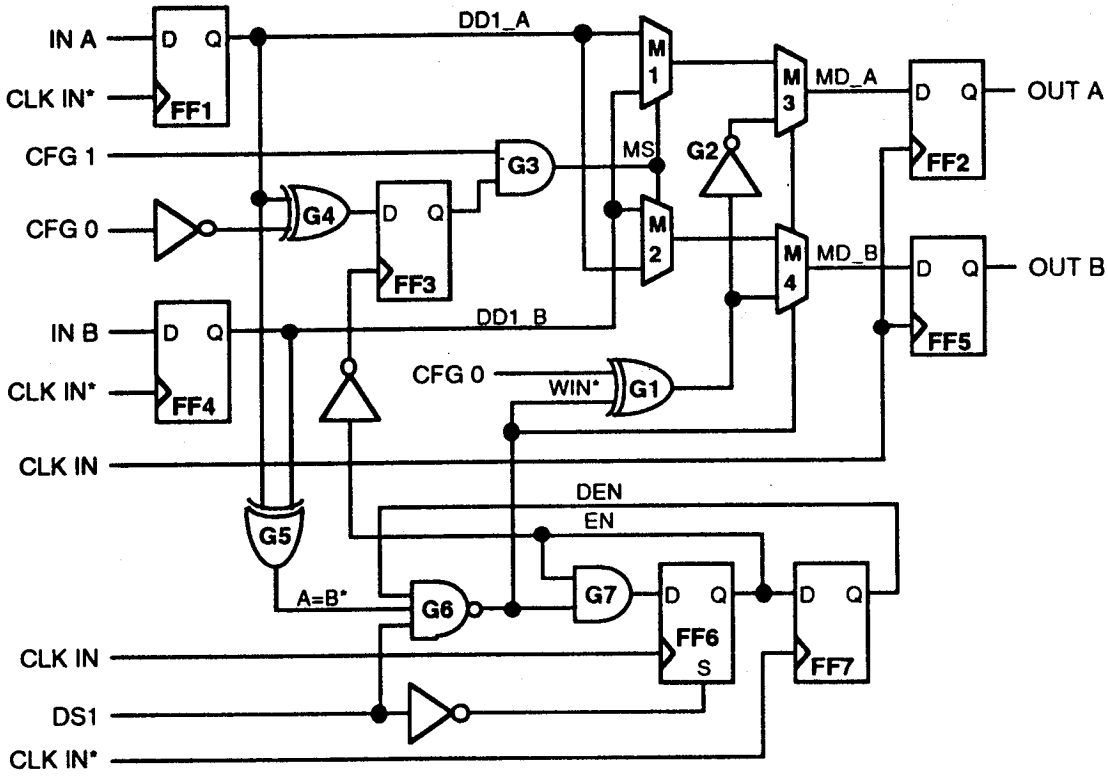

Figure 7-50. Batcher Sorting Cell Detailed Block Diagram

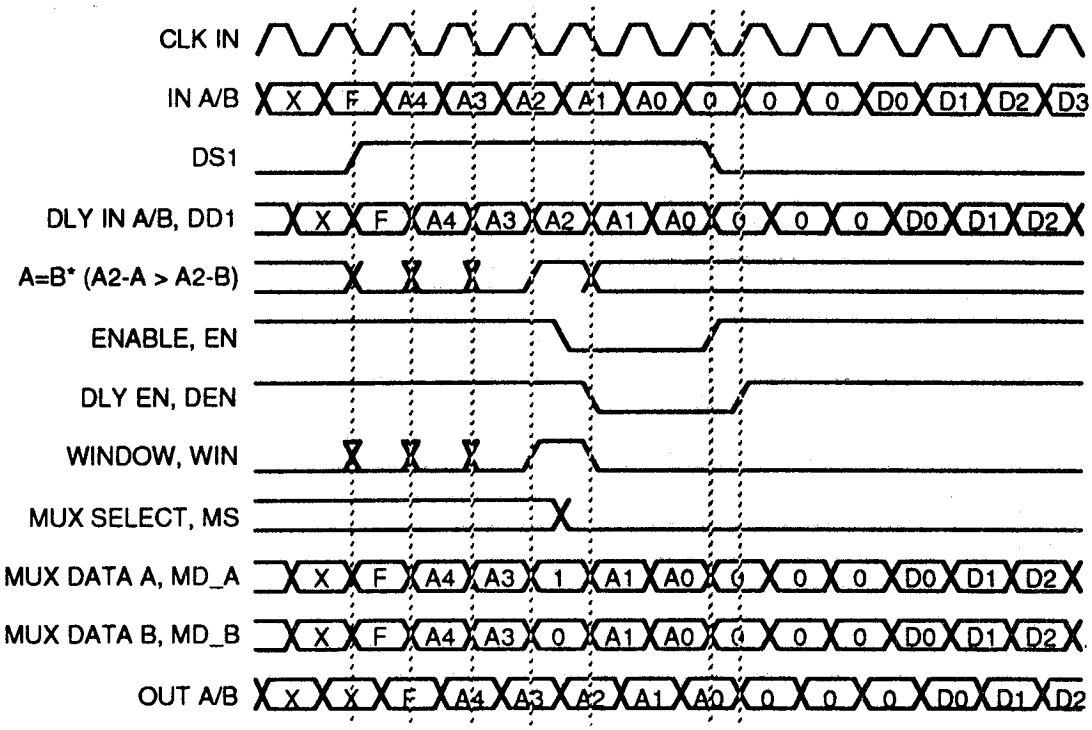

Figure 7-51. Batcher Sorting Cell Timing Diagram

Flip-flop FF1/4 mid-bit samples the input data, IN A/B. Multiplexor M1/2 provide the selection to output OUT A/B. Flip-flop FF2/5 reclocks the serial data for output to the next stage.

Flip-flop FF3 determines the switch configuration. When low, the switch is configured straight through (IN A to OUT A; IN B to OUT B); when high, the switch is configured as cross (IN A to OUT B; IN B to OUT A). The gates G5, G6, and G7 and the flip-flops FF6 and FF7 provide the circuitry to generate the control signals EN, DEN, and WIN shown on the timing diagram. When EN goes low, FF3 will latch the inversion of DD1_A. Therefore, if DD1_A is low at this time (implies DD1_B is high), FF3 will latch a high, causing the switch to be in the cross configuration which will route the A side input (lesser address) to the B side output (and the B side input to the A side output).

Multiplexors M1 and M2 and gates G1 and G2 provide the output bits when the multiplexor select bit, MS, is latched. When WIN is inactive, the inputs to FF2 and FF5 are from M1 and M2; when WIN is active, the inputs are selected from G1 and G2. G1 is present only for the two lower cells in the leftmost column; in all other cells, G1 is removed and G2 is fed directly by WIN*. When WIN is active, the inputs of FF2 and FF5 will be forced to a "1" and a "0", respectively, if CFG 0 is low or when G1 is not used; when CFG 0 is high and G1 is used, the inputs will be forced to a "0" and a "1", respectively.

Gate G3 is present only in the leftmost cells within the ASIC. With all other cells, the output of FF3 controls MS directly. When G3 is present, it is used to disable the output of FF3. When the MODE D configuration is used, the leftmost cells are forced into the

straight through configuration continuously. This is accomplished by bringing CFG 1 low, disabling the output of FF3 and forcing a low on MS continuously.

Gate G4 is present only in the bottom two cells of the leftmost column within the ASIC. With all other cells G4 is replaced by an inverter with input only from DD1_A. When present, G4 is used to determine the sorting direction of the cell. If CFG 0 is high, the cell will operate as all the other cells where the larger address is routed to the output OUT A. If CFG 0 is low, the larger address will be routed to the output OUT B, i.e., the cell will sort down instead of up.

Common Control/Timing. The input clock, CLK IN, is buffered and provided to each cell in true and inverted form. Figure 7-52 is a block diagram of the circuitry required to generate the strobe signal, DS1, for a given column of switching cells. The timing diagram is shown in Figure 7-53.

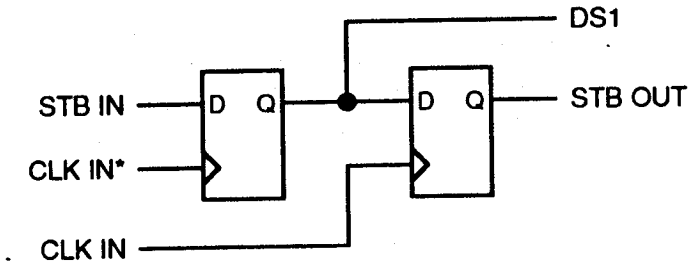

Figure 7-52. Strobe Generation Block Diagram

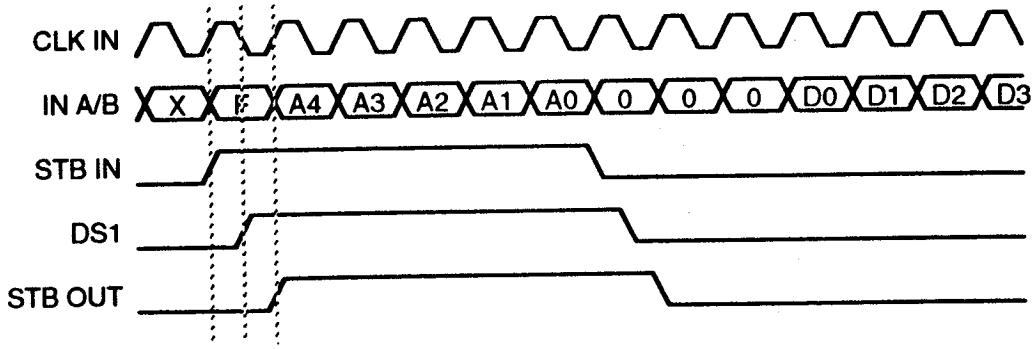

Figure 7-53. Strobe Generation Timing Diagram

The timing signals, DS1 and STB OUT, are generated from a string of two D-type flip-flops. DS1 is simply STB IN delayed by half of the clock and STB OUT is DS1 delayed by another half of the clock.

7.1.1.4.5 Batcher ASIC Characteristics

Batcher ASIC Gate Count. The estimated gate count for the Batcher ASIC is shown in Table 7-9.

Table 7-9. Batcher ASIC Estimated Gate Count

FUNCTION	GATES/FNCTN	QUANTITY	TOTAL GATES
D FLIP-FLOP	12	90	1080
LOGIC GATE	6	180	900
TOTAL + 20%			2400

The logic gate count includes inverters and 2 to 4 input and, nand, or, nor, and xor logic.

Batcher ASIC Power Consumption. The estimated power consumption for the Batcher ASIC is shown in Table 7-10.

Table 7-10. Batcher ASIC Estimated Power Consumption

FUNCTION	POWRE/DEVICE	QUANTITY	TOTAL POWER
INPUT I/O	6.4	13	0.083 W
OUTPUT I/O	7.81	10	0.078 W
D FLIP-FLOP	2.77	90	0.249 W
LOGIC GATE	1.2	150	0.180 W
TOTAL + 20%			0.7 W

7.1.1.4.6 Switch Fabric Characteristics

ASIC Count. The number of ASICs required for each of the three switch fabric configurations is shown in Table 7-11.

Table 7-11. Switch Fabric ASIC Count

	8 x 8	16 x 16	32 x 32
BANYAN	1	4	8
BATCHER	2	8	20
TOTAL	3	12	28

Switch Fabric Power Consumption. The power consumption of the Switch Fabric is summarized in Table 7-12 for the three configurations. Miscellaneous support circuitry is estimated at about 0.5 watts per port.

Table 7-12. Switch Fabric Power Consumption

	8 x 8	16 x 16	32 x 32
BANYAN MATRIX	.7 W	2.8 W	5.6 W
BATCHER MATRIX	1.4 W	5.6 W	14 W
MISC SUPPORT	4 W	8 W	16 W
TOTAL	6.1 W	16.4 W	35.6 W

7.1.1.5 Sorted Banyan Summary

7.1.1.5.1 Power Consumption

The power consumption of the Sorted Banyan Switch is summarized in Table 7-13 for the three switch configurations.

Table 7-13. Sorted Banyan Switch Power Consumption

	8 x 8	16 x 16	32 x 32
INPUT PORT	38.4 W	76.8 W	153.6 W
OUTPUT PORT	27.2 W	54.4 W	108.8 W
SWITCH FABRIC	6.1 W	16.4 W	35.6 W
TOTAL	71.7 W	147.6 W	298 W
TOTAL/PORT	9 W	9.2 W	9.3 W

7.1.1.5.2 Mass/Size

A circuit board of 30 square inches will support one Input Port and one Output Port. One board will be required for every port that the switch must service. The 8 x 8 and 16 x 16 Switch Fabric configurations can also fit on one board of this size; the 32 x 32 configuration will require two boards. In addition, one board will be required for the generation of common control and timing signals. This includes the following functions:

- Generation of Switch Fabric clocks.
- Generation and control of the Congestion Control Message.
- Generation of strobes to the Input Ports, Output Ports, and the Switch Fabric.

Therefore, the total number of modules required is 10, 18, and 35 for the 8 x 8, 16 x 16, and 32 x 32 switch configurations, respectively.

7.1.1.5.3 ASIC Design Complexity

There should be no problem in designing the three ASICs required for this approach. All ASICs have fairly low gate count and moderate power consumption. A disadvantage to this approach is that three ASIC designs are required as opposed to only two for other approaches.

The design of these ASICs did not include added circuitry which may be required to implement a viable redundancy scheme. However, the addition of this circuitry is not expected to significantly increase the complexity of the ASIC design.

7.1.1.5.4 Fault Tolerance

The fault tolerance of the basic Sorted Banyan Switch is relatively good compared to other switch architectures. A failure of an Input Port will not affect any other Input Port. This is also true for the Output Ports. A failure in a cell of the Switch Fabric will only affect 1/N of the paths through the fabric where N is the size of the Switch Fabric. For these reasons, it is much more efficient to implement a 1-for-N redundancy scheme for this switch than a 1-for-1 scheme, although it may be very difficult to implement in the Switch Fabric.

The Congestion Control Message path is a ring type topology. If any Input Port fails, the entire ring becomes ineffective, causing the switch to completely shut down. For this reason, it is imperative that this ring be redundant. Also, a detected failure in a given Input Port should cause the ring input to pass through unaffected to the ring output.

7.1.2 Self-Routing Crossbar-Based Network

A block diagram of the Self-Routing Crossbar-Based Network is shown in Figure 7-54 below.

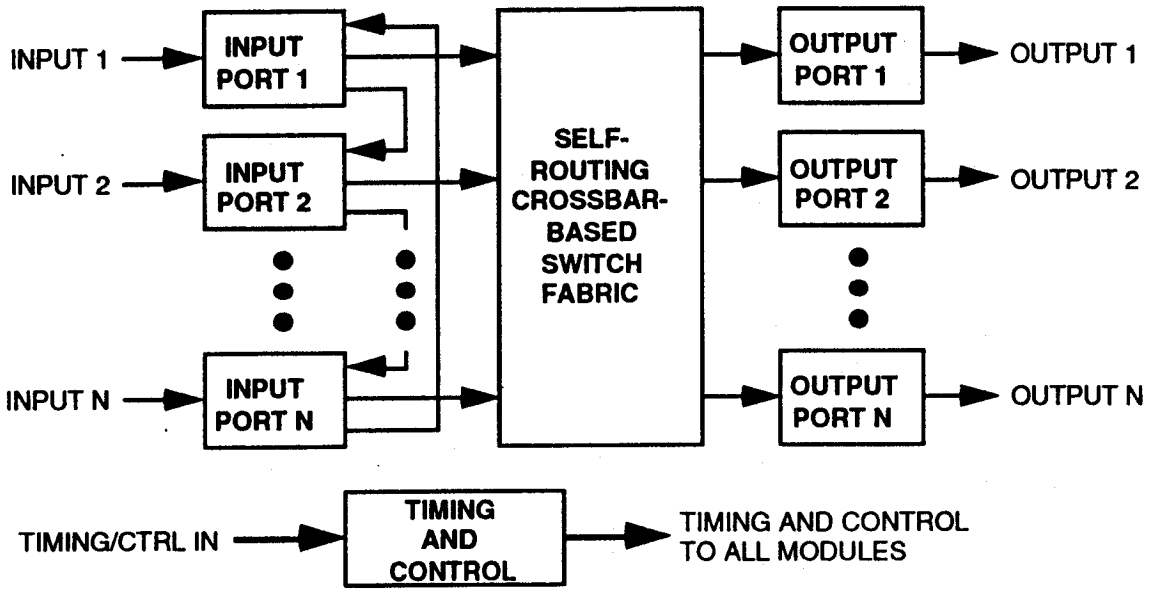

Figure 7-54. Self-Routing Crossbar-Based Network Block Diagram

Serial data from the demodulators is routed through the switch to the modulators in the following steps:

- The serial packet data enters the respective Input Port where it is stored in a 128 packet FIFO type buffer.
- As each packet is stored in the packet buffer, the routing tag associated with the packet is obtain from the routing map memory then stored with the packet in the buffer.
- The Input Port performs Output Port congestion control to ensure that there are no blocked packets at the Output Port. The Input Ports convey their desired destination to each other by means of a serial daisy chained message every packet slot time. When an Input Port reserves an Output Port for the next packet transmission, no other Input Port further down the daisy chain may select the reserved Output Port. The four oldest in the packet buffer will be examined for transmission through the Switch Fabric to minimize head-of-line blocking.
- The selected packet, if any, from each Input Port is sent to the Switch Fabric. The Switch Fabric will route packets from all Input Ports to the desired Output Ports as specified by the routing tag. The Switch Fabric will need to operate at a higher rate than the input serial data to compensate for the

added bits for the routing tag and for the lost packet transfers through the Switch Fabric due to head-of-line blocking.

- The Output Port will accept serial data from the Switch Fabric then store it in a 128 packet FIFO type buffer. The serial packet data will be output to the modulator for downlink transmission.

7.1.2.1 Formats

7.1.2.1.1 Routing Tag Format

The format for the routing tag is identical to that used for the Sorted Banyan Network described in Section 7.1.1.1.1.

7.1.2.1.2 Output Port Congestion Control Message Format

The format for the Output Port Congestion Control Message is identical to that used for the Sorted Banyan Network described in Section 7.1.1.1.2.

7.1.2.2 Input Port

The Input Port is identical to that used for the Sorted Banyan Network described in Section 7.1.1.2. A summary of the Input Port characteristics are shown below.

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is repeated here in Table 7-14.

Table 7-14. Input Port ASIC Estimated Gate Count

FUNCTION	GATES/FNCTN	QUANTITY	TOTAL GATES
D FLIP-FLOP	7	200	1400
LOGIC GATE	1.75	575	1006
TOTAL + 20%			2900

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is repeated here in Table 7-15.

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Input Port Power Characteristics. The Input Port will require one ASIC chip, two 8K by 8 bit memories, differential receivers and drivers, and minor support logic requiring an estimated ten square inches of board space. The total power consumption of the Input Port is summarized in Table 7-16.

Table 7-15. Input Port ASIC Estimated Power Consumption

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	7	4	0.580 W
OUTPUT I/O	0	18	0	4	0.530 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	125	30	25	20	0.095 W
LOGIC GATE	300	240	35	0	0.050 W
TOTAL + 20%					1.8 W

Table 7-16. Input Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.8 W
RAM	1.5 W
I/O & SUPPORT	1.5 W
TOTAL	4.8 W

7.1.2.3 Output Port

The Output Port is identical to that used for the Sorted Banyan Network described in Section 7.1.1.3. A summary of the Output Port characteristics are shown below.

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-14. This is the same ASIC as used on the Input Port and will require approximately 2900 gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-17. The values have been altered from the Input Port design to reflect the parts of the ASIC not used for the Output Port.

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Output Port Power Characteristics. The Output Port will require one ASIC chip, one 4K by 16 bit memory, differential receivers and drivers, and minor support logic requiring an estimated seven square inches of board space. The total power consumption of the Output Port is summarized in Table 7-18.

Table 7-17. Input Port ASIC Estimated Power Consumption For Output Port

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	4	4	0.440 W
OUTPUT I/O	0	18	0	1	0.345 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	105	30	25	20	0.095 W
LOGIC GATE	300	240	35	0	0.050 W
TOTAL + 20%					1.4 W

Table 7-18. Output Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.4 W
RAM	1.0 W
I/O & SUPPORT	1.0 W
TOTAL	3.4 W

7.1.2.4 Switch Fabric

7.1.2.4.1 Switch Fabric Configurations

The configurations for an 8 x 8, 16 x 16, and 32 x 32 Switch Fabric for the Crossbar-Based Network are shown in Figures 7-55 to 7-57. The Crossbar ASIC will accommodate thirty-two inputs and eight outputs. The 8 x 8, 16 x 16, and 32 x 32 Switch Fabric configurations will require one, two, and four ASICs, respectively.

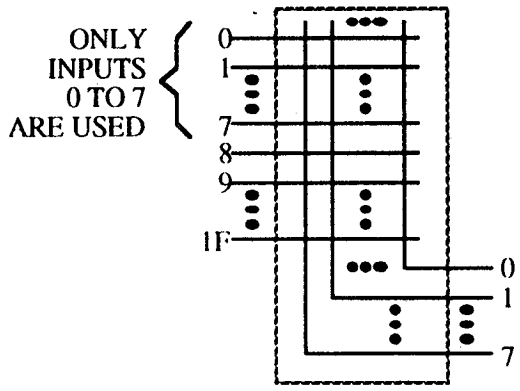

Figure 7-55. 8 x 8 Crossbar-Based Switch Fabric

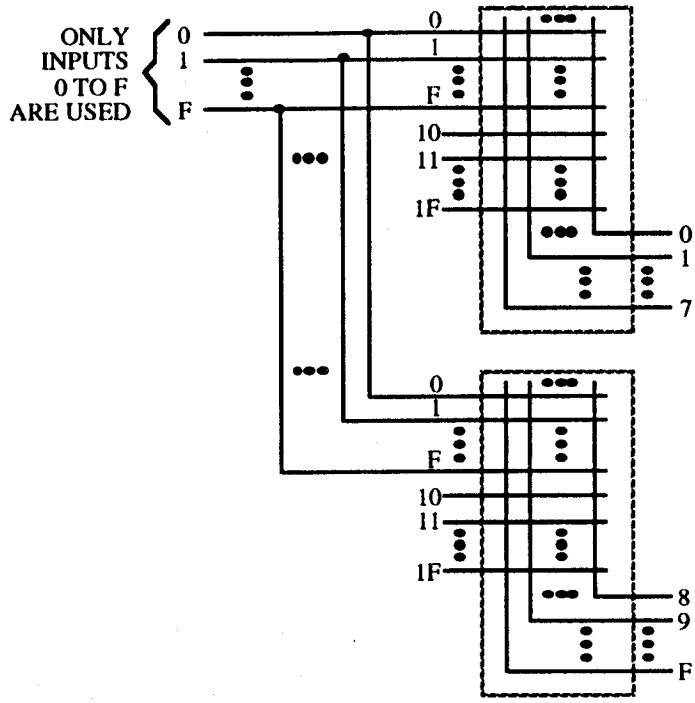

Figure 7-56. 16 x 16 Crossbar-Based Switch Fabric

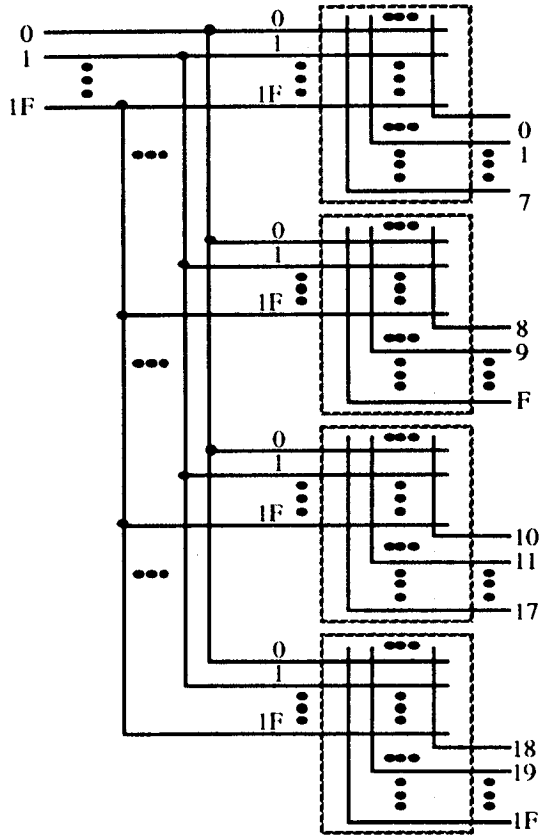

Figure 7-57. 32 x 32 Crossbar-Based Switch Fabric

7.1.2.4.2 Crossbar ASIC Design

A block diagram of the Crossbar ASIC is shown in Figure 7-58.

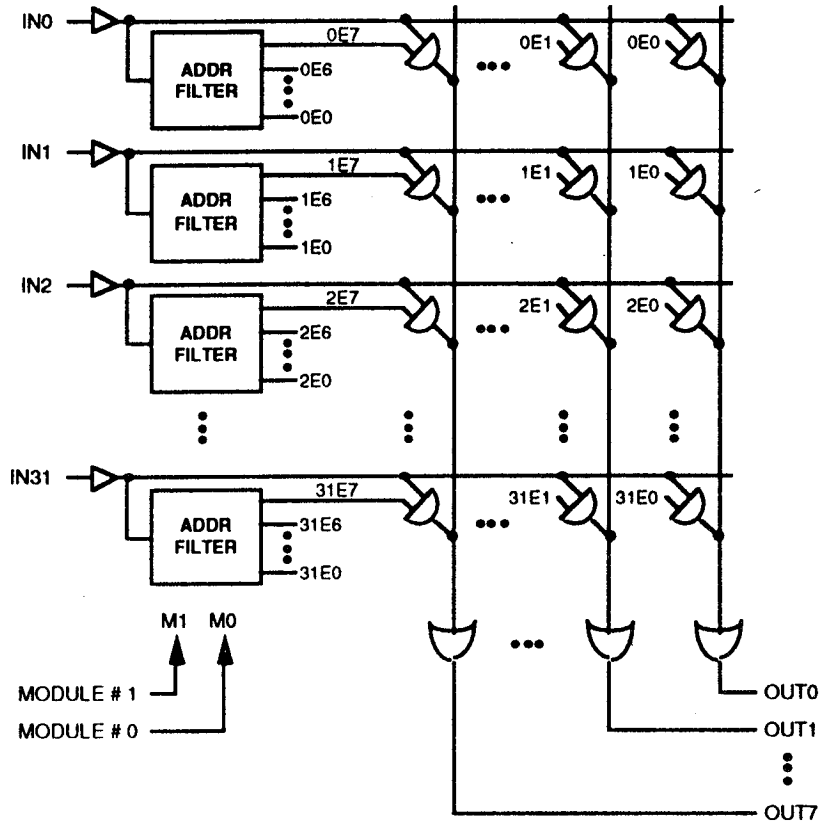

Figure 7-58. Crossbar ASIC Block Diagram

Mode Control, Module # 1 and Module #0 (M1, M0). M1 and M0 determine the definition of the eight outputs of the ASIC as shown if Table 7-19. These two bits are used by the ADDR FILTER module to determine if the desired output for the current packet exists in this ASIC.

Table 7-19. Definition Of M1 and M0

M1 M0	SWITCH FABRIC DEFINITION OF ASIC OUTPUTS 0 TO 7
0 0	FABRIC OUTPUTS 0 TO 7
0 1	FABRIC OUTPUTS 8 TO 15
1 0	FABRIC OUTPUTS 16 TO 23
1 1	FABRIC OUTPUTS 24 TO 31

Address Filter. The address filter is responsible for determining if the destination Output Port is associated with this ASIC, and, if so, which ASIC output is associated

with this port. A block diagram of this circuit is shown in Figure 7-59. A timing diagram is shown in Figure 7-60.

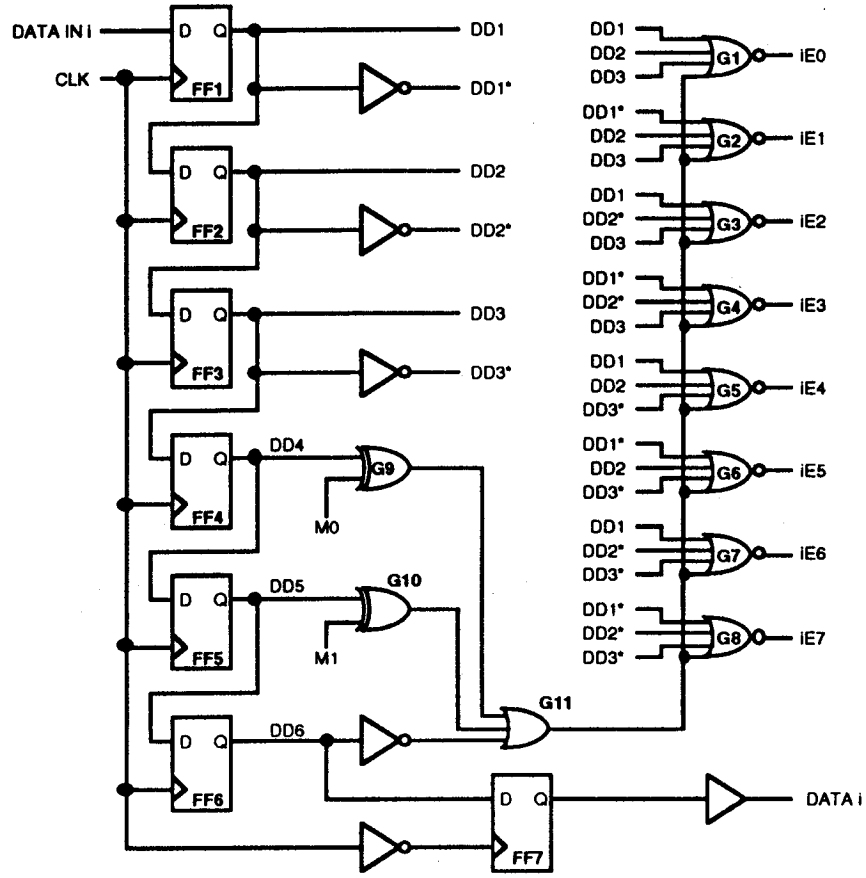

Figure 7-59. Address Filter Block Diagram

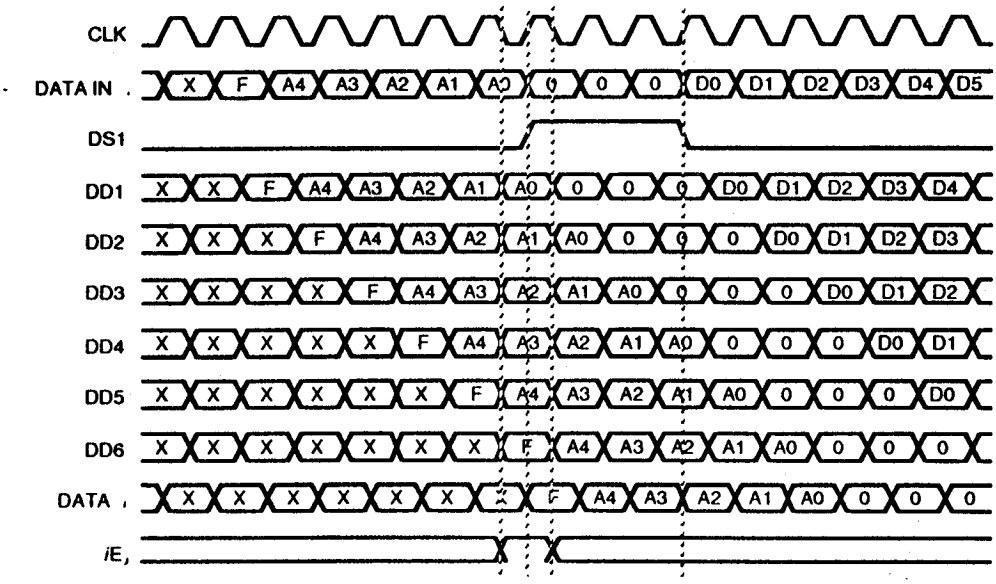

Figure 7-60. Address Filter Timing Diagram

Flip-flops FF1 to FF6 form a six bit shift register for the serial data in. When the strobe DS1 is active, FF6 contains the flag bit, FF5 (MSB) to FF1 (LSB) contain the routing tag. In order for this packet to have a valid output in this ASIC, the output of gate G11 must be low which will occur when the following three conditions are met:

- The flag bit (FF6) must be high.
- Mode bits M1 and M0 specify which group of eight addresses are valid for this ASIC per Table 7-19. These two bits must compare to the most significant two bits of the routing tag for this ASIC to respond to the packet. The comparison of the mode bit M1 to the most significant routing tag bit (FF5) is performed by the XOR gate G10. When low, the output of G10 signifies a match.
- The comparison of the mode bit M0 to the second most significant routing tag bit (FF4) is performed by the XOR gate G9. When low, the output of G9 signifies a match.

When the above three conditions are met, the output of OR gate G11 will go low, enabling the decoder NOR gates G1 to G8. Gates G1 to G8 form a 3-to-8 decoder for the outputs of flip-flops FF1 to FF3. The output of G1 will be high if the bits FF3-FF1 are "000", G2 if the bits are "001", G3 if the bits are "010", etc. Since this is a point-to-point switch, one and only one decoder output, if any, can be active at any given time.

Flip-flop FF7 reclocks and buffers the data for distribution to the Output Selection circuitry of all eight outputs of the ASIC.

Output Selection. One Output Selection circuitry is responsible for accepting the requests from all the address filters for a particular output; therefore, eight Output Selection modules exists in the ASIC. A block diagram and timing diagram are shown in Figures 7-61 and 7-62.

The 32-to-5 encoder will convert the 32 inputs, one from each address filter, to a five bit code. A simple algorithm was selected assuming that, at most, one input will be active at any time (proper operation is not guaranteed if more than one input is active). This condition should always exist since the Output Congestion Control circuitry in the Input Port prevents more than one source for any given destination Output Port. The five bit code output from the encoder will be unique for every selected input. If no input is selected, the NONE output of the encoder will be high.

The inputs to the encoder are valid a half clock before the rising edge of DS1 as shown in the timing diagram. When DS1 strobes the output of the encoder into the five bit latch formed by flip-flops FF1 to FF5, the latch will contain the code associated with the one valid input to the encoder. If there are no valid inputs to the encoder, the output of FF6 will be high, disabling the output of the multiplexor.

The output of the five bit latch is fed to the select inputs of the multiplexor to route the desired serial data to the output flip-flop FF7. If the multiplexor is disabled, the input to FF7 will be held low until the next occurrence of DS1.

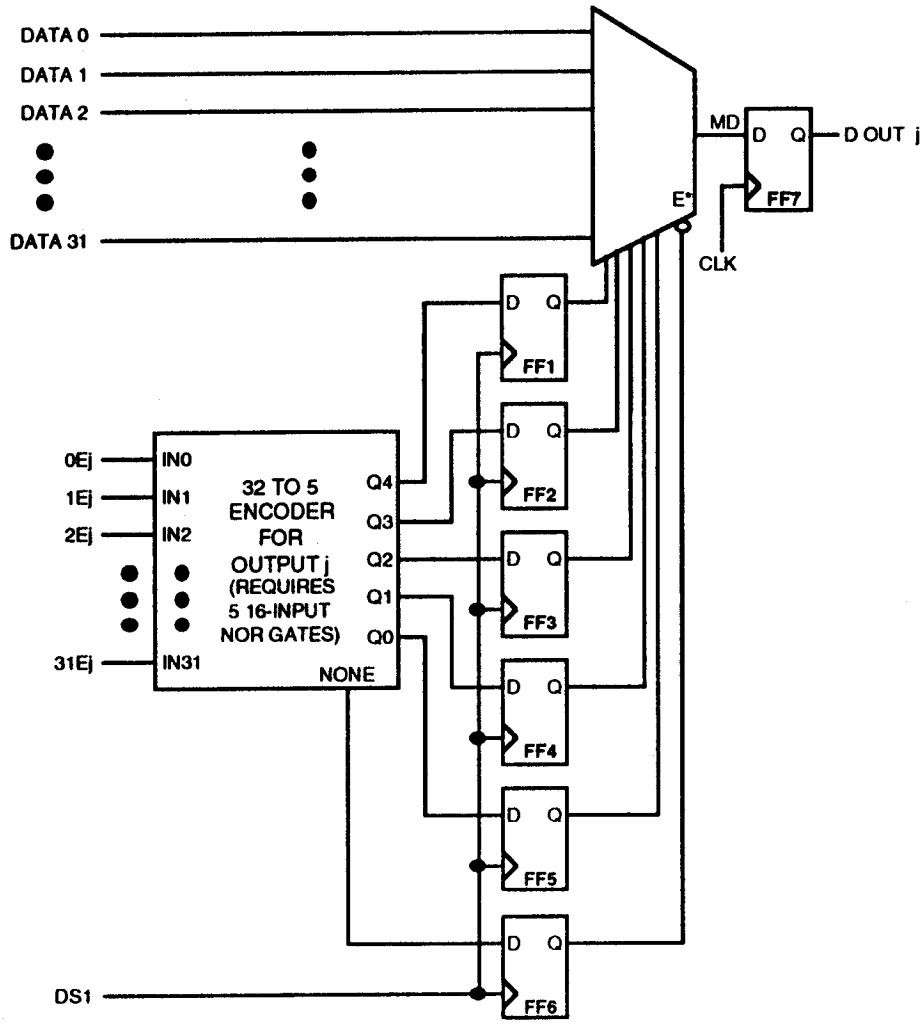

Figure 7-61. Output Selection Block Diagram

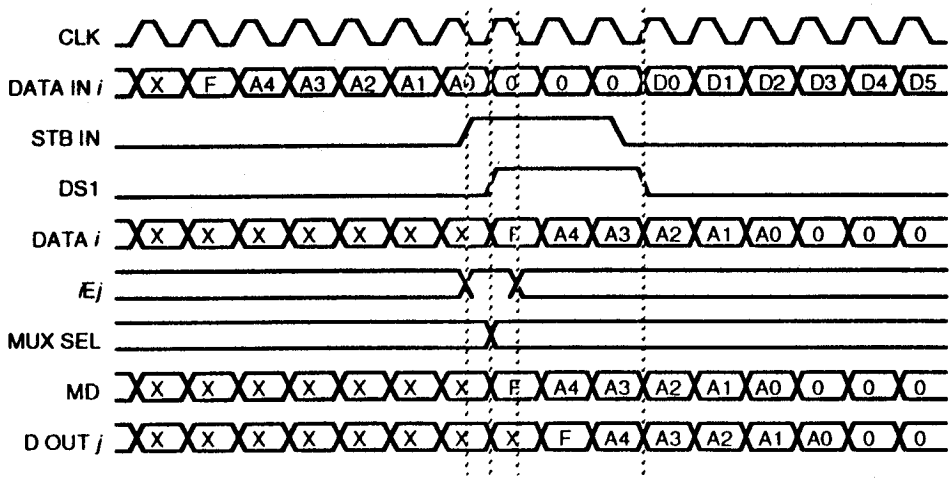

Figure 7-62. Output Selection Timing Diagram

Common Control/Timing. The input clock, CLK IN, is buffered and provided to each module in true and inverted form. Figure 7-63 is a block diagram of the circuitry required to generate the strobe signal DS1. The timing diagram is shown in Figure 7-64.

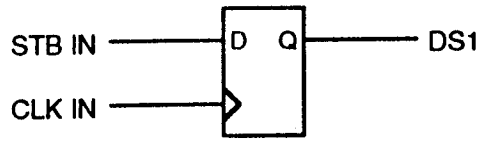

Figure 7-63. Strobe Generation Block Diagram

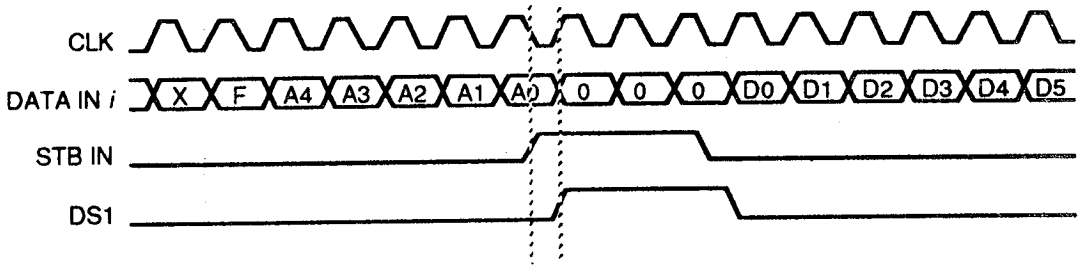

Figure 7-64. Strobe Generation Timing Diagram

The generation of DS1 is accomplished by simply mid-bit sampling the input strobe, STB IN.

7.1.2.4.3 Switch Fabric Characteristics

Crossbar ASIC Gate Count. The estimated gate count for the Crossbar ASIC is shown in Table 7-20.

Table 7-20. Crossbar ASIC Estimated Gate Count

FUNCTION	GATES/FNCTN	QUANTITY	TOTAL GATES
D FLIP-FLOP	12	300	3600
LOGIC GATE	7	850	5950
TOTAL + 20%			11,500

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Crossbar ASIC Power Consumption. The estimated power consumption for the Crossbar ASIC is shown in Table 7-21.

Table 7-21. Crossbar ASIC Estimated Power Consumption

FUNCTION	POWRE/DEVICE	QUANTITY	TOTAL POWER
INPUT I/O	6.4	36	0.230 W
OUTPUT I/O	7.81	8	0.062 W
D FLIP-FLOP	2.77	300	0.831 W
LOGIC GATE	1.4	850	1.19 W
TOTAL + 20%			2.4 W

Switch Fabric Power Consumption. The power consumption of the Switch Fabric is summarized in Table 7-22 for the three configurations. The 8 x 8, 16 x 16 and 32 x 32 configurations will require one, two, and four ASICs, respectively. Miscellaneous support circuitry is estimated at about 0.5 watts per port.

Table 7-22. Switch Fabric Power Consumption

	8 x 8	16 x 16	32 x 32
CROSSBAR MATRIX	2.4 W	4.8 W	9.6 W
MISC SUPPORT	4 W	8 W	16 W
TOTAL	6.4 W	12.8 W	25.6 W

7.1.2.5 Crossbar-Based Switch Summary

7.1.2.5.1 Power Consumption

The power consumption of the Self-Routing Crossbar-Based Switch is summarized in Table 7-23 for the three switch configurations.

Table 7-23. Crossbar-Based Switch Power Consumption

	8 x 8	16 x 16	32 x 32
INPUT PORT	38.4 W	76.8 W	153.6 W
OUTPUT PORT	27.2 W	54.4 W	108.8 W
SWITCH FABRIC	6.4 W	12.8 W	25.6 W
TOTAL	72 W	144 W	288 W
TOTAL/PORT	9 W	9 W	9 W

7.1.2.5.2 Mass/Size

A circuit board of 30 square inches will support one Input Port and one Output Port. One board will be required for every port that the switch must service. All Switch Fabric configurations can also fit on one board of this size. In addition, one board will

be required for the generation of common control and timing signals. This includes the following functions:

- Generation of Switch Fabric clocks.
- Generation and control of the Congestion Control Message.
- Generation of strobes to the Input Ports, Output Ports, and the Switch Fabric.

Therefore, the total number of modules required is 10, 18, and 34 for the 8 x 8, 16 x 16, and 32 x 32 switch configurations, respectively.

7.1.2.5.3 ASIC Design Complexity

There should be no problem in designing the two ASICs required for this approach. All ASICs have fairly low gate count and tolerable power consumption.

The power consumption of the Crossbar ASIC can be reduced by designing a 16 input and 8 output ASIC configuration. The gate count will be reduced to about 7,500 with a corresponding drop in power consumption to about 1.7 watts. However, the power and mass requirements will increase slightly for the 32 x 32 switch configuration since four additional ASICs are required for the Switch Fabric. The power consumption for this approach is shown in Table 7-24.

Table 7-24. Alternate Design Crossbar-Based Switch Power Consumption

	8 x 8	16 x 16	32 x 32
INPUT PORT	38.4 W	76.8 W	153.6 W
OUTPUT PORT	27.2 W	54.4 W	108.8 W
SWITCH FABRIC	5.7 W	11.4 W	29.6 W
TOTAL	71.3 W	142.6 W	292 W
TOTAL/PORT	8.9 W	8.9 W	9.1 W

One more possibility is to used the Multicast Self-Routing Crossbar-Based Switch Fabric. This approach will be explained in Section 7.2 for the multicast networks. This approach will destroy the routing tag as it passes through the Switch Fabric. The power consumption for this approach is shown in Table 7-25. The increased power for the Input Port is required to handle the multicast format for the routing tag. As can be seen, the power consumption per port is slightly less than 9 watts for all three configurations.

Table 7-25. Alternate Design #2 Crossbar-Based Switch Power Consumption

	8 x 8	16 x 16	32 x 32
INPUT PORT	39.2 W	78.4 W	156.8 W
OUTPUT PORT	27.2 W	54.4 W	108.8 W
SWITCH FABRIC	5.3 W	10.6 W	21.2 W
TOTAL	71.7 W	143.4 W	286.8 W
TOTAL/PORT	9 W	9 W	9 W

The design of these ASICs did not include added circuitry which may be required to implement a viable redundancy scheme. However, the addition of this circuitry is not expected to significantly increase the complexity of the ASIC design.

7.1.2.5.4 Fault Tolerance

The Crossbar-Based Network does not degrade as gracefully as the Sorted Banyan Network. Some of the shortcomings of this architectures are:

- The failure of one Input Port can affect other Input Ports. If the failed Input Port continuously sends the routing tag for the same Output Port, all other Input Ports will be prevented from accessing that Output Port.
- Only a few ASICs comprise the entire Switch Fabric. The complete loss of one ASIC will significantly decrease the operation of the Switch Fabric.

Some of the advantages of this architecture are:

- It is possible to implement a 1-for-N redundancy scheme for the Input Ports and Output Ports.
- Every input in the Crossbar ASIC has a dedicated data path to every output, therefore, a failure in one data path should have no affect on any other data path.
- A 1-for-N redundancy scheme can be implemented for the Switch Fabric ASICs. However, for and 8 x 8 Switch Fabric, N is equal to one. There are so few ASICs required in the Switch Fabric that a 1-for-1 redundancy scheme is also feasible.

The Congestion Control Message path is a ring type topology. If any Input Port fails, the entire ring becomes ineffective, causing the switch to completely shut down. For this reason, it is imperative that this ring be redundant. Also, a detected failure in a given Input Port should cause the ring input to pass through unaffected to the ring output.

7.1.3 Point-To-Point Switch Networks Summary

Table 7-26 below summarizes the various point-to-point switch approaches.

Table 7-26. Summary Of Point-To-Point Architectures

		SORTED BANYAN	CROSSBAR	CROSSBAR ALT 1	CROSSBAR ALT 2
Power Consumption (Total Switch / Per Port, in watts)	8 x 8	72/9	72/9	71/8.9	72/9
	16 x 16	148/9.2	144/9	143/8.9	143/9
	32 x 32	298/9.3	288/9	292/9.1	287/9
Switch Fabric ASIC Count	8 x 8	3	1	1	1
	16 x 16	12	2	2	2
	32 x 32	28	4	4	8
ASIC Designs	Number Required	3	2	2	2
	Gate Count (in K)	2.9/2.4/2.4	2.9/11.5	2.9/7.5	4.3/4.7
	Power (in watts)	1.8/0.7/0.7	1.8/2.4	1.8/1.7	1.9/1.3
Fault Tolerance	Graceful Degradation	Good	Moderate	Moderate	Moderate
	Redundancy	Difficult	Moderate	Moderate	Moderate

The Self-Routing Crossbar-Based Network, alternate design 1, is chosen as the optimal architecture for the current requirements. It has the lowest power consumption per port and the lowest ASIC count for the 8 x 8 and 16 x 16 Switch Networks. The power consumption and ASIC count for the 32 x 32 Switch Network is only marginally larger than the best approach for this size. There are only two ASICs to be designed, both of relatively low complexity. The overall fault tolerance of this approach is better than the Sorted Banyan Network.

The major disadvantages of the Sorted Banyan Network is the high ASIC count for larger switches, the requirement for three ASIC designs instead of two, and the complexity of implementing a 1-for-N redundancy scheme for the Switch Fabric. The Sorted Banyan Network does, however, have the advantage that it may be configured for switches larger than 32 x 32 without modification to the ASIC design.

7.2 Multicast Switch Networks

7.2.1 Self-Routing Multicast Banyan Network

A block diagram of the Self-Routing Multicast Banyan Network is shown in Figure 7-65 below.

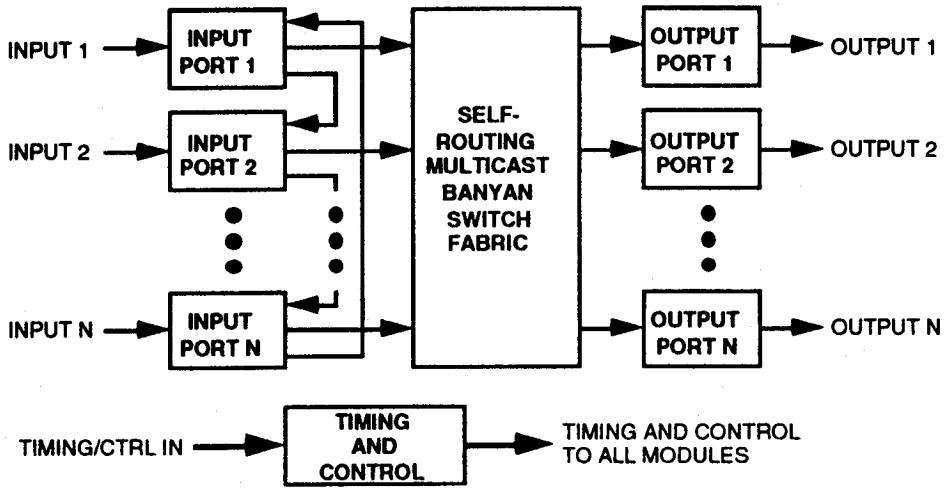

Figure 7-65. Self-Routing Multicast Banyan Network Block Diagram

Serial data from the demodulators is routed through the switch to the modulators in the following steps:

- The serial packet data enters the respective Input Port where it is stored in a 128 packet FIFO type buffer.
- As each packet is stored in the packet buffer, the routing tag associated with the packet is obtain from the routing map memory then stored with the packet in the buffer.
- The Input Port performs Output Port congestion control to ensure that there are no blocked packets at the Output Port. The Input Ports convey their desired destination(s) to each other by means of a serial daisy chained message every packet slot time. When an Input Port reserves an Output Port(s) for the next packet transmission, no other Input Port further down the daisy chain may select the reserved Output Port(s). The four oldest in the packet buffer will be examined for transmission through the Switch Fabric to minimize head-of-line blocking.
- The selected packet, if any, from each Input Port is sent to the Switch Fabric. The Switch Fabric will route packets from all Input Ports to the desired Output Ports as specified by the routing tag. The Switch Fabric will need to operate at a higher rate than the input serial data to compensate for the

added bits for the routing tag and for the lost packet transfers through the Switch Fabric due to head-of-line blocking.

- The Output Port will accept serial data from the Switch Fabric then store it in a 128 packet FIFO type buffer. The serial packet data will be output to the modulator for downlink transmission.

7.2.1.1 Formats

7.2.1.1.1 Routing Tag Format

The format of the routing tag appended to each packet is shown in Figure 7-66 below. When an 8 x 8 Switch Fabric is used, only the first 15 bits are required. When a 16 x 16 Switch Fabric is used, the first 31 bits are required. All 63 bits are required for the 32 x 32 Switch Fabric.

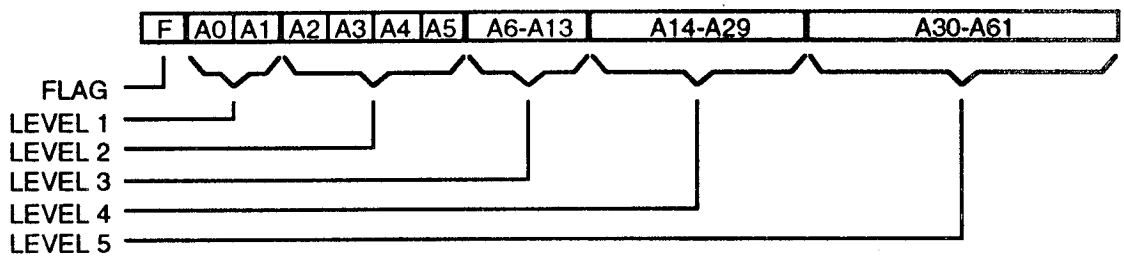

Figure 7-66. Multicast Banyan Routing Tag Format

Each of the bits above are used in pairs. A0 and A1 form a pair for the level 1 routing information. A2 and A3 form one pair and A4 and A5 form the second pair for the level 2 routing information. The definition of these bit pairs are given in Table 7-27.

Table 7-27. Routing Bit Pair Definition

EVEN BIT	ODD BIT	CONFIGURATION
0	0	Both Outputs
0	1	Even Output
1	0	Neither Output
1	1	Odd Output

The definition of the bits in the routing tag are given in Table 7-28.

Flag Bit. The flag bit must be high for the routing tag to be valid. If low, the packet is not a valid packet and will not be routed to any output.

Level 1 Routing Bits. The level 1 routing bits, A0 and A1, are used by the first level Batcher sorting network and the the first level Banyan switching cells.

Table 7-28. Routing Tag Bit Definitions

BIT	DEFINITION
F	Flag bit
A0, A1	Level 1 Routing Bits
A2, A3	Level 2 Routing Bits, Even Path
A4, A5	Level 2 Routing Bits, Odd Path
A6, A7	Level 3 Routing Bits, Even-Even Path
A8, A9	Level 3 Routing Bits, Even-Odd Path
A10, A11	Level 3 Routing Bits, Odd-Even Path
A12, A13	Level 3 Routing Bits, Odd-Odd Path
A14, A15	Level 4 Routing Bits, Even-Even-Even Path
A16, A17	Level 4 Routing Bits, Even-Even-Odd Path
A18, A19	Level 4 Routing Bits, Even-Odd-Even Path
A20, A21	Level 4 Routing Bits, Even-Odd-Odd Path
A22, A23	Level 4 Routing Bits, Odd-Even-Even Path
A24, A25	Level 4 Routing Bits, Odd-Even-Odd Path
A26, A27	Level 4 Routing Bits, Odd-Even-Even Path
A28, A29	Level 4 Routing Bits, Odd-Odd-Odd Path
A30, A31	Level 5 Routing Bits, Even-Even-Even-Even Path
A32, A33	Level 5 Routing Bits, Even-Even-Even-Odd Path
A34, A35	Level 5 Routing Bits, Even-Even-Odd-Even Path
A36, A37	Level 5 Routing Bits, Even-Even-Odd-Odd Path
A38, A39	Level 5 Routing Bits, Even-Odd-Even-Even Path
A40, A41	Level 5 Routing Bits, Even-Odd-Even-Odd Path
A42, A43	Level 5 Routing Bits, Even-Odd-Odd-Even Path
A44, A45	Level 5 Routing Bits, Even-Odd-Odd-Odd Path
A46, A47	Level 5 Routing Bits, Odd-Even-Even-Even Path
A48, A49	Level 5 Routing Bits, Odd-Even-Even-Odd Path
A50, A51	Level 5 Routing Bits, Odd-Even-Odd-Even Path
A52, A53	Level 5 Routing Bits, Odd-Even-Odd-Odd Path
A54, A55	Level 5 Routing Bits, Odd-Odd-Even-Even Path
A56, A57	Level 5 Routing Bits, Odd-Odd-Even-Odd Path
A58, A59	Level 5 Routing Bits, Odd-Odd-Odd-Even Path
A60, A61	Level 5 Routing Bits, Odd-Odd-Odd-Odd Path

Level 2 Routing Bits. Level 2 routing bits are used by the second level Batcher sorting networks and the second level Banyan switching cells. The bit pairs to be used will be determined by the path to the current cell as specified by Table 7-28. If the packet comes from the even output of the previous Banyan cell, then A2 and A3 are used; if from an odd output, A4 and A5 are used.

Level 3 Routing Bits. Level 3 routing bits are used by the third level Batcher sorting networks and the third level Banyan switching cells. The bit pairs to be used will be determined by the path to the current cell as specified by Table 7-28. For example, if the packet came from an even output of a level 1 cell and an even output of a level 2 cell, then A6 and A7 are used.

Level 4 Routing Bits. Level 4 routing bits are used by the fourth level Batcher sorting networks and the fourth level Banyan switching cells. The bit pairs to be used will be determined by the path to the current cell as specified by Table 7-28. For example, if the packet came from an even output of a level 1 cell, an even output of a level 2 cell, and an even output of a level 3 cell, then A14 and A15 are used.

Level 5 Routing Bits. Level 5 routing bits are used by the fifth level Batcher sorting networks and the fifth level Banyan switching cells. The bit pairs to be used will be determined by the path to the current cell as specified by Table 7-28. For example, if the packet came from an even output of a level 1 cell, an even output of a level 2 cell, an even output of a level 3 cell, and an even output of a level 4 cell, then A30 and A31 are used.

7.2.1.1.2 Output Port Congestion Control Message Format

The format for the Output Port Congestion Control Message is identical to that used for the Sorted Banyan Network described in Section 7.1.1.1.2. However, since this is a multicast switch, more than one Output Port may be reserved by the Input Port.

7.2.1.2 Input Port

The Input Port is very similar to the Input Port of the Sorted Banyan Network. The differences are listed below.

- The stored routing information in the data RAM is 32 bits, one bit for every Output Port. This affects the number of transfers required to and from the RAM.
- The Output Port Congestion Control circuitry must accommodate multiple Output Port destinations for a given packet (multicast packet).
- The Switch Fabric Side Processing must translate the 32 bit routing data stored in RAM to the 62 bit routing tag for transmission to the Switch Fabric.

A block diagram of the Input Port is shown in Figure 7-67.

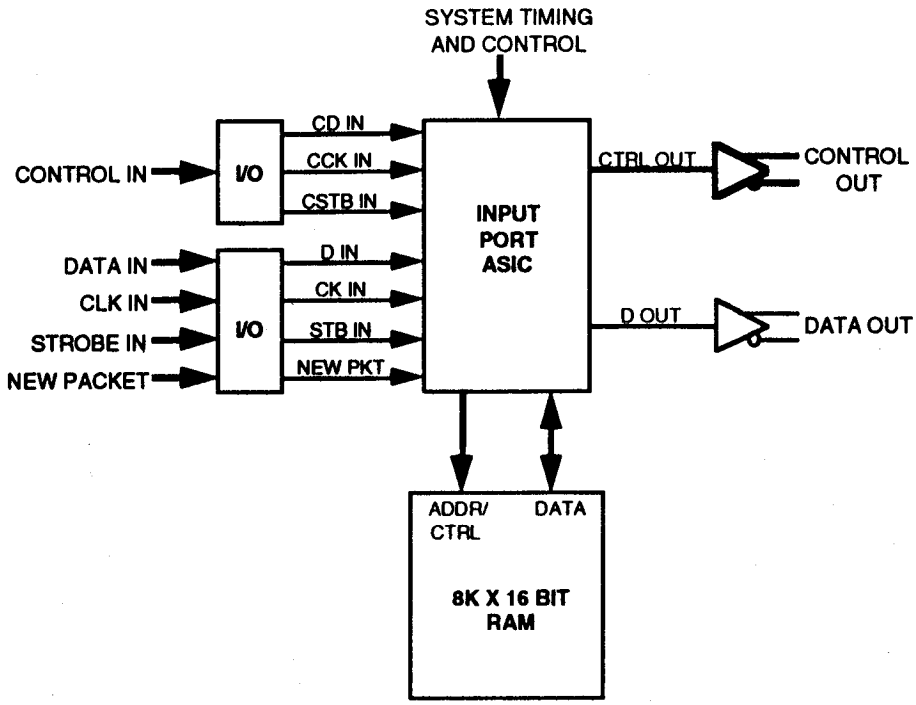

Figure 7-67. Multicast Banyan Input Port Block Diagram

The Input Port performs the following major tasks:

- Serial data enters the Input Port via the DATA IN signal. This data is aligned to the common control signals CLK IN and NEW PACKET. Transitions of the serial data occur on the rising edge of CLK IN. NEW PACKET determines the start of the received packet.

  The differentially received serial data and control lines are converted to single-ended ECL by the I/O modules.

- The Input Port ASIC will convert the serial data to parallel using the control signals.

- The parallel input data is then stored in the first half of the RAM which is configured as a 4K by 16 bit FIFO type buffer. The buffer is configured as 128 pages of 32 words (64 bytes). The 53 bytes of packet data is stored sequentially starting at word 0 of the next available page.

- As an input packet is being stored in the buffer, its routing tag is being obtained. The Input Port ASIC will examine up to sixteen bits of the destination address in the packet header. These sixteen bits will be used to obtain the routing tag by the mapping shown in Figure 7-68 below.

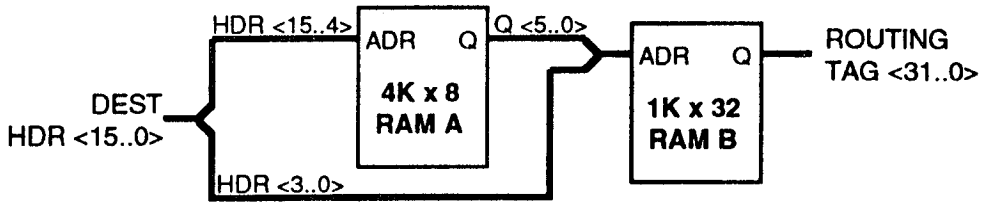

Figure 7-68. Routing Tag Memory Mapping

RAM A is actually implemented as a 2K x 16 RAM where the least significant bit of the address is used to determine if the high or low byte of the memory access is used. RAM B is also implemented as a 2K x 16 RAM where two word accesses are required to access the routing tag field. RAM A and RAM B occupy the lower half of the RAM memory.

Both RAM A and RAM B have been intialized by the signal processor. The most significant twelve bits of the destination address will be condensed by the mapping in RAM A into six bits. These six bits and the least significant four bits of the destination address will be used to fetch the routing tag from RAM B. This method provides a mapping of 64K addresses (2^{16}) to a minimum of 64 routing tags with only 16K bytes of memory.

- The thirty-two bit routing tag is stored in word 30 and 31 of the buffer page associated with the packet. The packet data occupy the first 26 1/2 words of the page with words 26 1/2 to 29 (7 bytes) not used.

- Output Port congestion control is performed. The congestion control information is contained on the CONTROL IN signals. These differential signals are converted to single-ended ECL by the I/O module. The message is received serially on the CD IN signal with the associated clock, CCK IN. The CSTB IN provides a strobe timing signal to delineate the fields in the message stream.

  The congestion control circuitry is initialized when a RESET mode is detected in the message. If the Input Port port number is greater than or equal to the filter address and the mode is not NOP, the congestion control circuity is enabled. The congestion control circuitry will check the token in the message associated with the desired Output Port(s). If there is at least one token for a desired Output Port not set, the congestion control circuitry will select this packet for transmission. All Output Ports which are available (tokens not previously set) will be selected and their tokens set. Output Ports which were desired but blocked because their token was already set are saved for retransmission in a subsequent packet slot period. If none of the above conditions are met, the congestion control message is passed through without modification.

  The congestion control message will be sent five times during each packet slot. Every Input Port will have an opportunity to check all four packets at the

head-of-line in the packet buffer. Only one packet is checked in an Input Port for each message. In this manner, the oldest packet in the buffer of all Input Ports are checked before subsequent packets in any buffer are checked.

- Once the packet to be transmitted to the Switch Fabric is determined, the packet, including the routing tag, must be fetched from the packet buffer. This fetch starts as the current packet transmission is ending. The STROBE IN signal identifies the packet slot boundaries within the Switch Fabric. The packets will be transmitted at a higher rate than the port data rate to compensate for the added bits in the routing tag (about 8%) and for the blocked packets due to head-of-line blocking (about 20%). Therefore, the minimum Switch Fabric rate must be 200 MHz (155 * 108% * 120%).

- The fetched 16 bit words from the packet buffer are converted to serial data to be sent to the Switch Fabric.

The Input Port ASIC is responsible for performing all the functions above. The four main functions of the ASIC are:

- Port Side Processing. This includes the input serial data conversion to parallel, generating the RAM packet buffer address to store the parallel word, and generating the RAM address to obtain the routing tag.

- Output Port Congestion Control. This includes all functions associated with the Output Port Congestion Control.

- Switch Fabric Side Processing. This includes storing of the start locations of the four packets at the head-of-line, generating the routing tag address for use in the Output Port Congestion Control, generating the address for the packet data to be transmitted to the Switch Fabric, and converting the parallel data to serial format.

- RAM Processing. This includes all address/data buffering and multiplexing and the generation of the control signals to the RAM memory.

7.2.1.2.1 Port Side Processing

A block diagram of the Port Side processing functions is shown in Figure 7-69.

Serial to Parallel Conversion. The "16 Bit Serial To Parallel SR" (S/P SR), the "4 Bit CTR" (BIT CTR), and the "16 Bit Data In Latch" (DATA LATCH) perform the serial to parallel conversion function.

- BIT CTR represents the bit count within a sixteen bit word boundary of the incoming serial data stream. BIT CTR is cleared on NEW PKT and is clocked on the data clock signal, CK IN. The terminal count output of the BIT CTR, TC, will occur sixteen bits after the start of the packet and every sixteen bits thereafter.

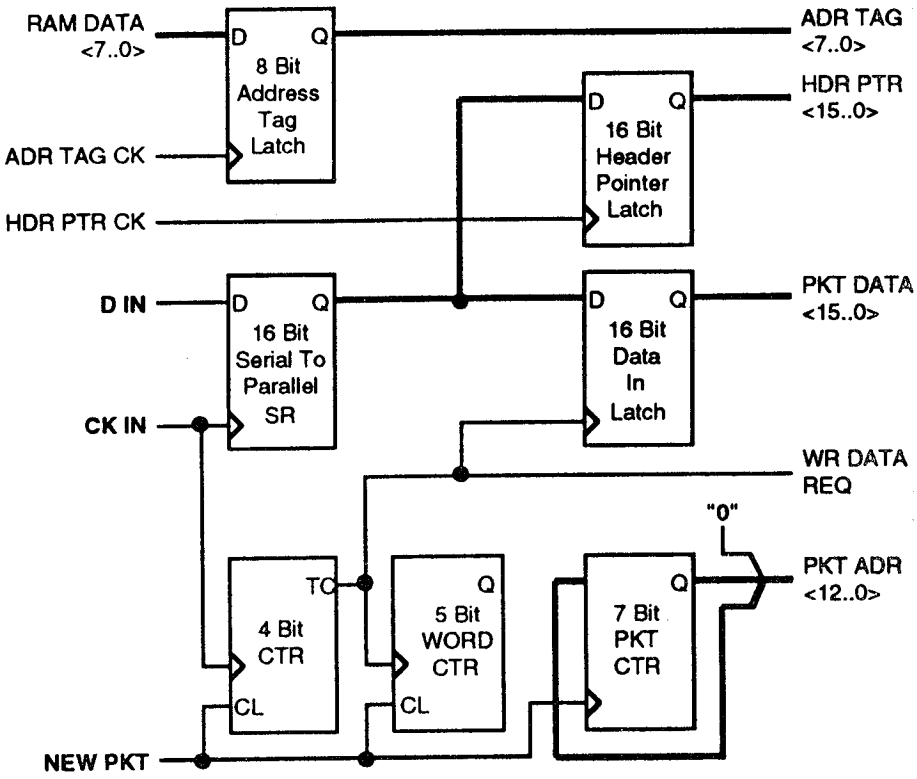

Figure 7-69. Port Side Processing Block Diagram

- Serial data, D IN, is clocked into the S/P SR by the data clock, CK IN to convert it to parallel data for storage in the RAM memory.

- The output of the S/P SR is latched into the DATA LATCH by the terminal count output of the BIT CTR at which time a request is made to the RAM Processing module to store the output of the DATA LATCH in the packet buffer. This request must be serviced before the next request is made sixteen clock periods later.

  The fifty-three bytes of the packet will be stored in the first twenty-seven locations of the page in memory associated with the packet.

Generate Buffer Address for Packet Data. The buffer address for the packet data is generated by the "4 Bit CTR" (BIT CTR), the "5 Bit WORD CTR" (WORD CTR), and the "7 Bit PKT CTR" (PKT CTR). The BIT CTR and WORD CTR are initialized at the start of every packet by NEW PKT. BIT CTR represents the bit count of the incoming serial data. The terminal count output, TC, of the BIT CTR goes active when a word of the packet has been latched in DATA LATCH. The WORD CTR therefore represents the word count of the incoming packet. The PKT CTR is incremented for every new packet by NEW PKT. The PKT CTR represents 1 of 128 packet pages within the packet buffer. The address for the parallel data is obtained by appending the 7 bit output of the PKT CTR to the 5 bit output of the WORD CTR, resulting in 12 bits (4K data buffer address

space). The most significant bit of the RAM address must be forced to "0" to access the data buffer portion of the RAM.

Generate Addresses for Obtaining Routing Tag. The RAM address for obtaining the routing tag is accomplished by the "8 Bit Address Tag Latch" (TAG LATCH) and the "16 Bit Header Pointer Latch" (HDR LATCH). The circuitry for obtaining the routing tag is able to map 16 bits (64K unique address) to a minimum of 64 routing tags. These sixteen bits, which should be two consecutive bytes within the packet, are latched into HDR LATCH.

The most significant twelve bits of HDR LATCH are used for the first access. The most significant eleven of the twelve bits are used for the least significant bits of the address; the two most significant bits of the address are set to "10" to select the third 2K block within the RAM memory. The least significant bit of the HDR LATCH determine if the low or high byte is to be used. The selected byte of the first fetch is latched into the TAG LATCH.

The second address is formed from the contents of the TAG LATCH and the least significant four bits of the HDR LATCH. The two most significant bits of the RAM address are "11" to select the last 2K block of the RAM memory. The next six bits will use the least significant six bits of the TAG LATCH. The next four bits of the address will be formed by the least significant four bits from the HDR LATCH. The least significant bit of the RAM address will determine if word 1 or word 2 of the thirty-two bit routing tag is being fetched.

The above process requires that the contents of the two specified maps be determined by the signal processor prior to the reception of the packet. The signal processor will be able to change one entry in the RAM at a time, i.e., the map is not a ping-pong type memory where the entire map can be changed offline then flipped online.

7.2.1.2.2 Output Port Congestion Control

A block diagram of the Output Port Congestion Control is shown in Figure 7-70 The Mode Circuitry and the Address Filter Circuitry are the same as for the Sorted Banyan Network with their block diagrams shown in Figures 7-8 and 7-9. The timing diagram for the Address Filter and Mode Circuitry is shown in Figure 7-10. A block diagram and timing diagram for the Output Port selection are shown in Figures 7-71 and 7-72.

The control message will experience a one clock delay through the circuitry. The serial data, CD IN, and strobe signal, CSTB IN, are clocked into flip-flop FF1 and FF2 on the negative edge of the input clock, CCK IN, to obtain the DIN 1 and STB 1 signals. Flip-flop FF3 reclocks STB 1 on the positive edge of CCK IN to compensate for the delay to the data through flip-flop FF4. Under most conditions, the or-gate, G1, will pass DIN 1 to flip-flop FF4 for transmission to the next Input Port. The only exception is when this Input Port wishes to set one of the tokens to reserve the associated Output Port.

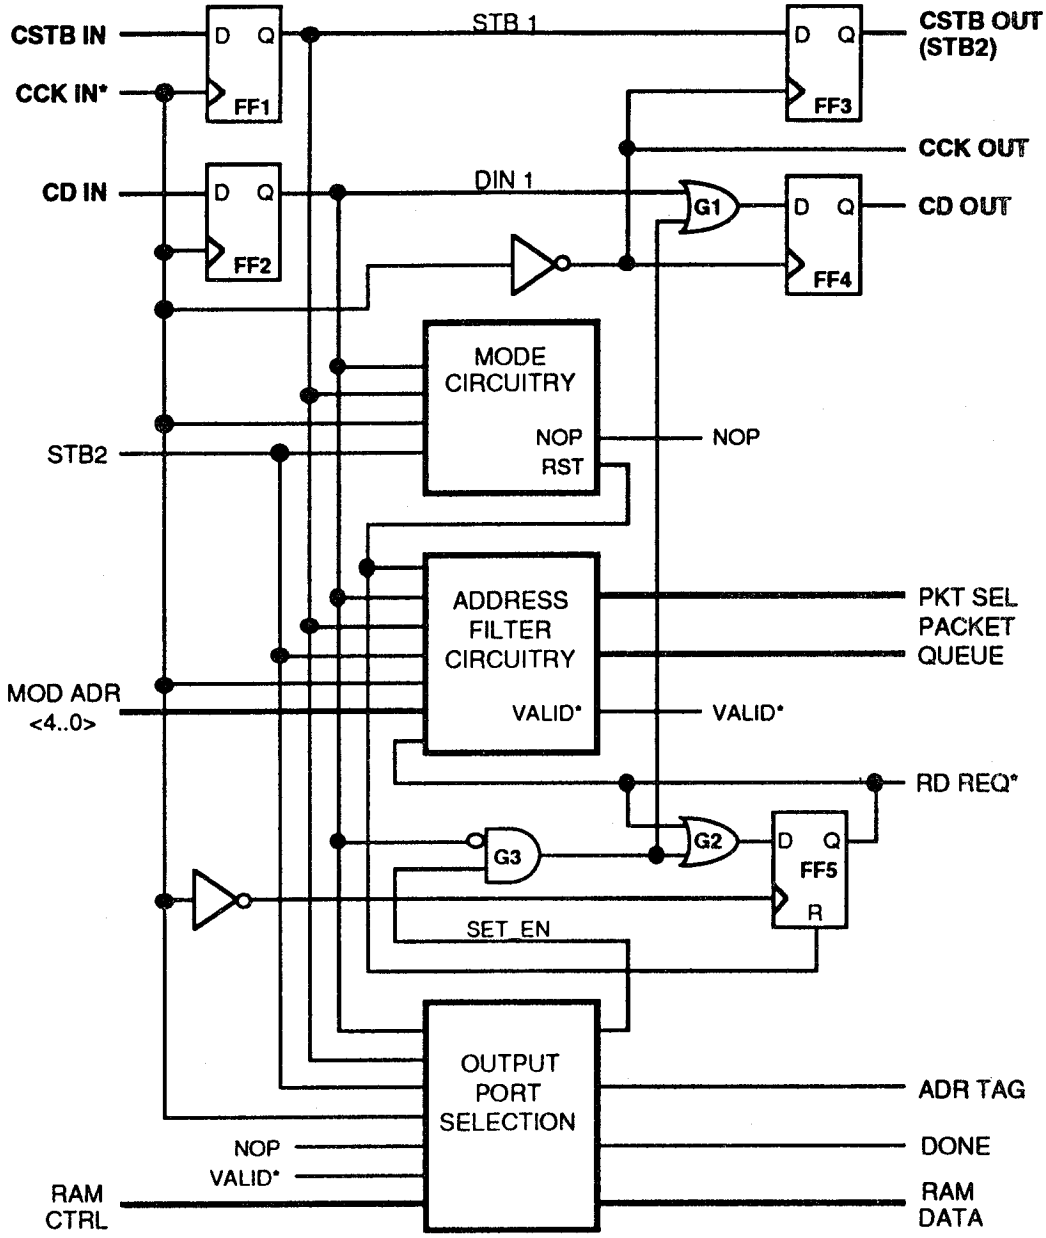

Figure 7-70. Output Port Congestion Control Block Diagram

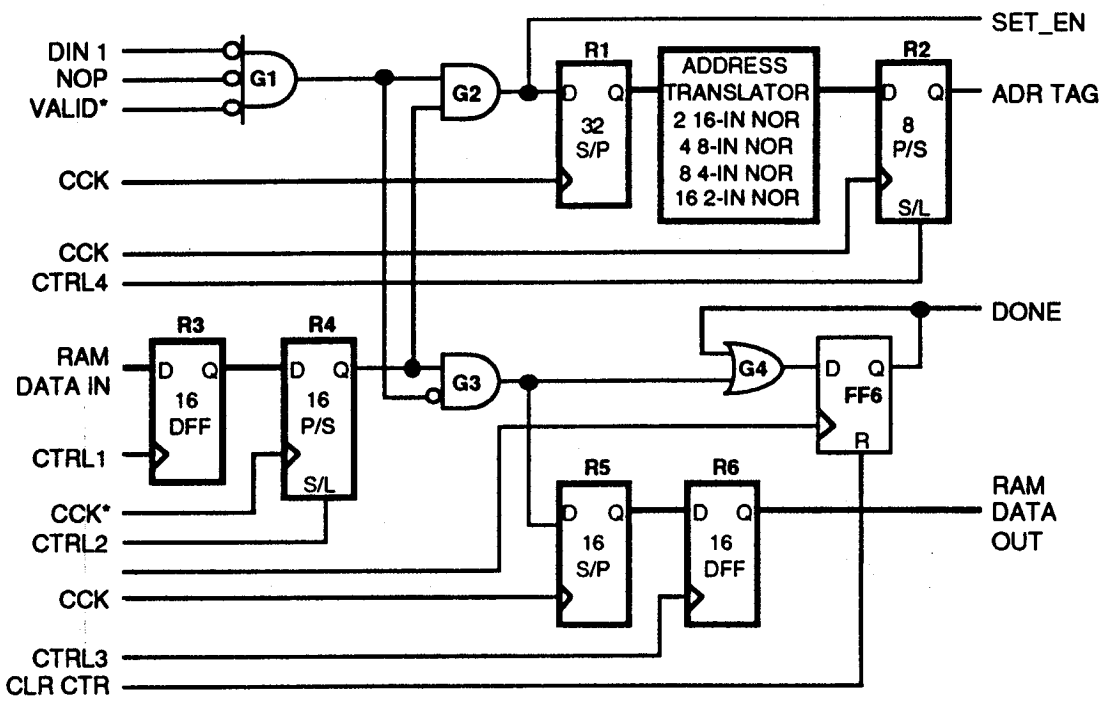

Figure 7-71. Output Port Selection Block Diagram

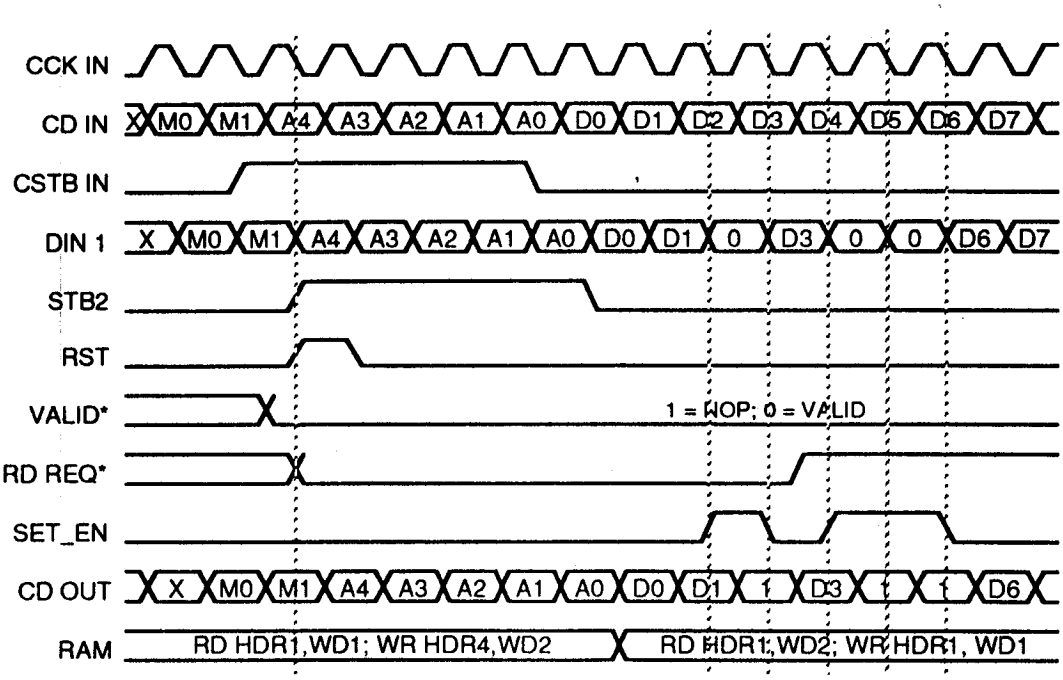

Figure 7-72a. Output Port Selection Timing Diagram (a)

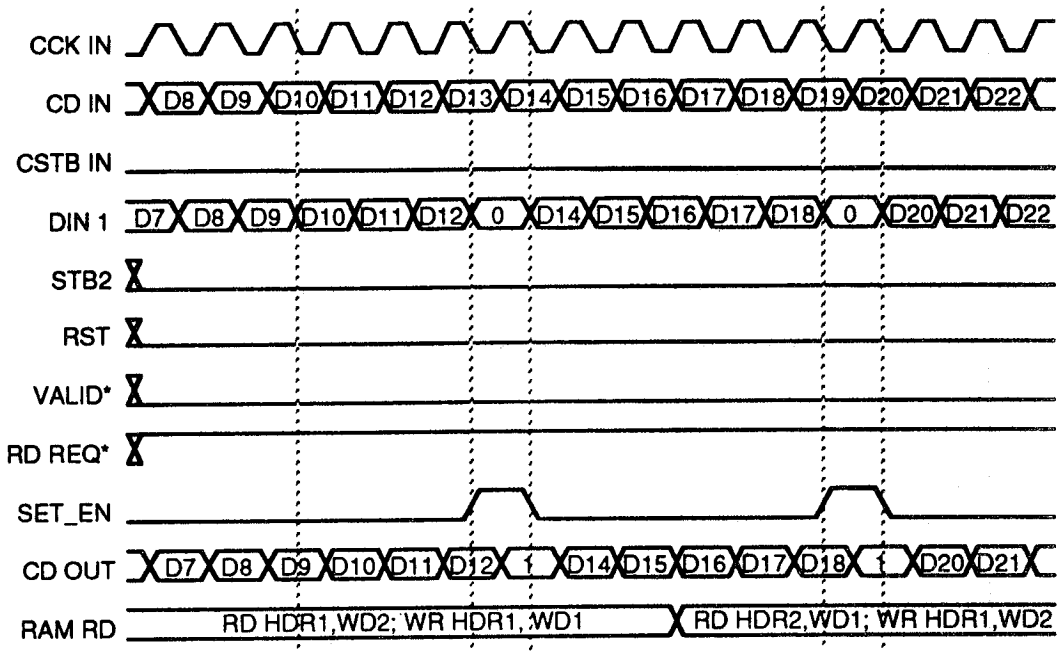

Figure 7-72b. Output Port Selection Timing Diagram (b)

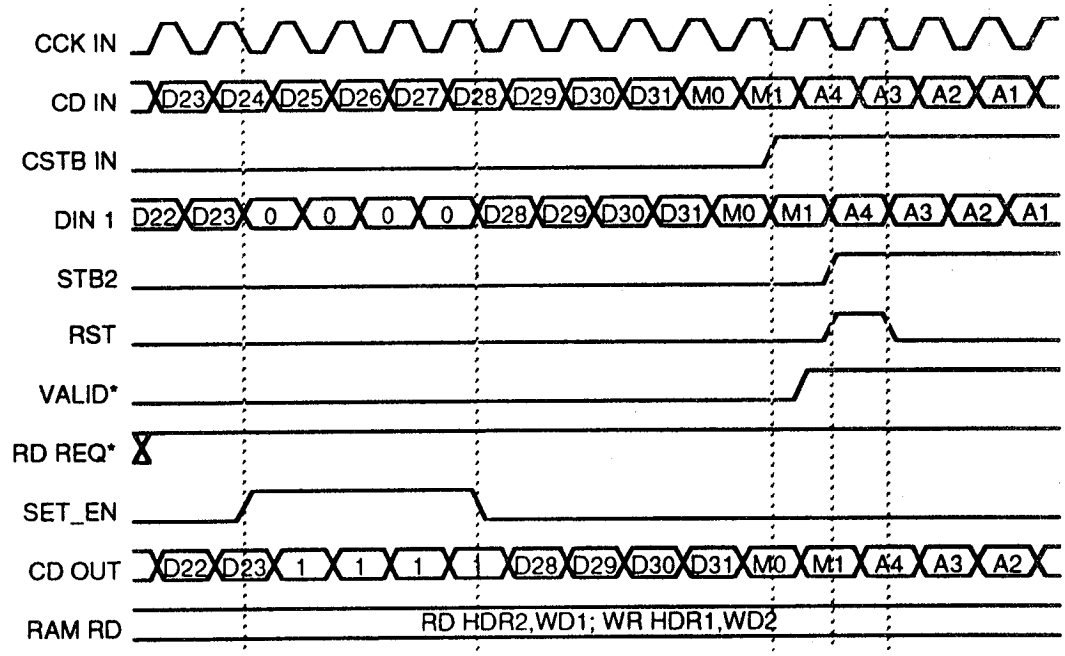

Figure 7-72c. Output Port Selection Timing Diagram (c)

The following conditions must be met for this Input Port to be allowed to set one of the Output Port token bits.

- This Input Port has not reserved an Output Port for a different packet. This information is present in flip-flop FF5. When set, it is an indication that a token has already been set in a previous message of the packet slot period. See final criteria below for a more detailed description of the operation of this flip-flop.

- The Mode Control bits at the beginning of the message must be a non-NOP code. The Mode Control block diagram is shown in Figure 7-8. The first two bits of the message are latched into flip-flops FF6 and FF7. These two bits are decoded to generate the signals for NOP, NORMAL, and RESET. The RESET signal is further waveshaped to provide the RST signal as shown in the timing diagram. See the timing diagram in Figure 7-10.

- The port number for this Input Port must be greater than or equal to the Filter Address in the control message. The Address Filter Circuitry block diagram is shown in Figure 7-9 The output of flip-flop FF8 indicates if the Input Port port number is greater than or equal to the Filter Address - a low indicates that the Input Port passed the address filter. The output of flip-flop FF9 is used to enable the 4-input and gate which feeds flip-flop FF8. The circuitry performs as follows:

  - When STB 1 goes active, the counter is cleared to zero, flip-flop FF8 is reset and flip-flop FF9 is set.

  - When STB 2 goes high, the check for the first bit is performed. The multiplexor will present port number bit A4 to the XOR gate. If this bit compares to the DIN 1 bit (A4 of the Filter Address), then no action is taken. If this bit differs from DIN 1, then the following tasks are implemented:

    1. FF9 is reset to disable all further checks

    2. If DIN 1 is high, it indicates that the Filter Address is larger than the port number so flip-flop FF8 is set and remains set until the next RST pulse.

    3. If DIN 1 is low, it indicates that the Filter Address is smaller than the port number so flip-flop FF8 remains low until the next RST pulse.

  - If no action was taken for the first bit check, then the second bit, A3, is checked on the next clock cycle. Again, if the two A3 bits compare, no action is taken; if they differ, the three actions in the above step are implemented.

- The above step is repeated until the the XOR detects a difference or the STB 2 signal goes low, indicating the end of the Filter Address field. If the STB 2 signal goes low with no differences detected, the port address is equal to the Filter Address and flip-flop FF8 should remain reset. Figure 7-10 shows an example where the port number was detected as smaller than the filter address on bit A2.

- The final criteria for the Input Port to be allowed to set a reservation token bit in the message is that the token has not been previously set by another Input Port. When set, flip-flop FF5 indicates that this Input Port has a packet to send to the Switch Fabric in the next packet slot period. This bit is reset on RST and if all conditions are satisfied, it will be set by the end of the current packet slot period. See Figure 7-71. Gate G2 is high if DIN 1, NOP, and VALID* are low and if the associated routing tag bit is high. When this occurs, this Output Port corresponding to this token bit is reserved and the token bit is set for transmission to the next Input Port. Serial-to-parallel shift register, R1, will serially shift in the output of G2. This register will be an indication of all bits set for this congestion control message. These thirty-two bits will go to a translator circuit to obtain the routing tag for the Switch Fabric. Registers R3 and R4 provide the routing tag bits for the comparison to G2. R4 loads sixteen bits of the routing tag, then shifts them out serially to G2. R3 is a latch which holds the subsequent sixteen bits to be used by R4. Gate G3 will pass the routing tag information from R4 except when G2 is active when a zero is forced on the output of G3. Register R5, which serially shifts in the output of G3, represents all routing tags which need to be serviced but were unable to reserve a token. Register R6 latches the output of R5 for buffering to the RAM memory store. Flip-flop FF6 is reset at the beginning of the message. It is set when the output of G3 is high. Therefore, if all outstanding routing tag bits reserved a token, FF6 will remain low, indicating that this packet will be completely serviced by the next packet slot transmission.

7.2.1.2.3 Switch Fabric Side Processing

A block diagram of the Switch Fabric side processing is shown in Figure 7-73.

Storing Start Packet Locations for the Four Head-Of-Line Packets. The "7 Bit PKT CTR" (PKT CTR) contains the start location for the next packet to be serviced from the data buffer. The four packet latches, "7 Bit Pkt *i* Ltch" (PKT*i* LTCH), contain the current start locations for the four packets at the head-of-line. PKT0 LTCH contains the start location for the packet at the head of the line followed by PKT1 LTCH, PKT2 LTCH, and PKT3 LTCH. A packet which is transmitted to the Switch Fabric can be associated with any one of the four packet latches. Once the packet is transmitted the contents of all packet latches down the line should move up by one and the new packet header should be latched into PKT3 LTCH. For example, if the packet that was transmitted was represented by PKT1 LTCH, then the contents of PKT2 LTCH should be moved to PKT1 LTCH, the contents of PKT3 LTCH should be moved to PKT2 LTCH,

and the PKT CTR contents should be latched into PKT3 LTCH. PKT0 LTCH remains unaffected. The four packet latches now contain the four packets at the head of the line in the same time sequence in which they entered the data buffer.

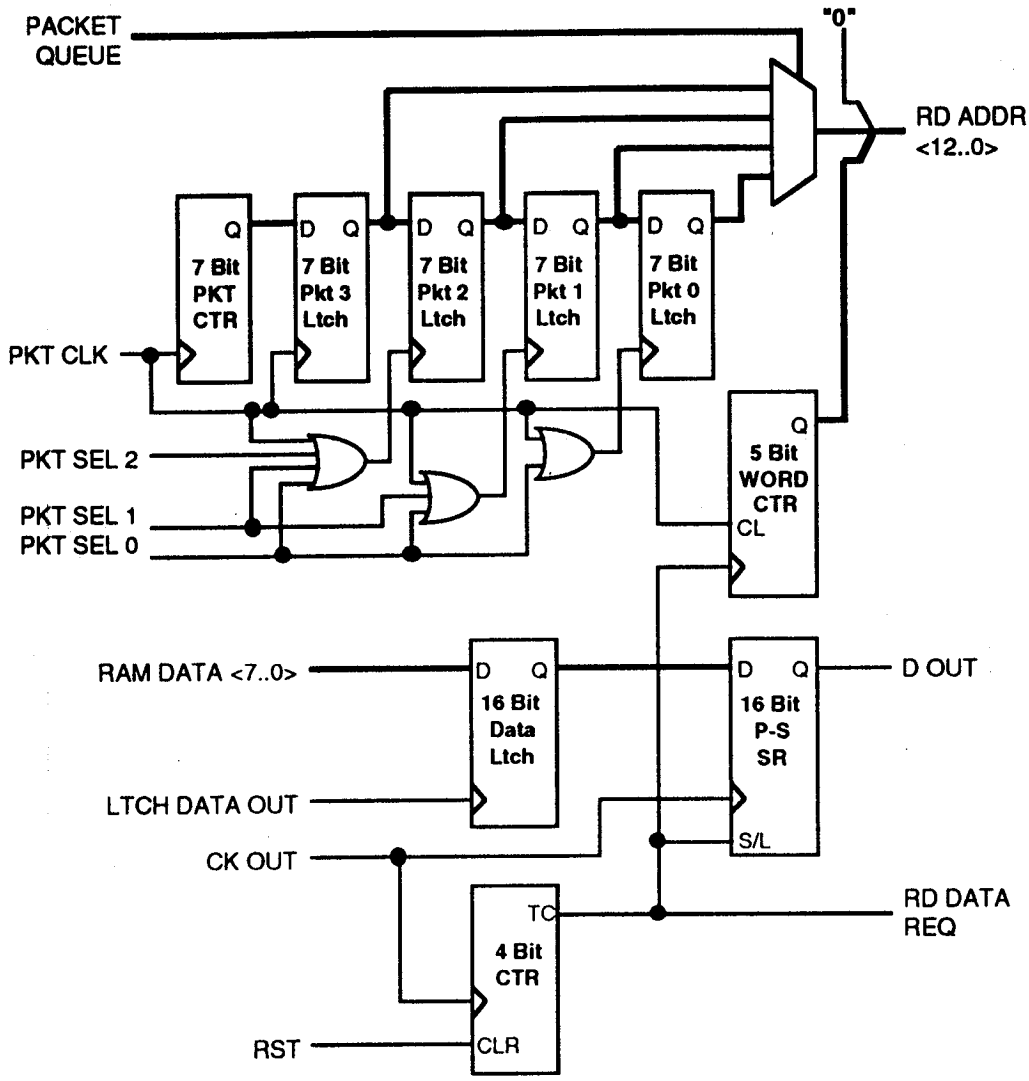

Figure 7-73. Switch Fabric Side Processing Block Diagram

The PKT CLK line increments the PKT CTR and latches the contents of the PKT CTR into PKT3 LTCH. If PKT2 LTCH is to be updated, PKT SEL 2 will be high to enable the transfer of PKT3 LTCH contents to PKT2 LTCH. Similarly, if PKT1 LTCH is to be updated, PKT SEL 1 will be high to enable the PKT2 LTCH and PKT1 LTCH for update. If PKT0 LTCH is to be updated, PKT SEL 0 will be high enabling all packet latches to be updated.

Generation of Routing Tag RAM Address. The routing tag is stored in the last two locations of the 32 word page in the data buffer. This page represents one packet and is specified by the seven bits held in the packet latches, PKT*i* LTCH. The RAM address to fetch the routing tag for the packet associated with PKT*i* LTCH is formed as follows:

- The most significant bit, A12 is forced to "0" to select the data buffer portion of the RAM memory.
- The next seven bits, A11 to A5, are the contents of the PKT*i* LTCH.
- The least significant five bits, A4 to A0, represent the word count and are forced to "11110" or "11111" to select the first or second word of the routing tag.

Generation of Packet Data RAM Address. The RAM address to fetch the packet data associated with the selected packet latch, PKT*i* LTCH, is formed as follows:

- The most significant bit, A12 is forced to "0" to select the data buffer portion of the RAM memory.
- The next seven bits, A11 to A5, are the contents of the PKT*i* LTCH.
- The least significant five bits, A4 to A0, represent the word count and are obtained from the word counter, "5 Bit WORD CTR" (WORD CTR). The WORD CTR is cleared at the beginning of every packet transmission. It is clocked from the terminal count output, TC, of the bit counter, "4 Bit CTR" (BIT CTR).

Parallel to Serial Conversion. Data fetched for output to the Switch Fabric is latched into "16 Bit Data Ltch" (DATA LATCH). The BIT CTR is reset at the beginning of the packet transmission by RST. When the terminal count of the BIT CTR is active every sixteen bit periods, the "16 Bit P-S SR" (P/S SR) will be parallel loaded. For the next fifteen clocks of the data clock, CK OUT, the data will be serially shifted out on D OUT of the P/S SR.

7.2.1.2.4 RAM Processing

A block diagram of the RAM Processing circuitry is shown in Figure 7-74.

Address Multiplexing. The multiplexor M1 selects the desired RAM address for the current access as determined by the RAM Control Logic. Sources of RAM addresses are:

- RD ADDR <12..0>. Address from the Switch Fabric Side Processing module. This address is used to fetch the routing tag for the Output Port Congestion Control and the packet data for transmission to the Switch Fabric.
- PKT ADDR <12..0>. Address from the Port Side Processing module. This address is used to store the incoming packet data and associated routing tag.
- "10" | HDR PTR <15..5>. Address from the Port Side Processing module. This address is the indirect address for fetching the routing tag. The most significant two bits are "10" to select the routing tag indirect address map portion of the RAM memory.

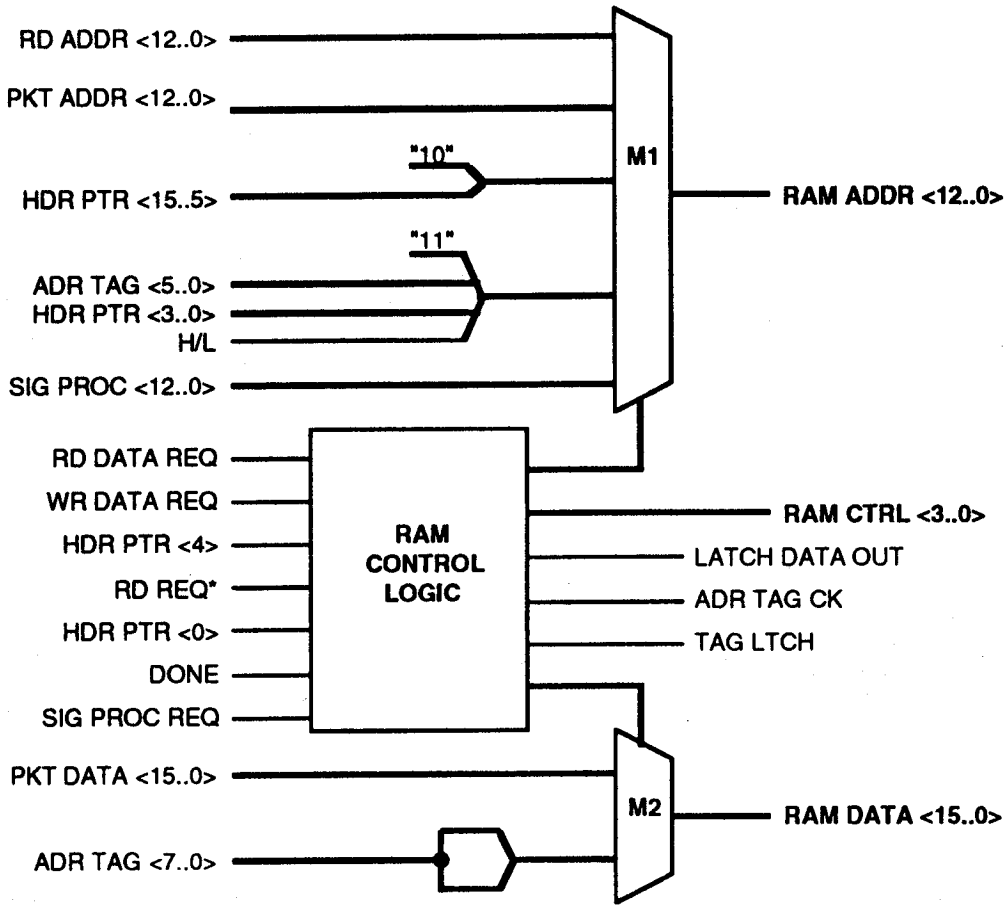

Figure 7-74. RAM Processing Block Diagram

- "11" | ADR TAG <5..0> | HDR PTR <3..0> | H/L. Address from the Port Side Processing module. This address will fetch the two word routing tag from the RAM memory. The most significant two bits are "11" to select the routing tag portion of the RAM memory.

- SIG PROC <12..0>. Address from the Signal Processor. This address is used by the Signal Processor for read and write accesses to the RAM memory.

Data Multiplexing. The multiplexor M2 selects the desired RAM write data for the current access as determined by the RAM Control Logic. Sources of RAM address are:

- PKT DATA <15..0>. Write data from the Port Side Processing module. This is the parallel input packet data to be stored in the data buffer.

- ADR TAG <7..0>. Address tag from the Port Side Processing module. This is the eight bit address tag used for routing the packet through the Switch Fabric.

RAM Control. The RAM Control Logic module will accept requests from all sources for accesses to the RAM memory. It will arbitrate the requests, then generate the proper

memory and register control signals to execute the desired access. Priority of accesses will be arbitrated as follows:

- RD DATA REQ. Requests for a read access to fetch the packet data to be transmitted to the Switch Fabric. This has the highest priority since it has the least amount of time between requests. The rate of requests is approximately 12.5 MHz (200 MHz / 16 Bits).
- WR DATA REQ. Requests for a write access to store the packet data to the data buffer. This has the second highest priority with a rate of request at approximately 9.7 MHz (155 MHz / 16 Bits).
- RD REQ*. Request from the Output Port Congestion Control for a read access to fetch the next address tag. This has the third highest priority with a rate of request at approximately 2.6 MHz (100 MHz / 39 Bits).
- Indirect Address Tag Fetch. Request from internal RAM Control Logic for a read access to the routing tag indirect address map. This has the fourth highest priority with a rate of request at approximately 183 KHz (9.7 MHz / 53 Wds/Pkt).
- Routing Tag Fetch. Request from internal RAM Control Logic for a read access to the routing tag memory. This has the fifth highest priority with a rate of request the same as the indirect address tag fetch of approximately 183 KHz.
- Routing Tag Write. Request from internal RAM Control Logic for a write access to the data buffer. This has the sixth highest priority with a rate of request the same as the indirect address tag fetch of approximately 183 KHz.
- SIG PROC REQ. Request from the Signal Processor for a read or write access to the RAM memory. This has the lowest priority since there is no real time requirement for these accesses.

7.2.1.2.5 Input Port Characteristics

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-29.

Table 7-29. Input Port ASIC Estimated Gate Count

FUNCTION	GATES/ FUNCTION	QUANTITY	TOTAL GATES
D FLIP-FLOP	7	325	2275
LOGIC GATE	1.75	750	1300
TOTAL + 20%			4300

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-30.

Table 7-30. Input Port ASIC Estimated Power Consumption

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	7	4	0.580 W
OUTPUT I/O	0	18	0	4	0.530 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	235	30	40	20	0.135 W
LOGIC GATE	370	240	100	20	0.080 W
TOTAL + 20%					1.9 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Input Port Power Characteristics. The Input Port will require one ASIC chip, two 8K by 8 bit memories, differential receivers and drivers, and minor support logic requiring an estimated ten square inches of board space. The total power consumption of the Input Port is summarized in Table 7-31.

Table 7-31. Input Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.9 W
RAM	1.5 W
I/O & SUPPORT	1.5 W
TOTAL	4.9 W

7.2.1.3 Output Port

The Output Port is identical to that used for the Sorted Banyan Network described in Section 7.2.1.2 except that the Multicast Banyan Input Port ASIC will be used. A summary of the Output Port characteristics are shown below.

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-29. This is the same ASIC as used on the Input Port and will require approximately 4300 gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-32. The values have been altered from the Input Port design to reflect the parts of the ASIC not used for the Output Port.

Table 7-32. Input Port ASIC Estimated Power Consumption For Output Port

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	4	4	0.440 W
OUTPUT I/O	0	18	0	1	0.345 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	105	30	25	20	0.095 W
LOGIC GATE	300	240	35	0	0.050 W
TOTAL + 20%					1.4 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Output Port Power Characteristics. The Output Port will require one ASIC chip, one 4K by 16 bit memory, differential receivers and drivers, and minor support logic requiring an estimated seven square inches of board space. The total power consumption of the Output Port is summarized in Table 7-33.

Table 7-33. Output Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.4 W
RAM	1.0 W
I/O & SUPPORT	1.0 W
TOTAL	3.4 W

7.2.1.4 Switch Fabric

7.2.1.4.1 Switch Fabric Configurations

The Self-Routing Multicast Banyan Switch requires both Banyan switching cells and standard Batcher Sorting Networks. The Batcher Sorting Networks are described in Section 7.1.1.4.1.2. The configurations for an 8 x 8, 16 x 16, and 32 x 32 Multicast Banyan Switch Fabric are shown in Figures 7-75 to 7-77.

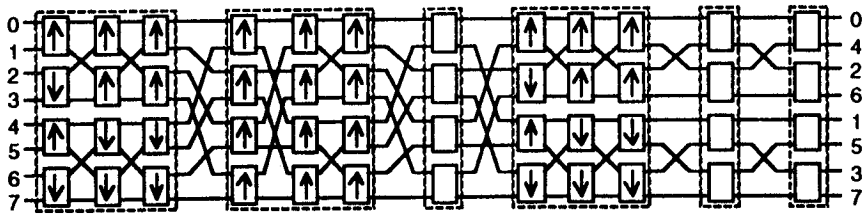

Figure 7-75. 8 x 8 Multicast Banyan Switch Fabric

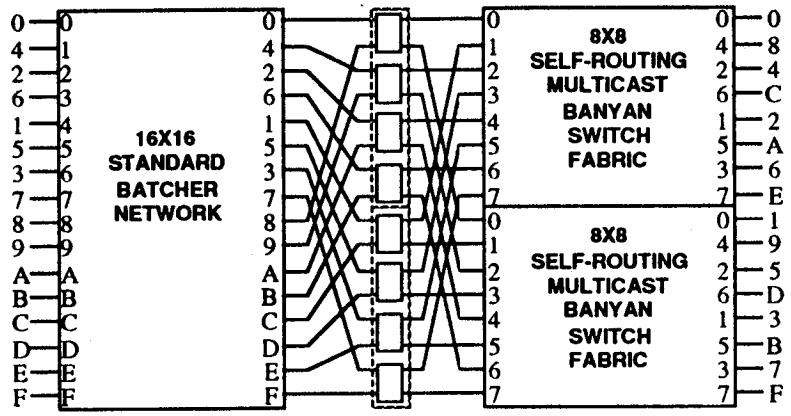

Figure 7-76. 16 x 16 Multicast Banyan Switch Fabric

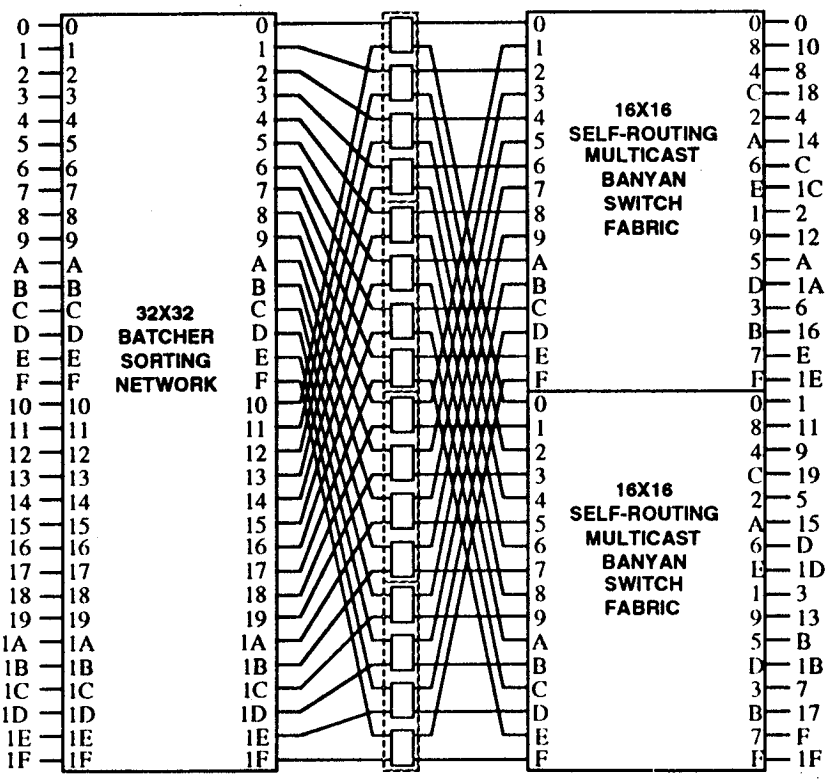

Figure 7-77. 32 x 32 Multicast Banyan Switch Fabric

7.2.1.4.2 Banyan ASIC Design

A block diagram of the Banyan ASIC is shown in Figure 7-78.

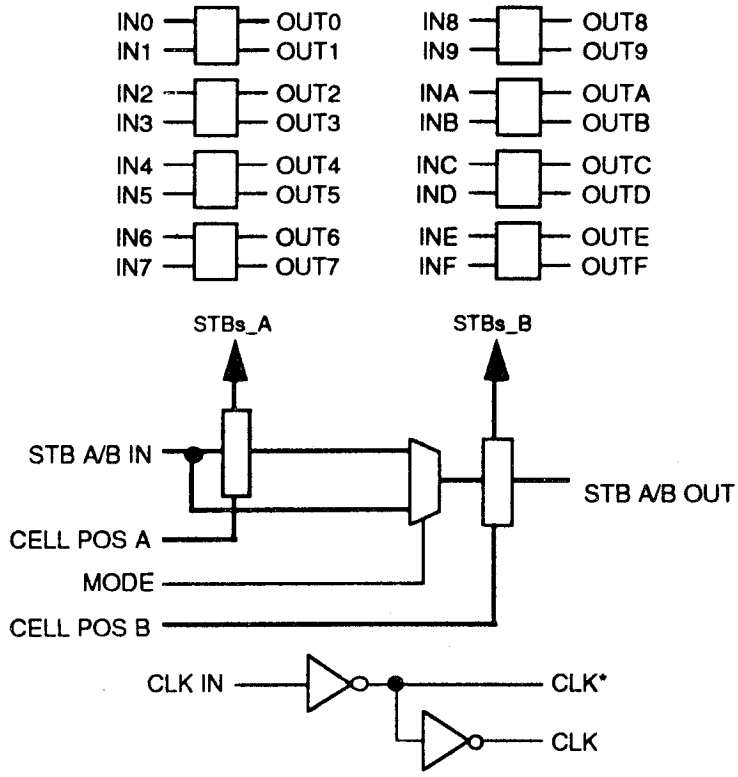

Figure 7-78. Multicast Banyan ASIC Block Diagram

Mode Control. Under most circumstances the MODE bit will be set such that the STB IN signal is routed to both strobe generation blocks. The STBs_A and STBs_B signals will therefore be identical. If, as in the case of the 8 x 8 Switch Fabric, two separate groups of four cells are required, the MODE bit will be set such that the input to the second block comes from the first block. This mode will allow the use of all eight cells in the 8 x 8 Switch Fabric.

Cell Position. The CELL POS signals are hardwired inputs to the ASIC. They inform the associated cells of the Switch Fabric size and the ASIC position within the Switch Fabric. This information is required by the Common Control logic to determine where the two address bits associated with the next switching cell are located within the routing tag. The required CELL POS values for the ASICs are shown in Figures 7-79 to 7-81 for the three Switch Fabric configurations. These figures show only the Banyan cells within the fabric; the Batcher networks interspersed throughout the Banyan cells are not shown for clarity. A description of how the CELL POS values are used is presented in the Common Control/Timing paragraph of this section.

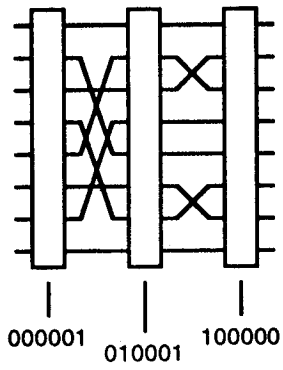

Figure 7-79. CELL POS Values for 8 x 8 Switch Fabric

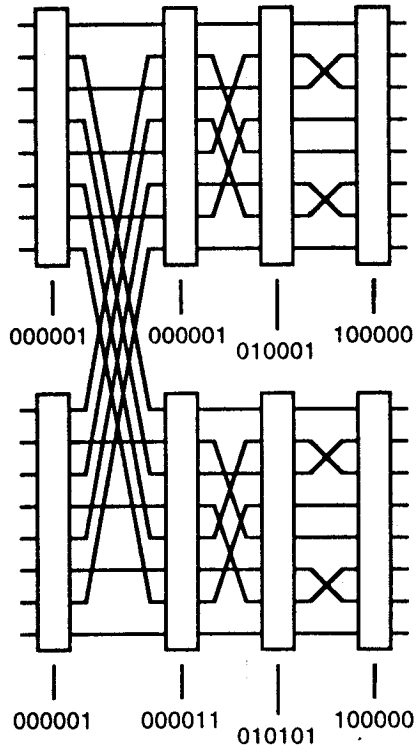

Figure 7-80. CELL POS Values for 16 x 16 Switch Fabric

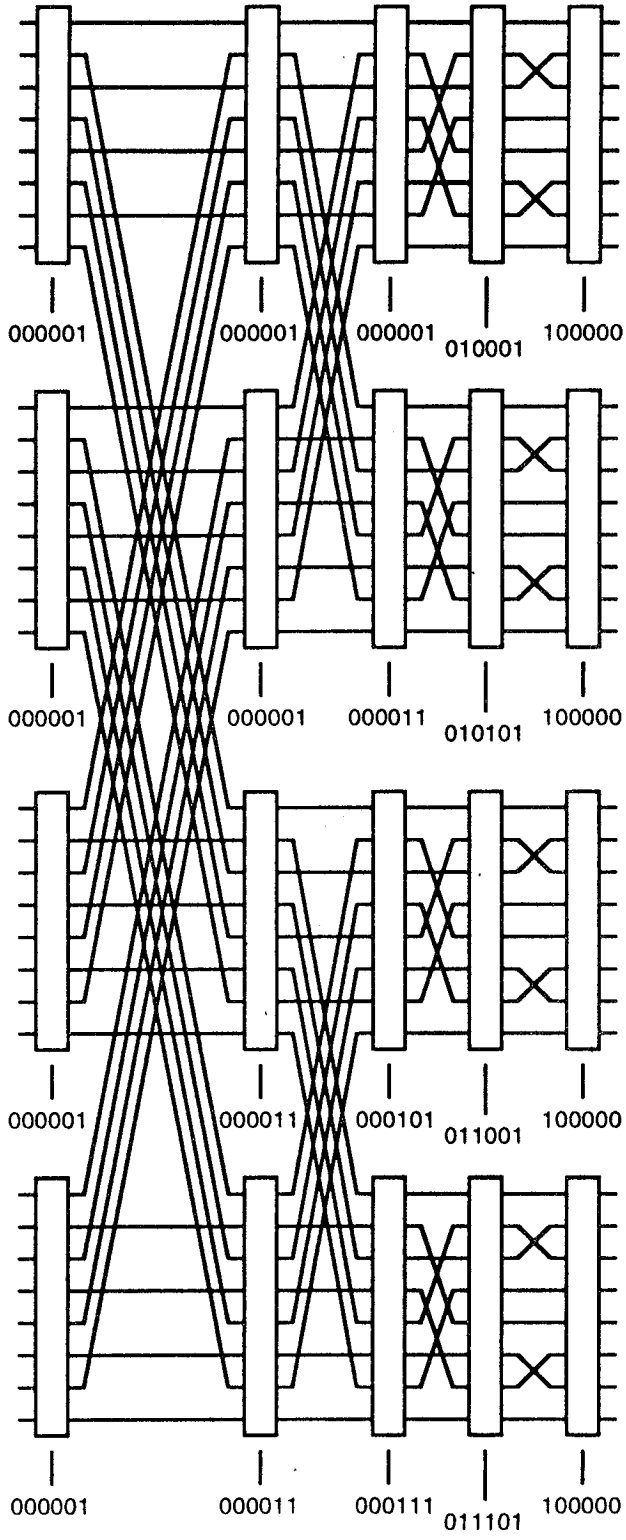

Figure 7-81. CELL POS Values for 32 x 32 Switch Fabric

Switching Cell. A block diagram of the Banyan switching cell is shown in Figure 7-82. All switch paths can be controlled independently with the only restriction that two inputs may not be routed to the same output. A detailed block diagram of the implementation is shown in Figure 7-83. A detailed timing diagram where the CELL POS value is "000001" is shown in Figure 7-84. In the figures, DS3* and DS4* are used only for the upper two cells of a given column; DS5* and DS6* are used only for the lower two cells.

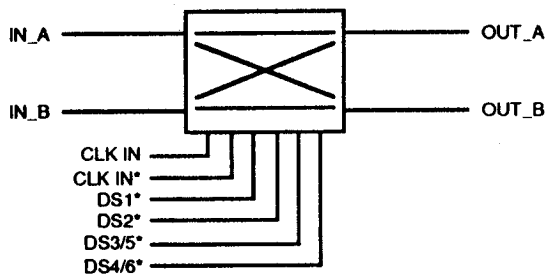

Figure 7-82. Multicast Banyan Switching Cell Block Diagram

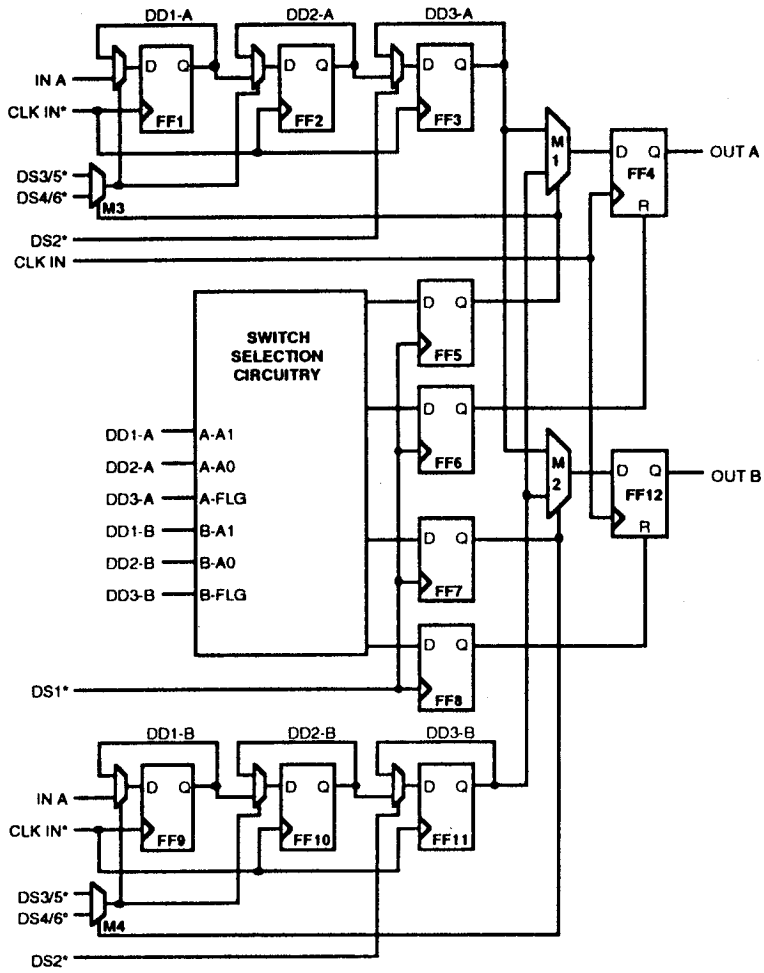

Figure 7-83. Multicast Banyan Switching Cell Detailed Block Diagram

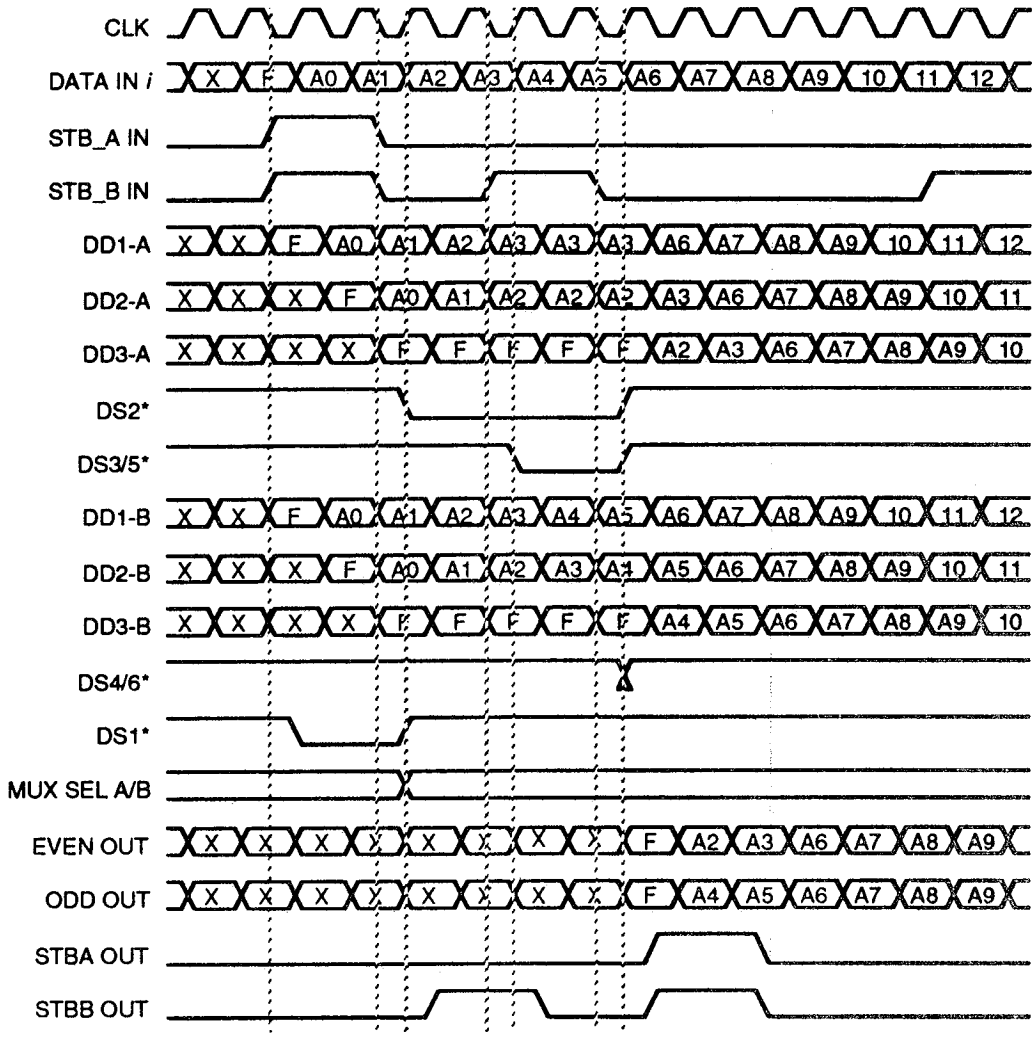

Figure 7-84. Multicast Banyan Switching Cell Timing Diagram Example

When the packet first enters the switching cell, the multiplexors for flip-flops FF1 to FF3 and FF9 to FF11 are set such that the lower input is passed to the output. Flip-flop FF1 mid-bit samples the incoming serial data for the A side input, IN A. Flip-flop FF2 and FF3 provide for a two bit storage such that, on the rising edge of DS1*, FF3, FF2, and FF1 contain the flag bit, the most significant routing bit for this cell, and the least significant routing bit for this cell, respectively. Flip-flops FF9, FF10, and FF11 provide the same function for the B side input, IN B. The outputs of these six flip-flops are fed into the switch selection circuitry which will follow the truth table specified in Table 7-34. The outputs of the switch selection circuitry are latched into flip-flops FF5 to FF8 on the rising edge of DS1*. When set, FF5/7 will cause multiplexor M1/2 to pass DD1-A to OUT A/B; when reset M1/2 will pass DDA-B to OUT A/B. When set, FF6/8 will keep the output flip-flop FF4/12 in reset (outputs at low); when reset, FF4/12 will pass the serial input data selected by multiplexor M1/2.

Table 7-34. Switch Selection Circuitry Truth Table

FLAG A	A1A0 (IN A)	FLAG B	A1A0 (IN B)	Q5	Q6	Q7	Q8
INVALID	X	INVALID	X	X	RST	X	RST
INVALID	X	VALID	BOTH	B	EN	B	EN
INVALID	X	VALID	EVEN	B	EN	X	RST
INVALID	X	VALID	BAD	X	RST	X	RST
INVALID	X	VALID	ODD	X	RST	B	EN
VALID	BOTH	INVALID	X	A	EN	A	EN
VALID	BOTH	VALID	BAD	A	EN	A	EN
VALID	EVEN	INVALID	X	A	EN	X	RST
VALID	EVEN	VALID	BAD	A	EN	X	RST
VALID	EVEN	VALID	ODD	A	EN	B	EN
VALID	BAD	INVALID	X	X	RST	X	RST
VALID	BAD	VALID	BOTH	B	EN	B	EN
VALID	BAD	VALID	EVEN	B	EN	X	RST
VALID	BAD	VALID	BAD	X	RST	X	RST
VALID	BAD	VALID	ODD	X	RST	B	EN
VALID	ODD	INVALID	X	X	RST	A	EN
VALID	ODD	VALID	EVEN	B	EN	A	EN
VALID	ODD	VALID	BAD	X	RST	A	EN

The key for the truth table above is given below.

- "VALID" Flag bit is when the F bit at the head of the routing tag is "1"; "INVALID" is when it is "0".
- A1A0 definitions are:
  - "BOTH" implies the selected input is switched to both outputs.
  - "EVEN" implies the selected input is switched to the even output (OUT A).
  - "BAD" implies that the selected input is an invalid packet and should be treated as though the Flag bit is not set.
  - "ODD" implies the selected input is switched to the odd output (OUT B)
- Q5 and Q7 definitions are:
  - "X" implies "Don't Care". Since the output will be disabled by Q6 or Q8, the setting of the multiplexor does not matter.

- "A" implies that the selected output will come from the A side input.
- "B" implies that the selected output will come from the B side input.

- Q6 and Q8 definitions are:
  - "EN" implies that the output flip-flop, FF4 or FF12, will be enabled to pass the output of the associated multiplexor.
  - "RST" implies that the associated output flip-flop, FF4 or FF12, will be held in reset until the reception of the next packet routing tag.

Once the selection of the output multiplexors has been determined, the DS2*, DS3/5*, and DS4/6* strobe signals will control the input multiplexors to FF1 to FF3 and FF9 to FF11 as shown in the timing diagrams. These strobes will be enabled to store the flag bit and the two address bits for transmission to the next cell. Multiplexor M3/4 will gate the proper strobes to the storage flip-flops. DS3* and DS4* are strobes for the packet to be sent out on the even output; DS5* and DS6* are strobes for the odd output. The timing of the strobe signals are dependant on the COL POS input to the strobe generation blocks which will be described below in the "Common Control/Timing" section.

Common Control/Timing. The input clock, CLK IN, is buffered and provided to each cell in true and inverted form. Figure 7-85 is a block diagram of the circuitry required to generate the DS1* strobe signal, STB_AOUT, and STB_BOUT. Figure 7-86 is a timing diagram of this circuit. Figures 7-87 and 7-88 show the block diagram and timing diagram for the generation of the DS2* to DS6* strobe signals.

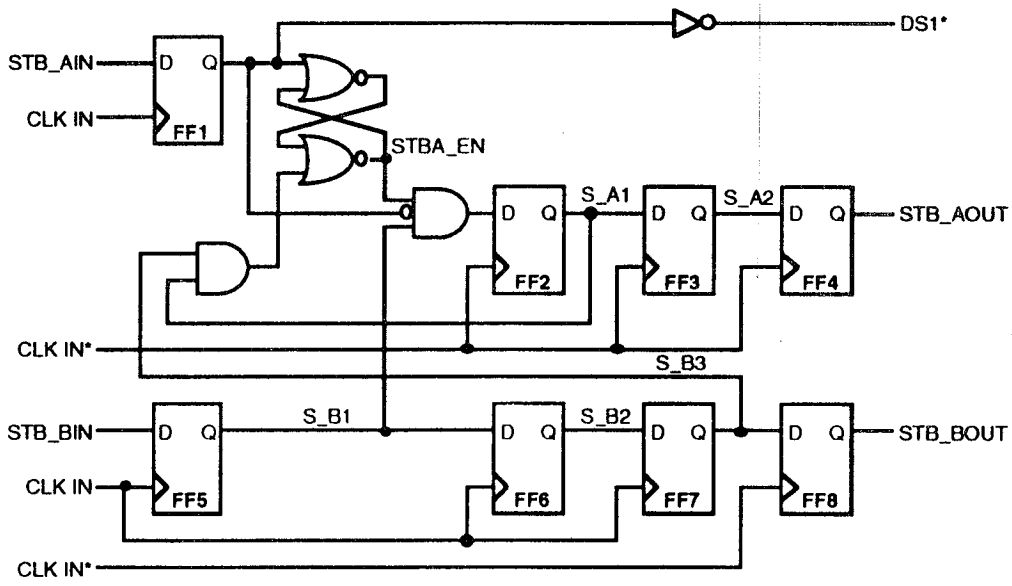

Figure 7-85. DS1*, STB_AOUT, STB_BOUT Generation Block Diagram

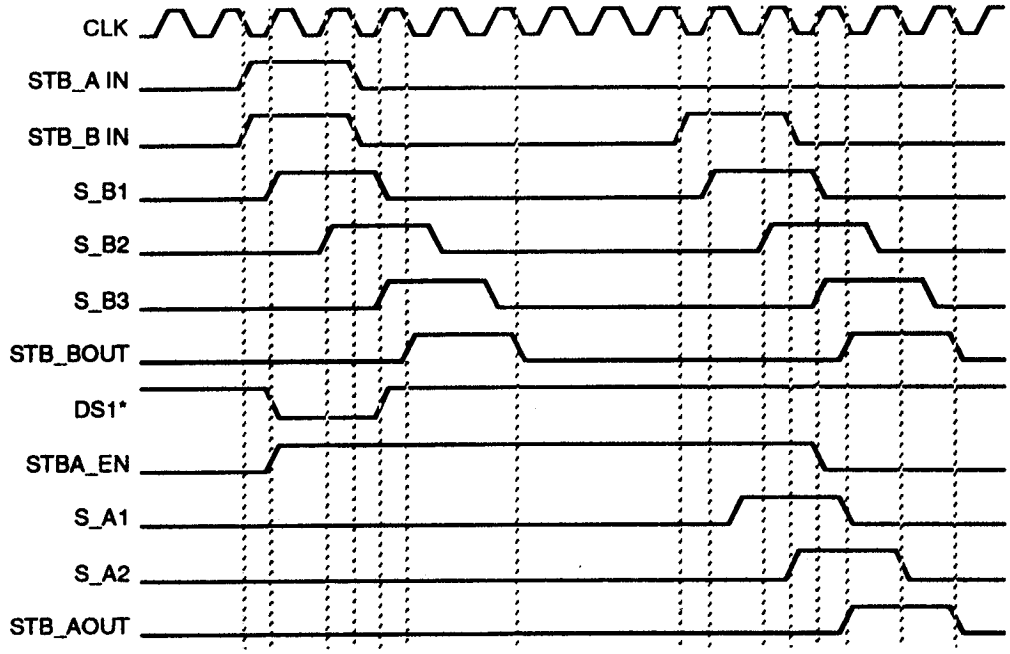

Figure 7-86. DS1*, STB_AOUT, STB_BOUT Generation Timing Diagram

Flip-flops FF5 to FF8 mid-bit sample the STB_B IN signal and provide for the required three clock delay to form STB_B OUT. Flip-flop FF1 mid-bit samples the STB_A IN signal. The output of FF5 is inverted to provide the required DS1* strobe signal. Flip-flops FF2 to FF4 and the gates provide the circuitry for generation of STB_A OUT which moves the STB_A IN signal to the next active STB_B IN signal then delays it by three clock periods.

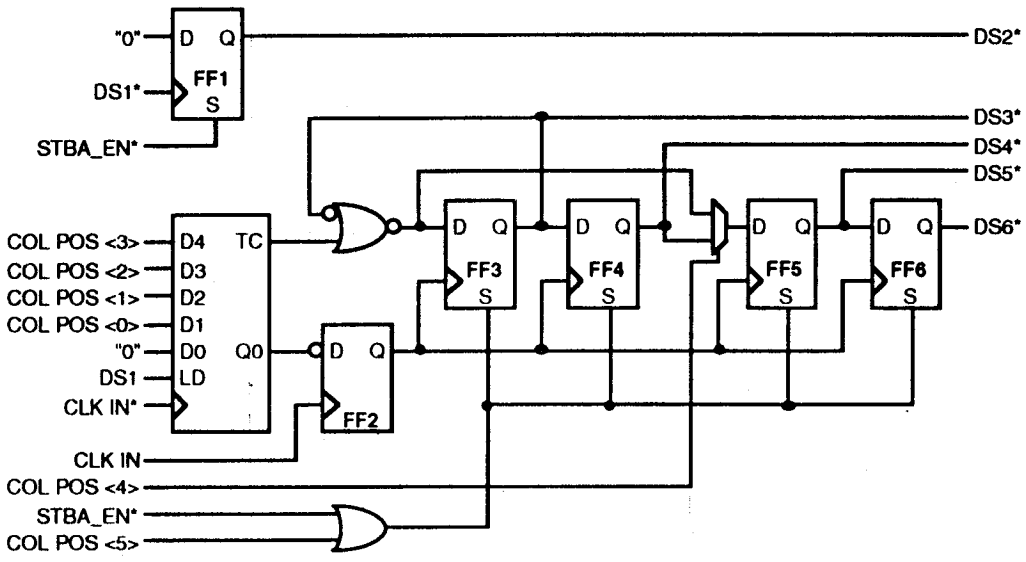

Figure 7-87. DS2* to DS6* Generation Block Diagram

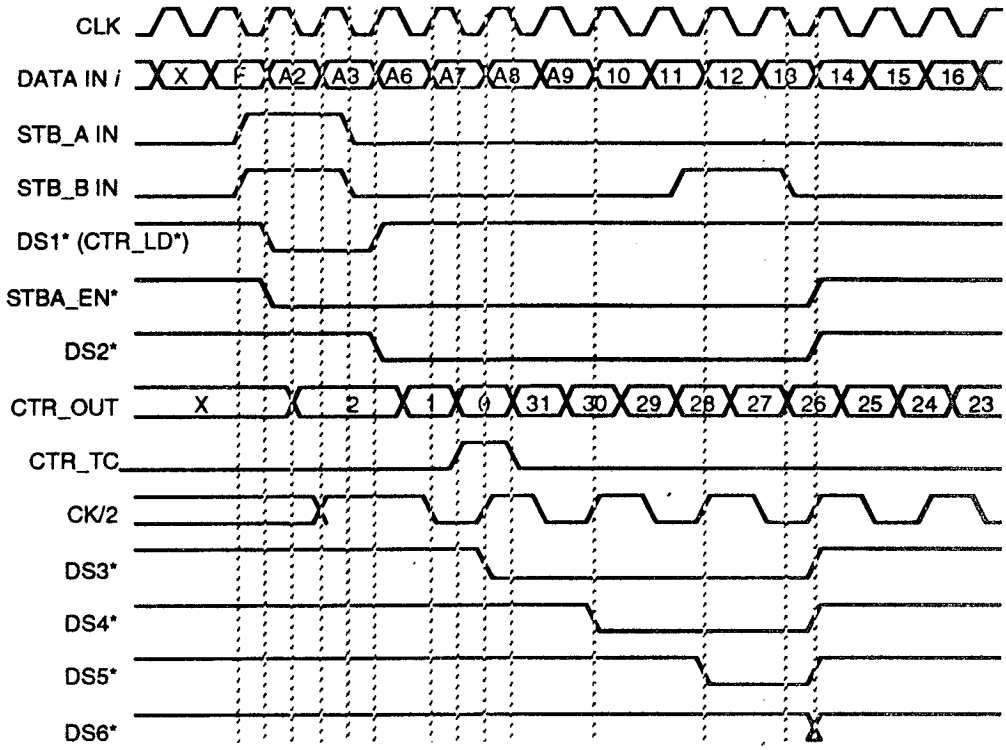

Figure 7-88a. DS2* to DS6* Generation Timing Diagram (Type 1)

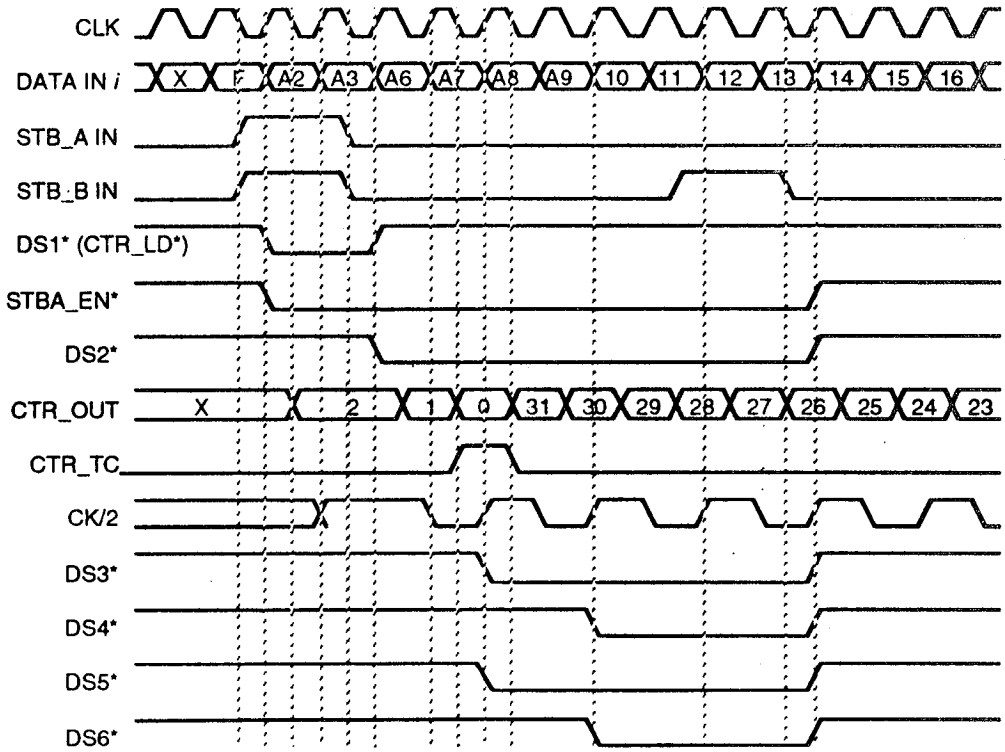

Figure 7-88b. DS2* to DS6* Generation Timing Diagram (Type 2)

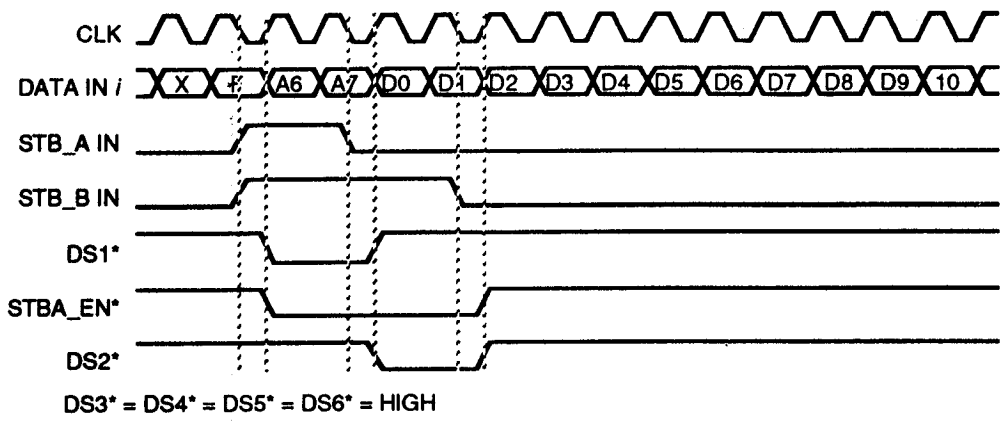

Figure 7-88c. DS2* to DS6* Generation Timing Diagram (Type 3)

Flip-flop FF1 is used to generate the DS2* strobe signal. The signal goes active on the rising edge of DS1* and is cleared on STBA_EN*. The five bit counter is used to count to the desired position within the routing tag to obtain the two address bits for the next switching cells. This counter is preloaded with the COL POS <3..0> code multiplied by two. When the terminal count TC of this down counter is reached, the DS3* strobe is generated by flip-flop FF3. Flip-flop FF2 is used to provide a synchronized divide by two of the symbol clock, CLK IN. Flip-flop FF4 generates the DS4* signal on the next rising edge of the output of FF2. When COL POS <4> is low, flip-flops FF5 and FF6 will generate DS5* and DS6* as shown in Figure 7-88b; when high DS5* and DS6* will be as shown in Figure 7-88a. When COL POS <5> is high, indicating a cell in the rightmost column, the DS3* to DS6* strobes are inhibited and DS2* is generated in the normal manner as shown in Figure 7-88c.

7.2.1.4.3 Multicast Banyan ASIC Characteristics

Multicast Banyan ASIC Gate Count. The estimated gate count for the Banyan ASIC is shown in Table 7-35.

Table 7-35. Multicast Banyan ASIC Estimated Gate Count

FUNCTION	GATES/ FUNCTION	QUANTITY	TOTAL GATES
D FLIP-FLOP	12	120	1440
LOGIC GATE	2	320	640
TOTAL + 20%			2500

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Multicast Banyan ASIC Power Consumption. The estimated power consumption for the Banyan ASIC is shown in Table 7-36.

Table 7-36. Banyan ASIC Estimated Power Consumption

FUNCTION	POWER/ DEVICE	QUANTITY	TOTAL POWER
INPUT I/O	6.4	19	0.122 W
OUTPUT I/O	7.81	18	0.141 W
D FLIP-FLOP	2.77	120	0.332 W
LOGIC GATE	.4	300	0.128 W
TOTAL + 20%			0.85 W

7.2.1.4.4 Batcher ASIC Characteristics

The Batcher ASIC is identical to the one used for the Sorted Banyan Network. The characteristics are repeated below.

Batcher ASIC Gate Count. The estimated gate count for the Batcher ASIC is shown in Table 7-37.

Table 7-37. Batcher ASIC Estimated Gate Count

FUNCTION	GATES/ FUNCTION	QUANTITY	TOTAL GATES
D FLIP-FLOP	12	90	1080
LOGIC GATE	6	180	900
TOTAL + 20%			2400

The logic gate count includes inverters and 2 to 4 input and, nand, or, nor, and xor logic.

Batcher ASIC Power Consumption. The estimated power consumption for the Batcher ASIC is shown in Table 7-38.

Table 7-38. Batcher ASIC Estimated Power Consumption

FUNCTION	POWER/ DEVICE	QUANTITY	TOTAL POWER
INPUT I/O	6.4	13	0.083 W
OUTPUT I/O	7.81	10	0.078 W
D FLIP-FLOP	2.77	90	0.249 W
LOGIC GATE	1.2	150	0.180 W
TOTAL + 20%			0.7 W

7.2.1.4.5 Switch Fabric Characteristics

ASIC Count. The number of ASICs required for each of the three switch fabric configurations is shown in Table 7-39.

Table 7-39. Switch Fabric ASIC Count

ASIC COUNT	8 x 8	16 x 16	32 x 32
BANYAN	2	4	10
BATCHER	3	14	48
TOTAL	5	18	58

Switch Fabric Power Consumption. The power consumption of the Switch Fabric is summarized in Table 7-40 for the three configurations. Miscellaneous support circuitry is estimated at about 0.5 watts per port.

Table 7-40. Switch Fabric Power Consumption

	8 x 8	16 x 16	32 x 32
BANYAN MATRIX	1.7 W	3.4 W	8.5 W
BATCHER MATRIX	2.1 W	9.8 W	33.6 W
MISC SUPPORT	4 W	8 W	16 W
TOTAL	7.8 W	21.2 W	58.1 W

7.2.1.5 Sorted Banyan Summary

7.2.1.5.1 Power Consumption

The power consumption of the Sorted Banyan Switch is summarized in Table 7-41 for the three switch configurations.

Table 7-41. Multicast Banyan Switch Power Consumption

	8 x 8	16 x 16	32 x 32
INPUT PORT	39.2 W	78.4 W	156.8 W
OUTPUT PORT	27.2 W	54.4 W	108.8 W
SWITCH FABRIC	7.8 W	201.2 W	58.1 W
TOTAL	74.2 W	154 W	323.7 W
TOTAL/PORT	9.3 W	9.6 W	10.1 W

7.2.1.5.2 Mass/Size

A circuit board of 30 square inches will support one Input Port and one Output Port. One board will be required for every port that the switch must service. The 8 x 8 and 16 x 16 Switch Fabric configurations can also fit on one board of this size; the 32 x 32 configuration will require two or three boards. In addition, one board will be required for the generation of common control and timing signals. This includes the following functions:

- Generation of Switch Fabric clocks.
- Generation and control of the Congestion Control Message.
- Generation of strobes to the Input Ports, Output Ports, and the Switch Fabric.

Therefore, the total number of modules required is 10, 18, and 35 (or 36) for the 8 x 8, 16 x 16, and 32 x 32 switch configurations, respectively.

7.2.1.5.3 ASIC Design Complexity

There should be no problem in designing the three ASICs required for this approach. All ASICs have fairly low gate count and moderate power consumption. A disadvantage to this approach is that three ASIC designs are required as opposed to only two for other approaches.

The design of these ASICs did not include added circuitry which may be required to implement a viable redundancy scheme. However, the addition of this circuitry is not expected to significantly increase the complexity of the ASIC design.

7.2.1.5.4 Fault Tolerance

The fault tolerance of the basic Multicast Banyan Switch is relatively good compared to other switch architectures. A failure of an Input Port will not affect any other Input Port. This is also true for the Output Ports. A failure in a cell of the Switch Fabric will only affect 1/N of the paths through the fabric where N is the size of the Switch Fabric. For these reasons, it is much more efficient to implement a 1-for-N redundancy scheme for this switch than a 1-for 1 scheme, although it may be very difficult to implement in the Switch Fabric.

The Congestion Control Message path is a ring type topology. If any Input Port fails, the entire ring becomes ineffective, causing the switch to completely shut down. For this reason, it is imperative that this ring be redundant. Also, a detected failure in a given Input Port should cause the ring input to pass through unaffected to the ring output.

7.2.2 Self-Routing Multicast Crossbar-Based Network

A block diagram of the Self-Routing Multicast Crossbar-Based Network is shown in Figure 7-89 below.

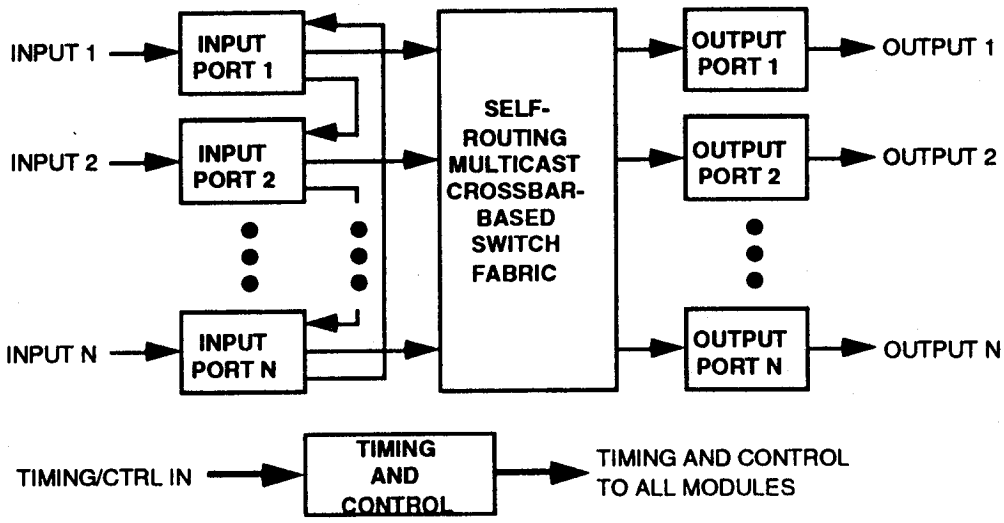

Figure 7-89. Self-Routing Multicast Crossbar-Based Network Block Diagram

Serial data from the demodulators is routed through the switch to the modulators in the following steps:

- The serial packet data enters the respective Input Port where it is stored in a 128 packet FIFO type buffer.
- As each packet is stored in the packet buffer, the routing tag associated with the packet is obtain from the routing map memory then stored with the packet in the buffer.
- The Input Port performs Output Port congestion control to ensure that there are no blocked packets at the Output Port. The Input Ports convey their desired destination to each other by means of a serial daisy chained message every packet slot time. When an Input Port reserves an Output Port for the next packet transmission, no other Input Port further down the daisy chain may select the reserved Output Port. The four oldest in the packet buffer will be examined for transmission through the Switch Fabric to minimize head-of-line blocking.
- The selected packet, if any, from each Input Port is sent to the Switch Fabric. The Switch Fabric will route packets from all Input Ports to the desired Output Ports as specified by the routing tag. The Switch Fabric will need to operate at a higher rate than the input serial data to compensate for the added bits for the routing tag and for the lost packet transfers through the Switch Fabric due to head-of-line blocking.

- The Output Port will accept serial data from the Switch Fabric then store it in a 128 packet FIFO type buffer. The serial packet data will be output to the modulator for downlink transmission.

7.2.2.1 Formats

7.2.2.1.1 Routing Tag Format

The format for the routing tag appended to each packet is shown in Figure 7-90.

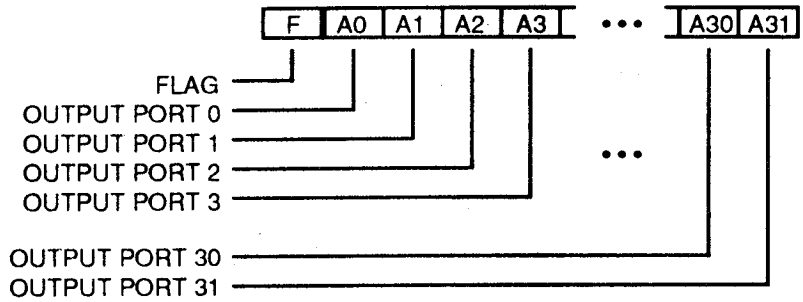

Figure 7-90. Multicast Crossbar Routing Tag Format

The first bit is the FLAG bit which must be set for a valid packet. The next thirty-two bits represent each of the Output Ports. Bit A0 is set if Output Port 0 is a desired destination, bit A1 for Output Port 1, bit A2 for Output Port 2, etc. For an 8 x 8 Switch Fabric, only the first eight bits (A0 to A7) are required; for a 16 x 16, the first sixteen bits (A0 to A15) are required; and for a 32 x 32, all thirty-two bits are required. For the current approach, thirty-two bits will always be used so that the Output Ports will more easily determine the starting location of the information bits.

7.2.2.1.2 Output Port Congestion Control Message Format

The format for the Output Port Congestion Control Message is identical to that used for the Sorted Banyan Network described in Section 7.1.1.1.2.

7.2.2.2 Input Port

The Input Port is very similar to that used for the Self-Routing Multicast Banyan Network described in Section 7.2.1.2. The only difference is that the Address Translator and the parallel to serial shift register R2 are not required. The address tag for the packet will be the contents of register R1. The difference in gate count and power consumption for this change is insignificant so the Multicast Banyan Network results will be used. A summary of these characteristics are shown below.

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-42.

Table 7-42. Input Port ASIC Estimated Gate Count

FUNCTION	GATES/ FUNCTION	QUANTITY	TOTAL GATES
D FLIP-FLOP	7	325	2275
LOGIC GATE	1.75	750	1300
TOTAL + 20%			4300

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-43.

Table 7-43. Input Port ASIC Estimated Power Consumption

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	7	4	0.580 W
OUTPUT I/O	0	18	0	4	0.530 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	235	30	40	20	0.135 W
LOGIC GATE	370	240	100	20	0.080 W
TOTAL + 20%					1.9 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Input Port Power Characteristics. The Input Port will require one ASIC chip, two 8K by 8 bit memories, differential receivers and drivers, and minor support logic requiring an estimated ten square inches of board space. The total power consumption of the Input Port is summarized in Table 7-44.

Table 7-44. Input Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.9 W
RAM	1.5 W
I/O & SUPPORT	1.5 W
TOTAL	4.9 W

7.2.2.3 Output Port

The Output Port is identical to that used for the Multicast Banyan Network described in Section 7.1.1.3 except that the Multicast Banyan Input Port ASIC will be used. A summary of the Output Port characteristics are shown below.

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-42. This is the same ASIC as used on the Input Port and will require approximately 4300 gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-45. The values have been altered from the Input Port design to reflect the parts of the ASIC not used for the Output Port.

Table 7-45. Input Port ASIC Estimated Power Consumption For Output Port

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	4	4	0.440 W
OUTPUT I/O	0	18	0	1	0.345 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	105	30	25	20	0.095 W
LOGIC GATE	300	240	35	0	0.050 W
TOTAL + 20%					1.4 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Output Port Power Characteristics. The Output Port will require one ASIC chip, one 4K by 16 bit memory, differential receivers and drivers, and minor support logic requiring an estimated seven square inches of board space. The total power consumption of the Output Port is summarized in Table 7-46.

Table 7-46. Output Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.4 W
RAM	1.0 W
I/O & SUPPORT	1.0 W
TOTAL	3.4 W

7.2.2.4 Switch Fabric

7.2.2.4.1 Switch Fabric Configurations

The configurations for an 8 x 8, 16 x 16, and 32 x 32 Switch Fabric for the Multicast Crossbar-Based Network are shown in Figures 7-91 to 7-93. The Multicast Crossbar ASIC will accommodate thirty-two inputs and eight outputs. The 8 x 8, 16 x 16, and 32 x 32 Switch Fabric configurations will require one, two, and four ASICs, respectively.

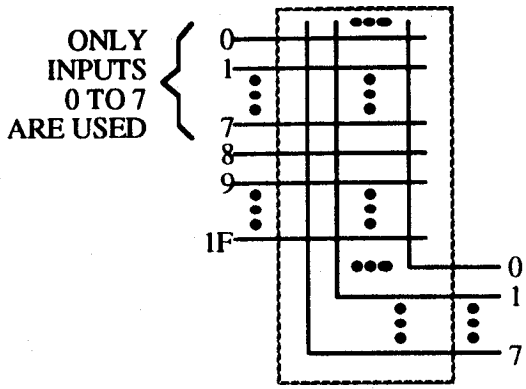

Figure 7-91. 8 x 8 Multicast Crossbar-Based Switch Fabric

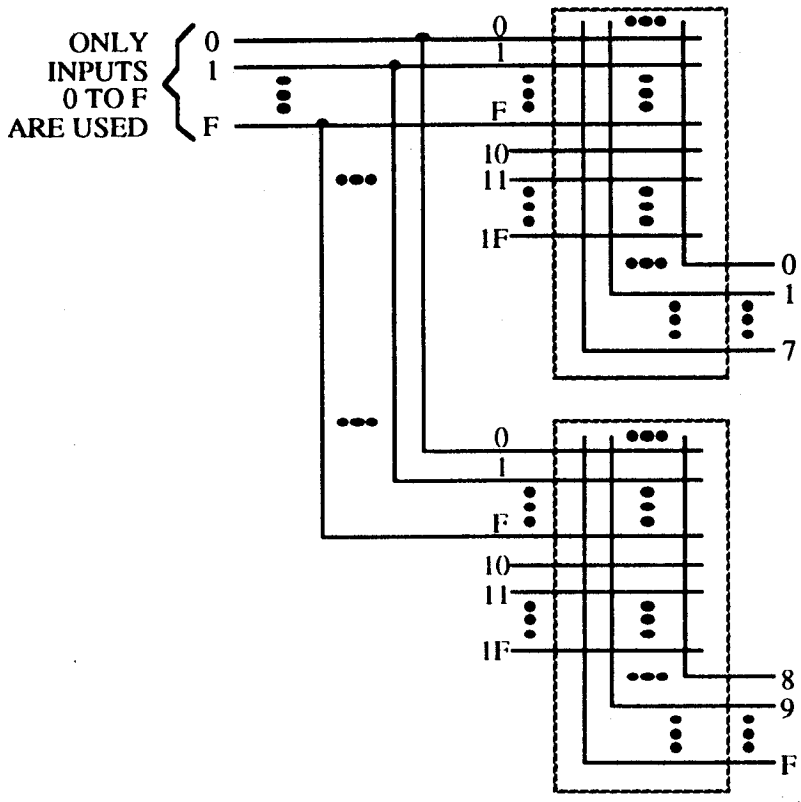

Figure 7-92. 16 x 16 Multicast Crossbar-Based Switch Fabric

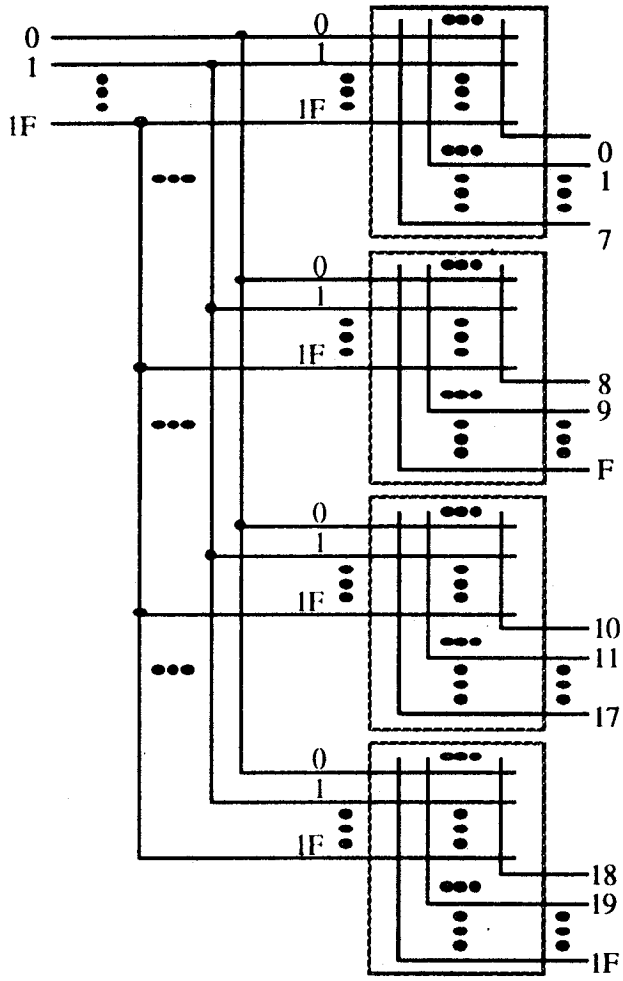

Figure 7-93. 32 x 32 Multicast Crossbar-Based Switch Fabric

7.2.2.4.2 Multicast Crossbar ASIC Design

A block diagram of the Multicast Crossbar ASIC is shown in Figure 7-94.

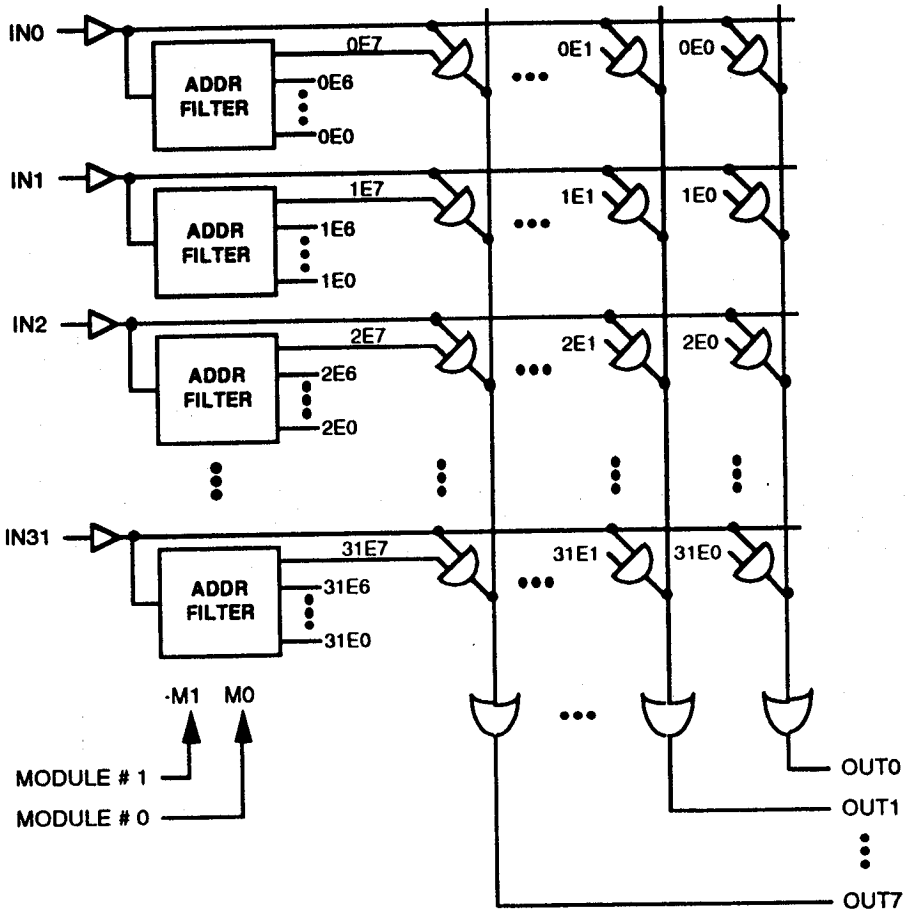

Figure 7-94. Multicast Crossbar ASIC Block Diagram

Mode Control, Module # 1 and Module # 0 (M1, M0). M1 and M0 determine the definition of the eight outputs of the ASIC as shown in Table 7-47. These two bits are used by the ADDR FILTER module to determine if the desired output for the current packet exists in this ASIC.

Table 7-47. Definition Of M1 and M0

M1 M0	SWITCH FABRIC DEFINITION OF ASIC OUTPUTS 0 TO 7
0 0	FABRIC OUTPUTS 0 TO 7
0 1	FABRIC OUTPUTS 8 TO 15
1 0	FABRIC OUTPUTS 16 TO 23
1 1	FABRIC OUTPUTS 24 TO 31

Address Filter. The address filter is responsible for determining if the destination Output Port is associated with this ASIC, and, if so, which ASIC output(s) is associated with this port. A block diagram of this circuit is shown in Figure 7-95. A timing diagram is shown in Figure 7-96.

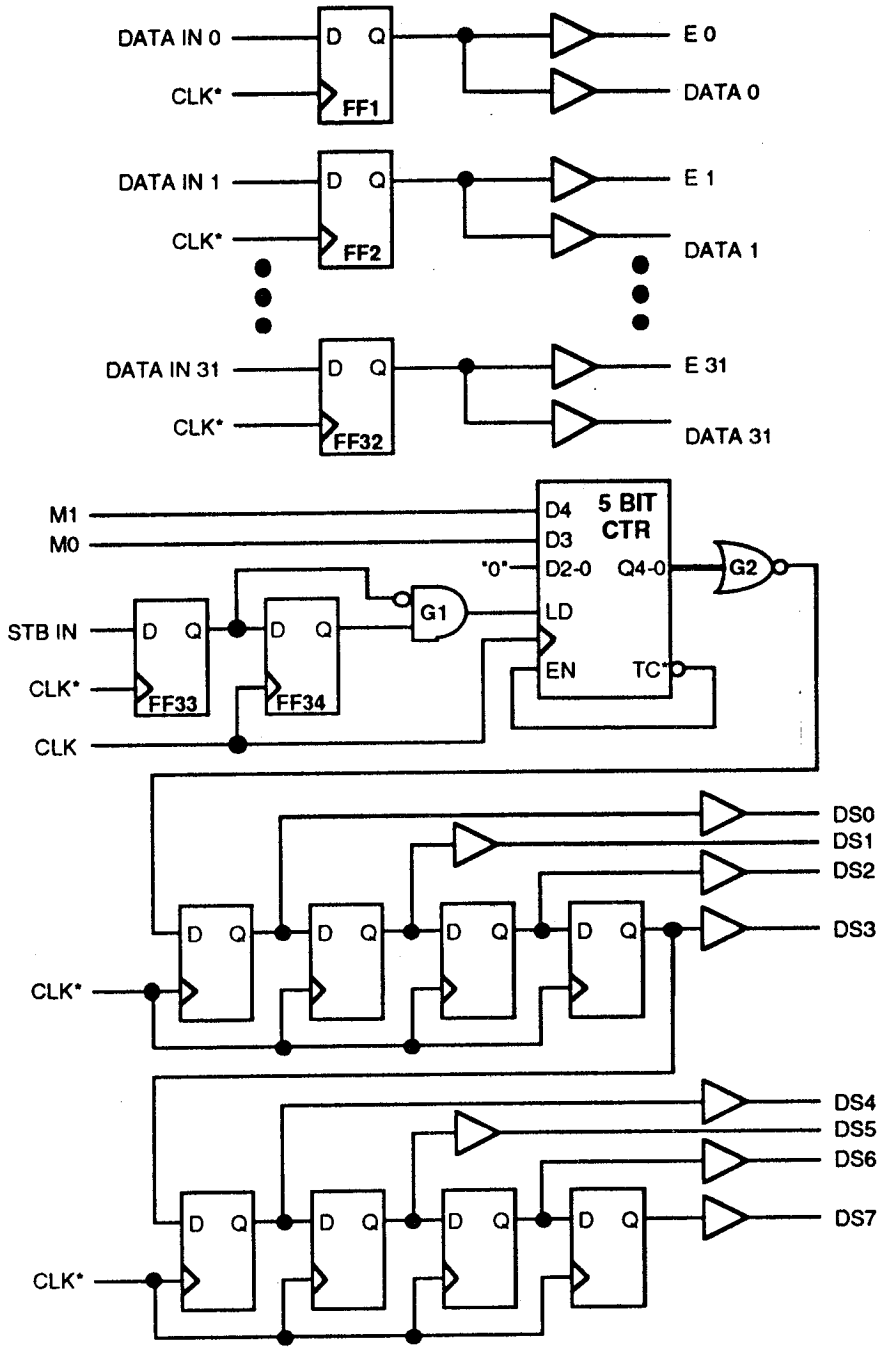

Figure 7-95. Address Filter Block Diagram

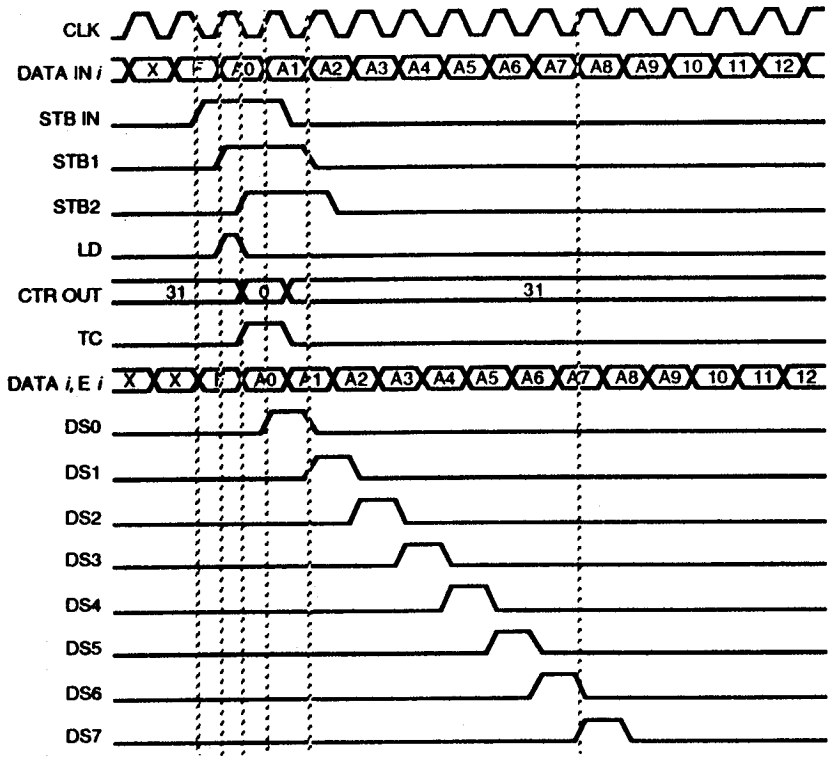

Figure 7-96a. Address Filter Timing Diagram (M1M0 = 00)

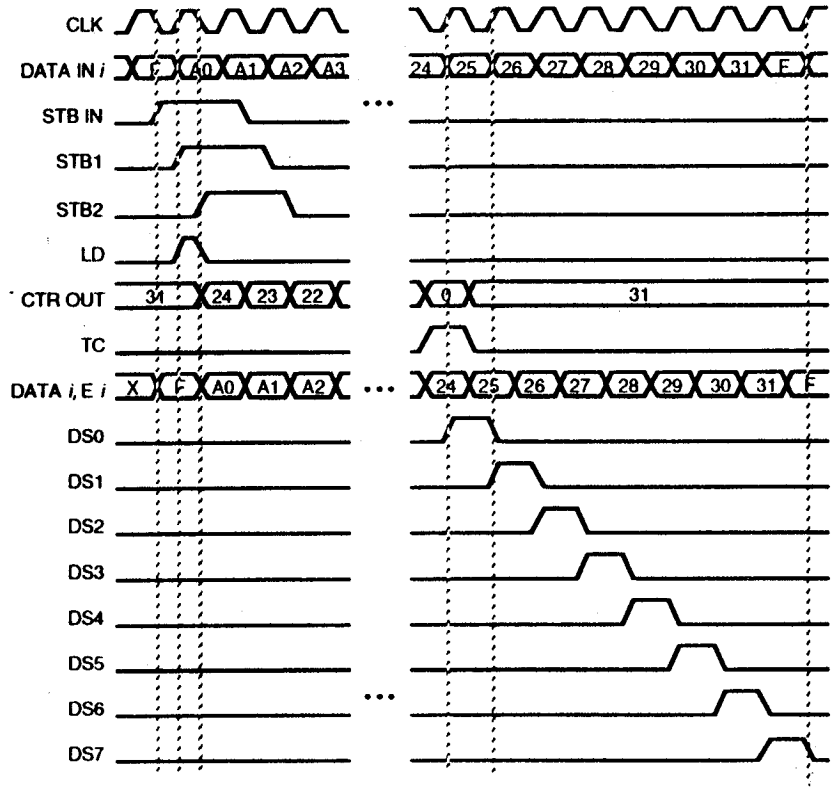

Figure 7-96b. Address Filter Timing Diagram (M1M0 = 01)

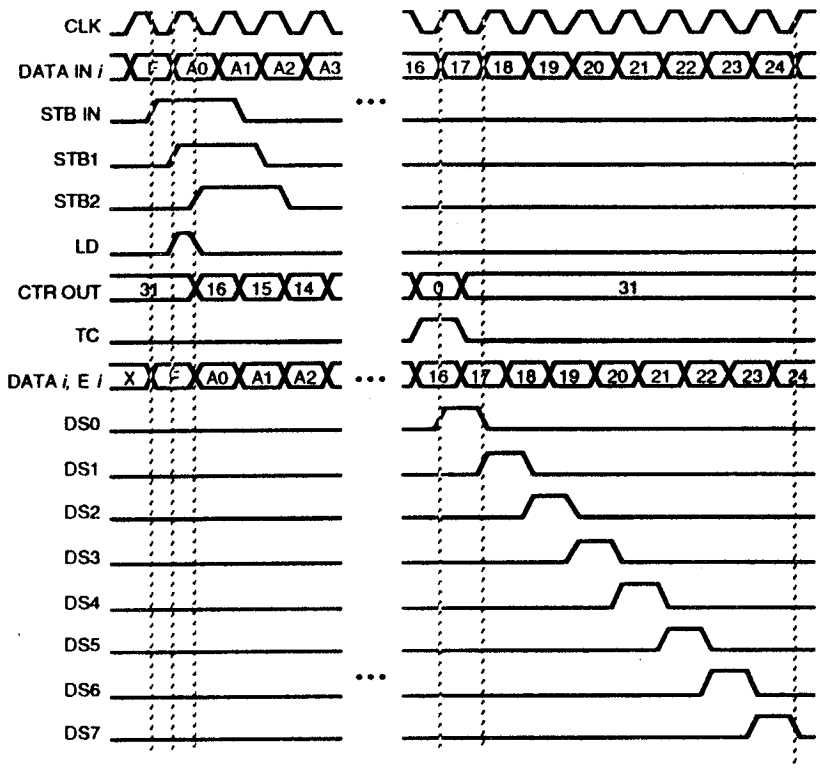

Figure 7-96c. Address Filter Timing Diagram (M1M0 = 10)

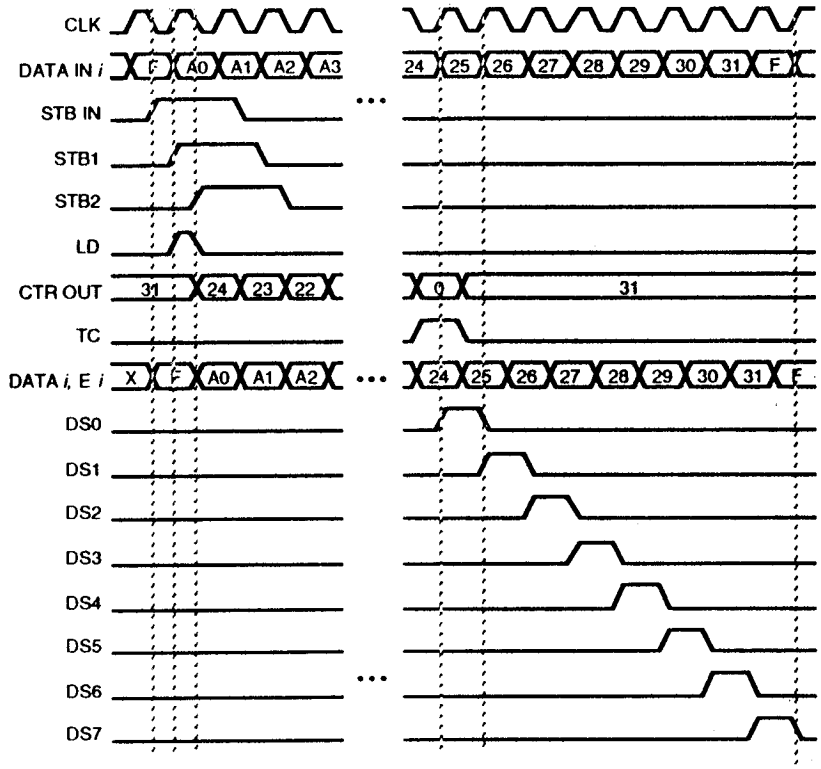

Figure 7-96d. Address Filter Timing Diagram (M1M0 = 11)

Flip-flops FF1 to FF32 mid-bit sample the thirty-two input signals. This information is provided to the Output Selection Circuitry.

Flip-flops FF33 and FF34 and gate G1 generate a load pulse for the counter at the beginning of each packet transmission as shown in the timing diagrams. The number loaded into the counter is dependant on the state of M1 and M0 as shown in Table 7-48.

Table 7-48. Definition Of Mode Bits M1, M0

M1	M0	CTR PRESET (DEC)	Valid ASIC Outputs
0	0	0	OUT 0 to OUT 7
0	1	8	OUT 8 to OUT 15
1	0	16	OUT 16 to OUT 23
1	1	24	OUT 24 to OUT 31

The gate G2 detect the all zero condition on the counter. When this occurs, a pulse is passed through the lower eight flip-flops to produce the DS0 to DS7 signals shown in the timing diagrams of Figure 7-96. The signals DS*i* therefore represents when the *i*th input is valid for sampling by the Output Selection Circuitry.

Output Selection. One Output Selection circuitry is responsible for accepting the requests from all the address filters for a particular output. Eight Output Selection modules exists in the ASIC, one for each output. A block diagram and timing diagram are shown in Figures 7-97 and 7-98.

The 32-to-5 encoder will convert the 32 inputs, one from each address filter, to a five bit code. A simple algorithm was selected assuming that, at most, one input will be active at any time (proper operation is not guaranteed if more than one input is active). This condition should always exist since the Output Congestion Control circuitry in the Input Port prevents more than one source for any given destination Output Port. The five bit code output from the encoder will be unique for every selected input. If no input is selected, the NONE output of the encoder will be high.

Each of the eight Output Selection circuits will be strobed on their respective DS*i* signal. This rising edge of this signal corresponds to the mid-bit position of the address tag bit associated with this output. The output of the encoder is latched on this rising edge.

The output of the five bit latch is fed to the select inputs of the multiplexor to route the desired serial data to the output flip-flop FF7. If the multiplexor is disabled, the input to FF7 will be held low until the next occurrence of DS1.

Since each of the thirty-two output multiplexors will be selected at different times and there is only a one clock delay from the input data to the output data, the routing tag information will not be valid after it passes through the Switch Fabric. If the routing

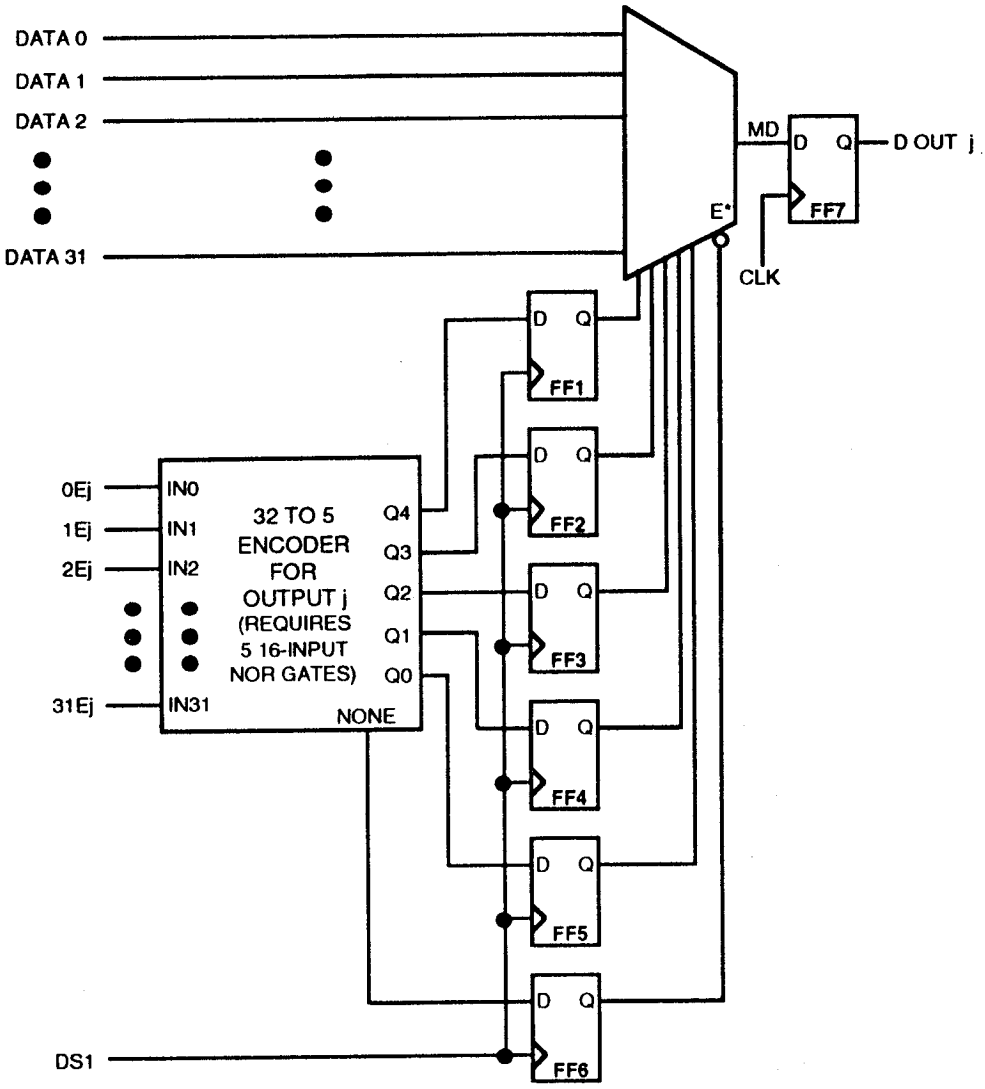

Figure 7-97. Output Selection Block Diagram

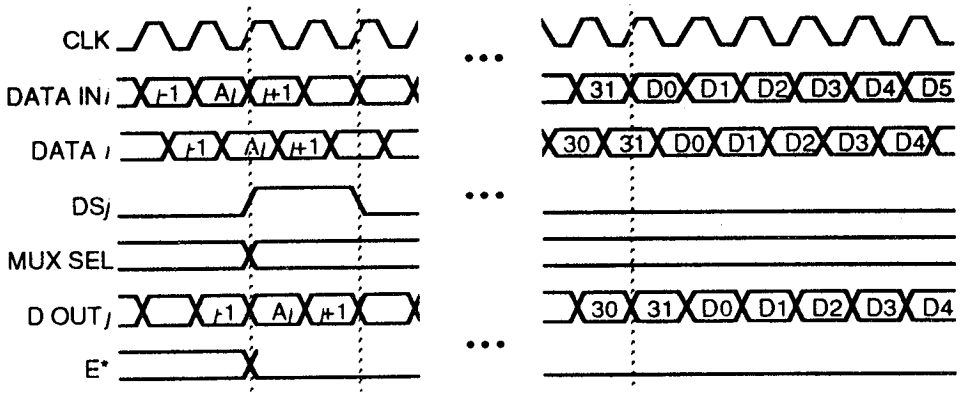

Figure 7-98. Output Selection Timing Diagram

tag is to be preserved, a minimum of 60 flip-flops would need to be added to each cell, increasing the power consumption dramatically.

7.2.2.4.3 Switch Fabric Characteristics

Multicast Crossbar ASIC Gate Count. The estimated gate count for the Multicast Crossbar ASIC is shown in Table 7-49.

Table 7-49. Multicast Crossbar ASIC Estimated Gate Count

FUNCTION	GATES/ FUNCTION	QUANTITY	TOTAL GATES
D FLIP-FLOP	12	105	1260
LOGIC GATE	8.5	315	2680
TOTAL + 20%			4700

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Multicast Crossbar ASIC Power Consumption. The estimated power consumption for the Multicast Crossbar ASIC is shown in Table 7-50.

Switch Fabric Power Consumption. The power consumption of the Switch Fabric is summarized in Table 7-51 for the three configurations. The 8 x 8, 16 x 16 and 32 x 32 configurations will require one, two, and four ASICs, respectively. Miscellaneous support circuitry is estimated at about 0.5 watts per port.

Table 7-50. Multicast Crossbar ASIC Estimated Power Consumption

FUNCTION	POWER/ DEVICE	QUANTITY	TOTAL POWER
INPUT I/O	6.4	36	0.230 W
OUTPUT I/O	7.81	8	0.062 W
D FLIP-FLOP	2.77	105	0.291 W
LOGIC GATE	1.5	315	0.473 W
TOTAL + 20%			1.3 W

Table 7-51. Switch Fabric Power Consumption

	8 x 8	16 x 16	32 x 32
CROSSBAR MATRIX	1.3 W	2.6 W	5.2 W
MISC SUPPORT	4 W	8 W	16 W
TOTAL	5.3 W	10.6 W	21.2 W

7.2.2.5 Multicast Crossbar-Based Switch Summary

7.2.2.5.1 Power Consumption

The power consumption of the Self-Routing Multicast Crossbar-Based Switch is summarized in Table 7-52 for the three switch configurations.

Table 7-52. Multicast Crossbar-Based Switch Power Consumption

	8 x 8	16 x 16	32 x 32
INPUT PORT	39.2 W	78.4 W	156.8 W
OUTPUT PORT	27.2 W	54.4 W	108.8 W
SWITCH FABRIC	5.3 W	10.6 W	21.2 W
TOTAL	71.7 W	143.4 W	286.8 W
TOTAL/PORT	9 W	9 W	9 W

7.2.2.5.2 Mass/Size

A circuit board of 30 square inches will support one Input Port and one Output Port. One board will be required for every port that the switch must service. All Switch Fabric configurations can also fit on one board of this size. In addition, one board will be required for the generation of common control and timing signals. This includes the following functions:

- Generation of Switch Fabric clocks.
- Generation and control of the Congestion Control Message.
- Generation of strobes to the Input Ports, Output Ports, and the Switch Fabric.

Therefore, the total number of modules required is 10, 18, and 34 for the 8 x 8, 16 x 16, and 32 x 32 switch configurations, respectively.

7.2.2.5.3 ASIC Design Complexity

There should be no problem in designing the two ASICs required for this approach. All ASICs have fairly low gate count and tolerable power consumption.

The design of these ASICs did not include added circuitry which may be required to implement a viable redundancy scheme. However, the addition of this circuitry is not expected to significantly increase the complexity of the ASIC design.

7.2.2.5.4 Fault Tolerance

The Multicast Crossbar-Based Network does not degrade as gracefully as the Multicast Banyan Network. Some of the shortcomings of this architectures are:

- The failure of one Input Port can affect other Input Ports. If the failed Input Port continuously sends the routing tag for the same Output Port, all other Input Ports will be prevented from accessing that Output Port.
- Only a few ASICs comprise the entire Switch Fabric. The complete loss of one ASIC will significantly decrease the operation of the Switch Fabric.

Some of the advantages of this architecture are:

- It is possible to implement a 1-for-N redundancy scheme for the Input Ports and Output Ports.
- Every input in the Multicast Crossbar ASIC has a dedicated data path to every output, therefore, a failure in one data path should have no affect on any other data path.
- A 1-for-N redundancy scheme can be implemented for the Switch Fabric ASICs. However, for and 8 x 8 Switch Fabric, N is equal to one. There are so few ASICs required in the Switch Fabric that a 1-for-1 redundancy scheme is also feasible.

The Congestion Control Message path is a ring type topology. If any Input Port fails, the entire ring becomes ineffective, causing switch to completely shut down. For this reason, it is imperative that this ring be redundant. Also, a detected failure in a given Input Port should cause the ring input to pass through unaffected to the ring output.

7.2.3 Point-To-Point Switch Fabric With Multicast Output Port

A block diagram of a Point-To-Point Switch Fabric using Multicast Output Ports is shown in Figure 7-99.

Serial data from the demodulators is routed through the switch to the modulators in the following steps:

- The serial packet data enters the respective Input Port where it is stored in a 128 packet FIFO type buffer.
- As each packet is stored in the packet buffer, the routing tag associated with the packet is obtain from the routing map memory then stored with the packet in the buffer.

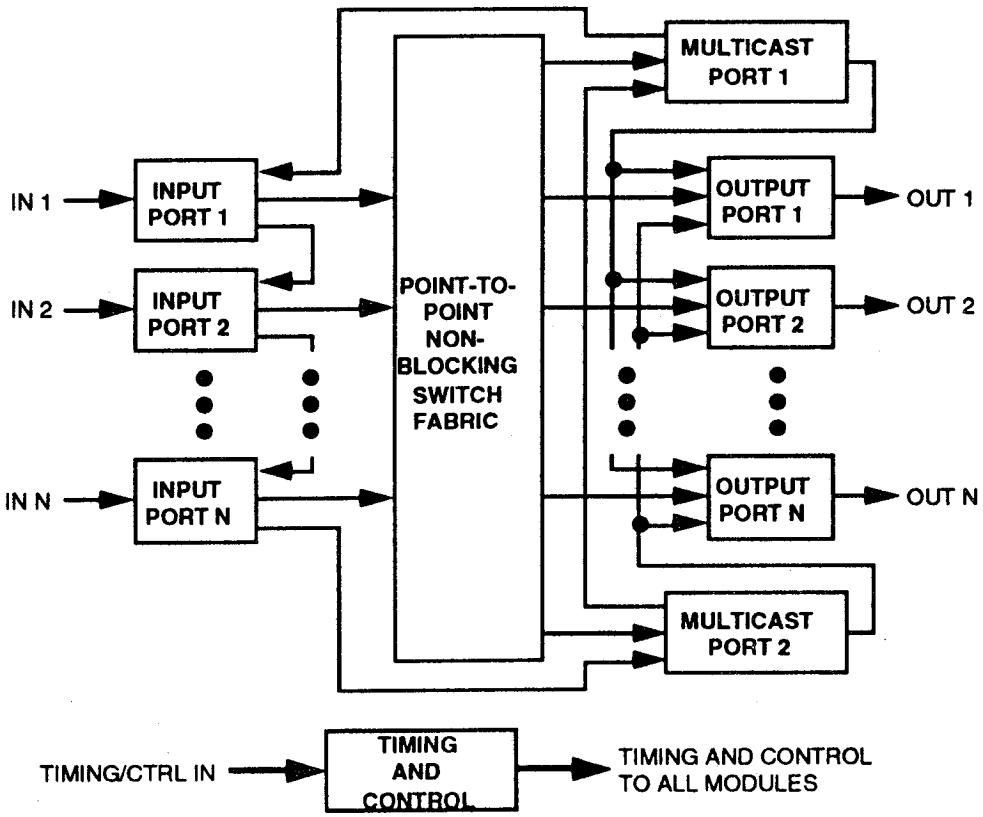

Figure 7-99. Point-To-Point Switch Fabric With Multicast Output Ports Block Diagram

- The Input Port performs Output Port congestion control to ensure that there are no blocked packets at the Output Port. The Input Ports convey their desired destination to each other by means of a serial daisy chained message every packet slot time. When an Input Port reserves an Output Port for the next packet transmission, no other Input Port further down the daisy chain may select the reserved Output Port. The four oldest in the packet buffer will be examined for transmission through the Switch Fabric to minimize head-of-line blocking.

- For point-to-point packets, the selected packet, if any, from each Input Port is sent to the Switch Fabric. The Switch Fabric will route packets from all Input Ports to the desired Output Ports as specified by the routing tag. The Switch Fabric will need to operate at a higher rate than the input serial data to compensate for the added bits for the routing tag and for the lost packet transfers through the Switch Fabric due to head-of-line blocking.

- For multicast packets, the selected packet will be sent to one of the Multicast Output Ports where it will be buffered in a 128 packet FIFO type buffer. The Multicast Output Port will then compete with the Input Ports for the desired Output Ports associated with the multicast packet. On the next packet slot period, the Multicast Output Port will transfer the packet to those Output

Ports which were reserved. Any desired Output Ports which were not reserved will be tried in subsequent slots until all have been serviced.

- The Output Port will accept serial data from the Switch Fabric then store it in a 128 packet FIFO type buffer. The serial packet data will be output to the modulator for downlink transmission.

7.2.3.1 Formats

7.2.3.1.1 Routing Tag Format

The format for the routing tag appended to each packet is shown in Figure 7-100.

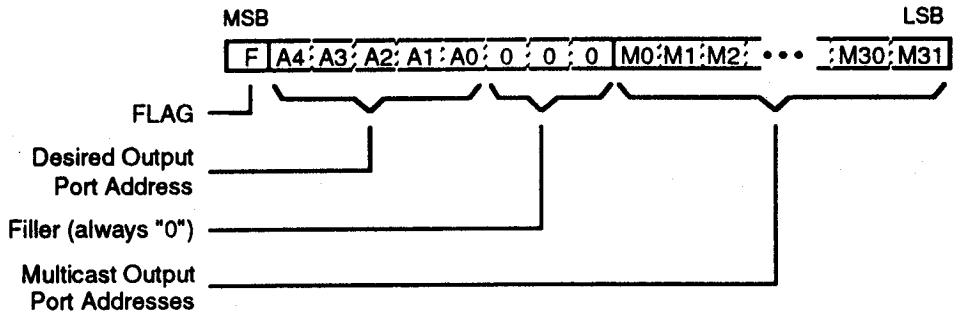

Figure 7-100. Multicast Output Port Scheme Routing Tag Format

The first bit is the FLAG bit which must be set for a valid packet. The next five bits are the address bits which the Switch Fabric uses to route the packet to the desired Output/Multicast Port. The last thirty-two bits are only used by the Multicast Output Ports and represent the desired destination Output Ports. Bit M0 is set if Output Port 0 is a desired destination, bit M1 for Output Port 1, bit M2 for Output Port 2, etc. This forty-one bit format will be used for all switch sizes so that the Output/Multicast Ports can easily determine the start position of the information bits.

7.2.3.1.2 Output Port Congestion Control Message Format

The format for the Output Port Congestion Control Message is identical to that used for the Sorted Banyan Network described in Section 7.1.1.1.2.

7.2.3.2 Input Port

The Input Port is very similar to that used for the Self-Routing Multicast Banyan Network described in Section 7.2.1.2 with the following exceptions:

- The Address Translator and the parallel to serial shift register R2 are not required. The address tag for the packet will be the contents of register R1. The difference in gate count and power consumption for this change is insignificant compared to the Multicast Banyan Network results.

- The addition required for the Output Port to process the data from the two Multicast Output Ports must be added. This will be explained in Section 7.2.3.3 where the Output Ports are presented. This addition is expected to add approximately 15 D-type flip-flops and 50 gates at 150 MHz.

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-53.

Table 7-53. Input Port ASIC Estimated Gate Count

FUNCTION	GATES/ FUNCTION	QUANTITY	TOTAL GATES
D FLIP-FLOP	7	340	2380
LOGIC GATE	1.75	800	1400
TOTAL + 20%			4500

The logic gate count includes inverters and 2 to 4 input nand and nor gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-54. The power consumption is not increased due to the added circuitry since the Input Port will not utilize these functions.

Table 7-54. Input Port ASIC Estimated Power Consumption

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	7	4	0.580 W
OUTPUT I/O	0	18	0	4	0.530 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	235	30	40	20	0.135 W
LOGIC GATE	370	240	100	20	0.080 W
TOTAL + 20%					1.9 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at the given frequency.

Input Port Power Characteristics. The Input Port will require one ASIC chip, two 8K by 8 bit memories, differential receivers and drivers, and minor support logic requiring an estimated ten square inches of board space. The total power consumption of the Input Port is summarized in Table 7-55.

Table 7-55. Input Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.9 W
RAM	1.5 W
I/O & SUPPORT	1.5 W
TOTAL	4.9 W

7.2.3.3 Output Port

7.2.3.3.1 ASIC Description (Addition of Input Data Selector)

The Output Port is identical to that used for the Multicast Banyan Network described in Section 7.2.1.3 except that the above Input Port ASIC will be used. The difference between this ASIC and that of the Multicast Banyan Network is that the new part must accommodate packets received not only from the Switch Fabric, but also from the two Multicast Output Ports. A block diagram of this function is shown in Figure 7-101. This circuitry is placed before the normal serial input from the Switch Fabric. D IN is the input from the Switch Fabric and D OUT feeds the input to the original Input Port ASIC.

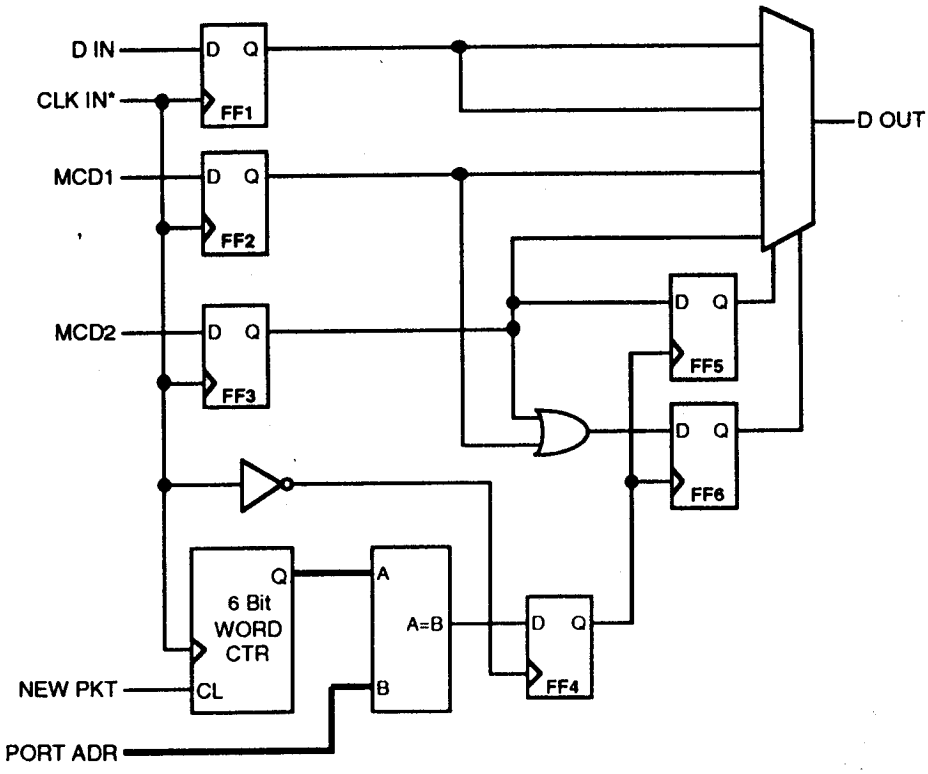

Figure 7-101. Input Data Selector Block Diagram

Flip-flop FF2 and FF3 mid-bit sample the data from the two Multicast Output Ports. Flip-flop FF1 mid-bit samples the data from the Switch Fabric to compensate for the delay of the Multicast Output Port data. The counter will be reset when a new packet arrives by NEW PKT. This counter represents the Output Port associated with the bit received by FF2 and FF3. The output of this counter is compared to the Output Port address. When these two values compare, the output of FF2 and FF3 contain the address bit for this Output Port. The output of the comparator is mid-bit sampled by FF4 to remove any glitches, then is used as the clock to FF5 and FF6 to sample the information on FF2 and FF3 as follows:

- If both address bits are low, indicating no multicast data for this port, the data from the Switch Fabric is passed to the output, D OUT.
- If only one of FF2 or FF3 is high, FF5 and FF6 will select the serial data associated with the active flip-flop to pass to D OUT.
- If both FF2 and FF3 are high, data received by MCD2 will be selected to pass to D OUT. This condition should never occur since the Output Port Congestion Control should prevent more than one source of data to any given destination Output Port.

The above circuitry is expected to add approximately 15 D-type flip-flops and 50 gates running at 150 MHz.

7.2.3.3.2 Output Port Characteristics

Input Port ASIC Gate Count. The estimated gate count for the Input Port ASIC is shown in Table 7-53. This is the same ASIC as used on the Input Port and will require approximately 4500 gates.

Input Port ASIC Power Consumption. The estimated power consumption for the Input Port ASIC is shown in Table 7-56. The values have been altered from the Input Port used in the Multicast Crossbar Network to include the new functions required of the ASIC.

Table 7-56. Input Port ASIC Estimated Power Consumption For Output Port

FUNCTION	QUANTITY AT 25 MHz	QUANTITY AT 50 MHz	QUANTITY AT 150 MHz	QUANTITY AT 200 MHz	TOTAL POWER
INPUT I/O	0	0	4	4	0.440 W
OUTPUT I/O	0	18	0	1	0.345 W
BIDIR I/O	0	16	0	0	0.250 W
D FLIP-FLOP	105	30	50	20	0.110 W
LOGIC GATE	300	240	85	0	0.065 W
TOTAL + 20%					1.5 W

The power consumption per gate used was 5.5 μW per MHz and per I/O was 1.25 mW per MHz. It was assumed that 20% of the logic gates transition at the given frequency and that 25% of the I/O transition at given frequency.

Output Port Power Characteristics. The Output Port will require one ASIC chip, one 4K by 16 bit memory, differential receivers and drivers, and minor support logic requiring an estimated seven square inches of board space. The total power consumption of the Output Port is summarized in Table 7-57.

Table 7-57. Output Port Estimated Power Consumption

FUNCTION	POWER
ASIC	1.5 W
RAM	1.0 W
I/O & SUPPORT	1.0 W
TOTAL	3.5 W

7.2.3.4 Multicast Output Port

The Multicast Output Port will perform the same functions as the Input Port. The only difference is that the input operates at the Switch Fabric data rate and the output operates at the user data rate. However, the characteristics for this module should be identical to the that of the Input Port.

7.2.3.5 Switch Fabric

The point-to-point Switch Fabric selected is the Sorted Banyan Network since the Crossbar Network can already accommodate multicast packets. The summary for this Switch Fabric is repeated here.

ASIC Count. The number of ASICs required for each of the three switch fabric configurations is shown in Table 7-58.

Table 7-58. Switch Fabric ASIC Count

ASIC COUNT	8 x 8	16 x 16	32 x 32
BANYAN	1	4	8
BATCHER	2	8	20
TOTAL	3	12	28

Switch Fabric Power Consumption. The power consumption of the Switch Fabric is summarized in Table 7-59 for the three configurations. Miscellaneous support circuitry is estimated at about 0.5 watts per port.

Table 7-59. Switch Fabric Power Consumption

	8 x 8	16 x 16	32 x 32
BANYAN MATRIX	.7 W	2.8 W	5.6 W
BATCHER MATRIX	1.4 W	5.6 W	14 W
MISC SUPPORT	4 W	8 W	16 W
TOTAL	6.1 W	16.4 W	35.6 W

7.2.3.6 Multicast Output Port Switch Summary

7.2.3.6.1 Power Consumption

The power consumption of the Self-Routing Multicast Crossbar-Based Switch is summarized in Table 7-60 for the three switch configurations.

Table 7-60. Multicast Crossbar-Based Switch Power Consumption

	8 x 8	16 x 16	32 x 32
INPUT PORT	39.2 W	78.4 W	156.8 W
OUTPUT PORT	28 W	56 W	112 W
MULITCAST PORT	4.9 W	9.8 W	9.8 W
SWITCH FABRIC	6.1 W	16.4 W	35.6 W
TOTAL	78.2 W	160.6 W	314.2 W
TOTAL/PORT	11.2 W	11.5 W	10.5 W

7.2.3.6.2 Mass/Size

A circuit board of 30 square inches will support one Input Port and one Output Port. One board will be required for every port that the switch must service. In addition one of these boards will be required for the 8 x 8 Multicast Output Ports and two will be require for the 16 x 16 and 32 x 32 configurations. The 8 x 8 and 16 x 16 Switch Fabric configurations can fit on one board of this size with the 32 x 32 configuration requiring two boards. In addition, one board will be required for the generation of common control and timing signals. This includes the following functions:

- Generation of Switch Fabric clocks.
- Generation and control of the Congestion Control Message.
- Generation of strobes to the Input Ports, Output Ports, and the Switch Fabric.

Therefore, the total number of modules required is 10, 18, and 35 for the 7 x 7, 14 x 14, and 30 x 30 switch configurations, respectively.

7.2.3.6.3 ASIC Design Complexity

There should be no problem in designing the three ASICs required for this approach. All ASICs have fairly low gate count and tolerable power consumption. A disadvantage to this approach is that three ASIC designs (Input Port, Banyan, and Batcher) are required as opposed to only two for other approaches.

The design of these ASICs did not include added circuitry which may be required to implement a viable redundancy scheme. However, the addition of this circuitry is not expected to significantly increase the complexity of the ASIC design.

7.2.3.6.4 Fault Tolerance

The fault tolerance of this approach is relatively good since the Sorted Banyan Switch Fabric is used. This network will have the advantages of the Sorted Banyan Switch which include

- A failure of an Input Port will not affect any other Input Port.
- A failure of an Output Port will not affect any other Output Port.
- A failure of a cell in the Switch Fabric will only affect 1/N of the paths through the fabric where N is the size of the Switch Fabric.

In addition the following advantages are obtained from the Multicast Output Ports:

- A failure of one Multicast Output Port can be overcome by the other Mulitcast Output Port.
- A failure in the Switch Fabric can be overcome by routing the packet through the Multicast Output Port. Since no single cell can affect both Multicast Output Ports, all affected paths can be routed to the Output Ports via the good Multicast Output Port(s).

The Congestion Control Message path is a ring type topology. If any Input Port fails, the entire ring becomes ineffective, causing switch to completely shut down. For this reason, it is imperative that this ring be redundant. Also, a detected failure in a given Input Port should cause the ring input to pass through unaffected to the ring output.

7.2.4 Multicast Switch Networks Summary

Table 7-61 below summarizes the various point-to-point switch approaches.

Table 7-61. Summary Of Multicast Architectures

		MULTICAST BANYAN	MULTICAST CROSSBAR	MULTICAST OUTPUT PORT
Power Consumption (Total Switch / Per Port, in watts)	8 x 8	74/9.3	72/9	78/11.2
	16 x 16	154/9.6	143/9	161/11.5
	32 x 32	324/10.1	287/9	314/10.5
Switch Fabric ASIC Count	8 x 8	5	1	3
	16 x 16	18	2	12
	32 x 32	58	4	28
ASIC Designs	Number Required	3	2	3
	Gate Count (in K)	4.3/2.5/2.4	4.3/4.7	4.5/2.4/2.4
	Power (in watts)	1.9/0.85/0.7	1.9/1.3	1.9/0.7/0.7
Fault Tolerance	Graceful Degradation	Good	Moderate	Good
	Redundancy	Difficult	Moderate	Moderate

The Self-Routing Multicast Crossbar-Based Network is chosen as the optimal architecture for the current requirements. It has the lowest power consumption per port and the lowest ASIC count for the Switch Fabric, especially for the larger switches. Only two ASIC designs are required, both of relatively low complexity.

The major disadvantage of the Multicast Banyan Network is the high ASIC count for larger switches, the requirement for three ASIC designs, and the complexity of implementing a 1-for-N redundancy scheme for the Switch Fabric.

The major disadvantage of the Point-To-Point Switch Fabric with Multicast Output Ports is the extremely high power cost paid for smaller switches. Even for the 32 x 32 configuration, the power consumption per port is over 10 watts. An advantage of this approach is that it has very attractive fault tolerance and simple, built-in redundancy for a wide class of failures.

7.3 Fault tolerance

A single point failure in any of the switch networks will cause a degradation in performance of the network. The severity of this degradation will depend on the type of failure and the switch architecture selected. For each of the switch architectures presented in the previous sections, a brief description was given for the fault tolerant characteristics of the switch.

A method to decrease the effects of a failure is to add redundancy to the switch network so that the functions performed by the failed area can be restored by a redundant replacement. Several redundancy schemes are possible. The optimal scheme for a particular applications will be determined by the two conflicting requirements of performance versus cost, i.e., the better the redundancy coverage, the higher the cost to power, space, and complexity of the system. Since added circuitry is required for the detection of the failure and the switchover mechanism, it must be recognized that this added circuitry will generated a new class of failures which will affect the switch performance.

This section will discuss 1-for-1 redundancy approaches for the switch as well as 1-for-N redundancy for modules within the switch. The following topics will be presented:

- 1-for-1 Switch Redundancy
- 1-for-N Input Port Redundancy
- 1-for-N Switch Fabric Path Redundancy
- 1-for-N Switch Fabric ASIC Redundancy
- 1-for-N Output Port Redundancy

7.3.1 Switch Network Redundancy

Conceptually, the simplest implementation for redundancy is a 1-for-1 redundancy scheme for the entire Switch Network. In this approach, any failure detected in the online network will cause the redundant network to switch online. Of course, if a failure exists in both the online and redundant network, switch degradation would result. A block diagram of this approach is shown in Figure 7-102.

Each input to the Switch Network is sent to both the online switch (Switch A) and the offline redundant switch (Switch B). The outputs from Switch A and Switch B are sent to a multiplexor which is under the control of the signal processor. The output of the multiplexors provide the output of the Switch Network.

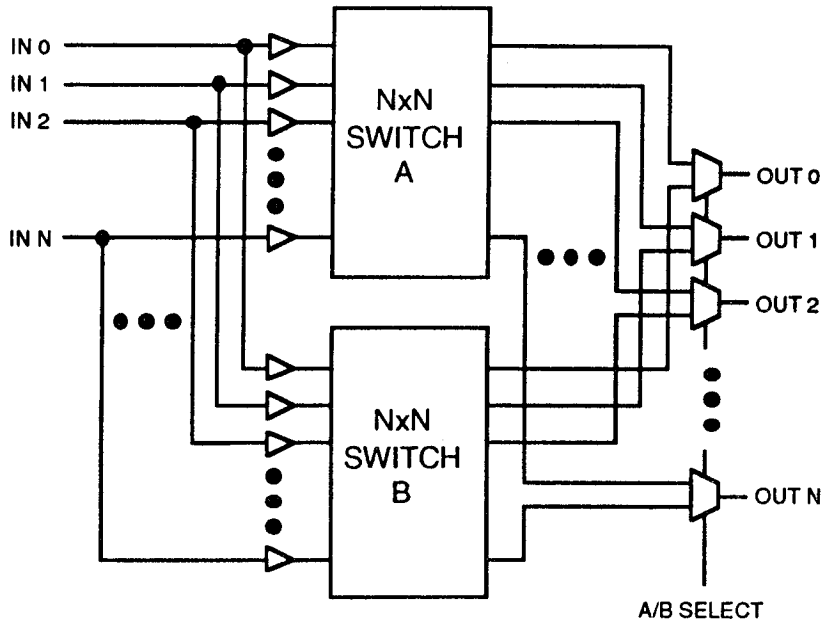

Figure 7-102. 1-For-1 Switch Network Redundancy Block Diagram

7.3.1.1 Fault Detection

Fault detection in this scheme would be very straight forward. Each of the output multiplexors will compare their two inputs. If they do not match, then a failure exists in either the online switch or the redundant switch. Once a problem is detected, the signal processor will initiate a fault isolation algorithm. Fault detection will therefore occur on a non-interference basis and require no user bandwidth through the online switch.

7.3.1.2 Fault Isolation

Fault isolation can also be performed on a non-interference basis. The multiplexor which detected the failure will report it to the signal processor. From the information received from the multiplexor, the signal processor will be able to determine the source and destination ports involved in the failure. The signal processor will initiate the fault isolation procedure as follows:

- The multiplexor will be placed in the fault isolation mode. In this mode, the online and offline inputs do not have to match.
- The fault isolation will involve only the offline redundant switch so that the online switch operation will be unaffected.
- The signal processor will cause the affected Input Port of the offline redundant switch to send a test packet to the affected Output Port.

- Since the contents of this packet are known, the multiplexor will be able to verify if the packet is received correctly.
- This signal processor will determine the criteria for an operational switch. The signal processor may wish to send multiple packets for verification or it may wish to verify all paths are operational before declaring that the offline redundant switch is operational.

7.3.1.3 Redundancy Switching

If the signal processor determines that the offline redundant switch has a failure, it will not put it online.

If the signal processor has determined that the offline redundant switch is operational, it will synchronize the offline switch to the online switch then force it online. The previously online switch will now be offline so that it may be checked for failures.

The signal processor can also decide which switch to put online when both have failures. Since it can exhaustively test all paths in the switch network, it will know which switch has less problems and can force it online.

7.3.1.4 Summary

The 1-for-1 Switch Network redundancy has several advantages.

- Since there is a complete 1-for-1 redundancy, the offline switch can be in hot standby ready to be put online at any time.
- Fault detection is simply implemented.
- Fault isolation can be performed on the offline switch so that online traffic will not be affected.

The disadvantages for this approach are as follows.

- The entire Switch Network must be duplicated.
- The switch will incur more than double the power consumption. The offline redundant switch can be powered down when not in use but this would elimination the current approach for fault detection.
- The switch will incur more than double the volume of a non-redundant network.

7.3.2 Input Port Redundancy (1-for-N)

A complete 1-for-1 Switch Network redundancy is very costly in space and power. An alternative is to provide a 1-for-N redundant scheme for modules within the switch.

The Input Ports in the Switch Network can be configured in this manner. A block diagram of this approach is shown in Figure 7-103.

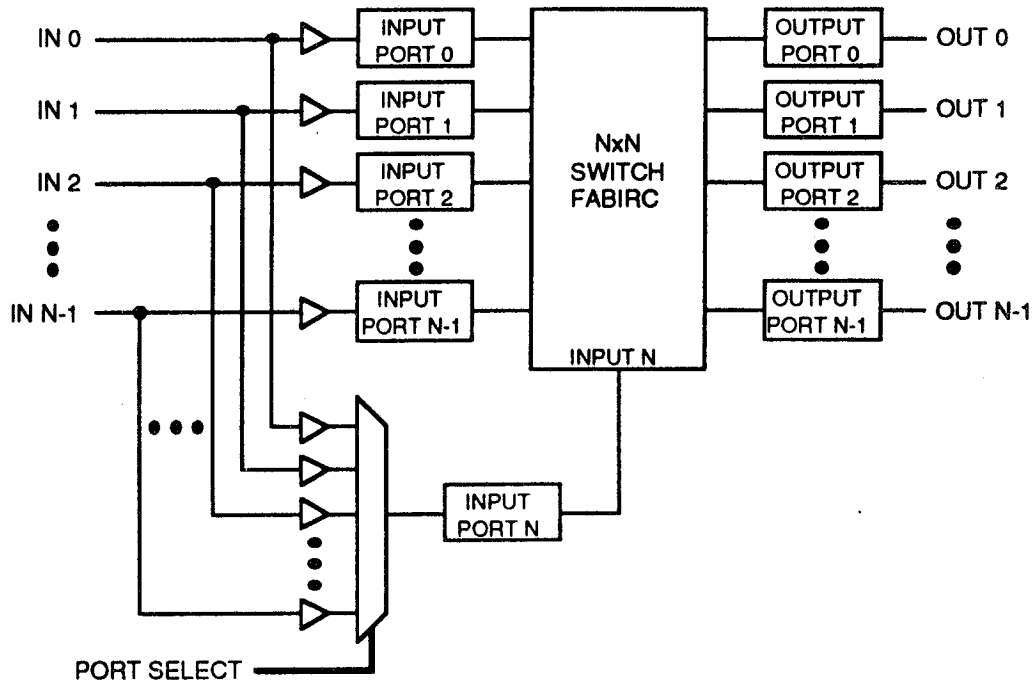

Figure 7-103. 1-For-N Input Port Redundancy Block Diagram

Each input the Switch Network is sent to their respective Input Port and to the Input Port multiplexor. When a failure of an Input Port is detected, the redundant port, Input Port N, will take over its operation. The Input Port multiplexor will select the input to the bad port to pass to Input Port N. Input Port N will be updated with all mapping information for the failed port. It will then enable itself online, effectively replacing the failed port.

7.3.2.1 Fault Detection

Fault detection for the failed Input Port will be more complicated than that for the 1-for-1 Switch Network Redundancy. One possible approach is to compare the online Input Ports to the redundant Input Port in the following manner:

- The signal processor will check one Input Port at a time. This check will be performed on a non-interference basis to the user traffic.
- The Input Port multiplexor will be set to select the input for the port under test.
- The redundant Input Port will be synchronized to the port under test so that they will be outputting the same data for a given packet slot period.

- The address tag will be modified to send all packets from the port under test to output N (not used output since the redundant Input Port is using input N) in addition to the normally selected output(s).
- A comparison circuit will determine if output N of the Switch Fabric is the same as the output of the redundant Input Port. If not, a failure exists in the port under test or the redundant Input Port. If a failure is not detected, the next Input Port will be checked.
- This check will be performed continuously as part of background testing.

7.3.2.2 Fault Isolation

Once a failure has been detected, it must be isolated to the failed module. The modules which could have caused the above fault detection are:

- The Input Port under test.
- The redundant Input Port.
- The Switch Fabric which routes the port under test to the Switch Fabric output.

Redundant Input Port. The redundant Input Port can be checked by sending a known test packet through it and verifying that its output is correct.

Input Port Under Test. Once the redundant port has been verified, it can be swapped with the suspected Input Port. This Input Port will now be offline and can be checked by sending a known test packet through it then verifying that its output is correct.

Switch Fabric. If both of the Input Ports verify properly and the fault is still being detected, the problem must be in the Switch Fabric. Switch Fabric redundancy will be presented in Sections 7.3.3 and 7.3.4.

7.3.2.3 Redundancy Switching

If the redundant Input Port has a problem, it will not be put online.

If the redundant port is working properly, it will be synchronized to the port under test then put online to replace this port. The port under test can now be checked on a non-interference basis to ongoing user traffic. If this port is determined to be failure free, it will be put back online.

If both the redundant port and the port under test are determined to failure free, the Switch Fabric will be checked. Switch Fabric redundancy will be presented in Sections 7.3.3 and 7.3.4.

7.3.2.4 Summary

The 1-for-N Input Port redundancy has the following advantages.

- Only one redundant port required for any size switch. This will greatly reduce the space and power requirements over the 1-for-1 Switch Network redundancy approach.
- If more protection is desired, the scheme is easily expanded to a i-for-N configuration. The cost penalty here, in addition to the space and power of the added redundancy port(s), is that each redundant port will require an input to the Switch Fabric reducing the operational size of the switch.
- Fault detection possible in background mode. Online user traffic will not be affected.
- Fault detection easily implemented.

The disadvantages for this approach are as follows.

- More complex fault isolation algorithm than in the 1-for-1 Switch Network redundancy approach.
- Operational size of the Switch Network is reduced by each redundant port added.
- More complex interconnect required.

7.3.3 Switch Fabric Path Redundancy (1-for-N)

A complete 1-for-1 Switch Network redundancy is very costly in space and power. An alternative is to provide a 1-for-N redundant scheme for modules within the switch. The Switch Fabric paths in the Switch Network can be configured in this manner. A block diagram of this approach is shown in Figure 7-104.

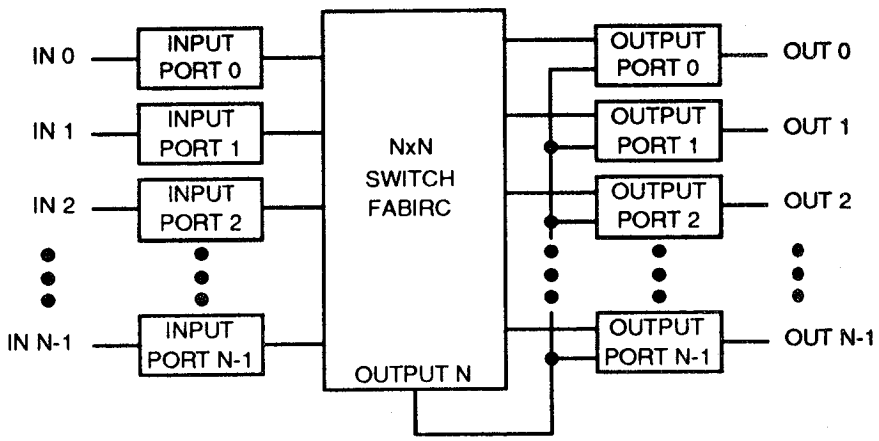

Figure 7-104. 1-For-N Switch Fabric Path Redundancy Block Diagram

When a failed path through the Switch Fabric is found, that path is re-routed from the affected Input Port to Output N of the Switch Fabric. The affected Output Port is instructed to enable the reception of data from Output N as well as its normal output. Output N will have priority over the normal output so that if packets arrive the the Output Port simultaneously, Output N will be selected.

7.3.3.1 Fault Detection

Fault detection can be implemented in the Output Ports on a non-interference basis to online user traffic. The packets from a given Input Port will be routed to Output N of the Switch Fabric in addition to the normal destination(s). Each Output Port will compare their normal input from the Switch Fabric to Output N. These results will be given to the signal processor. Since the signal processor can determine which were the destination port(s) it will be able to determine if there is a problem with the Switch Fabric. The Input Port to be routed to Output N will be changed periodically so that the coverage for the Switch Fabric paths will be as complete as possible.

7.3.3.2 Fault Isolation

The signal processor will be primarily responsible for the determination of the faulted module. It will accept test results from all the Output Ports which should contain enough information to determine the failed part. To eliminate the paths to Output N as failed, a test packet can be sent to Output N from the Input Ports and verified in the Output Ports. Once Output N has been verified as failure free, the suspect path can be easily determined.

7.3.3.3 Redundancy Switching

When the failed path has been determined, it will be a simple task to bypass it. The signal processor will change the mapping of the affected Input Port such that Output N will replace the affected Output Port in the routing tag field. The affected Output Port will be instructed to enable its input for Switch Fabric Output N. When a packet from the affected Input Port is destined for the affected Output Port, it will be routed through Output N, bypassing the failed path in the Switch Fabric. The Output Port will give priority to Output N in case of parallel reception from the normal Switch Fabric path.

If desired, multiple redundant paths may be used. Each redundant path will require an additional output from the Switch Fabric and the ability for the Output Port to accept an additional redundant input.

7.3.3.4 Summary

Advantages for the 1-for-N Switch Fabric Path redundancy approach are listed below.

- Only one output of the Switch Fabric needed to implement this approach. Some circuitry additions required in the Output Ports.

- The output of the Switch Fabric used can be the same as used for the 1-for-N Input Port redundancy scheme.
- Fault detection performed on non-interference basis to online traffic.
- Fault isolation easily implemented.

Disadvantages to this approach are as follows.

- More complex fault isolation algorithm than in the 1-for-1 Switch Network redundancy approach.
- Operational size of the Switch Network is reduced by each redundant path added.
- More complex interconnect required.
- Failure in the Switch Fabric will usually affect the entire part causing more than just one path to fail.

7.3.4 Switch Fabric ASIC Redundancy (1-for-N)

A complete 1-for-1 Switch Network redundancy is very costly in space and power. An alternative is to provide a 1-for-N redundant scheme for modules within the switch. The Switch Fabric ASICs in the Crossbar-Based Switch Network can be configured in this manner. The Banyan/Batcher Networks are not conducive to this approach. A block diagram is shown in Figure 7-105.

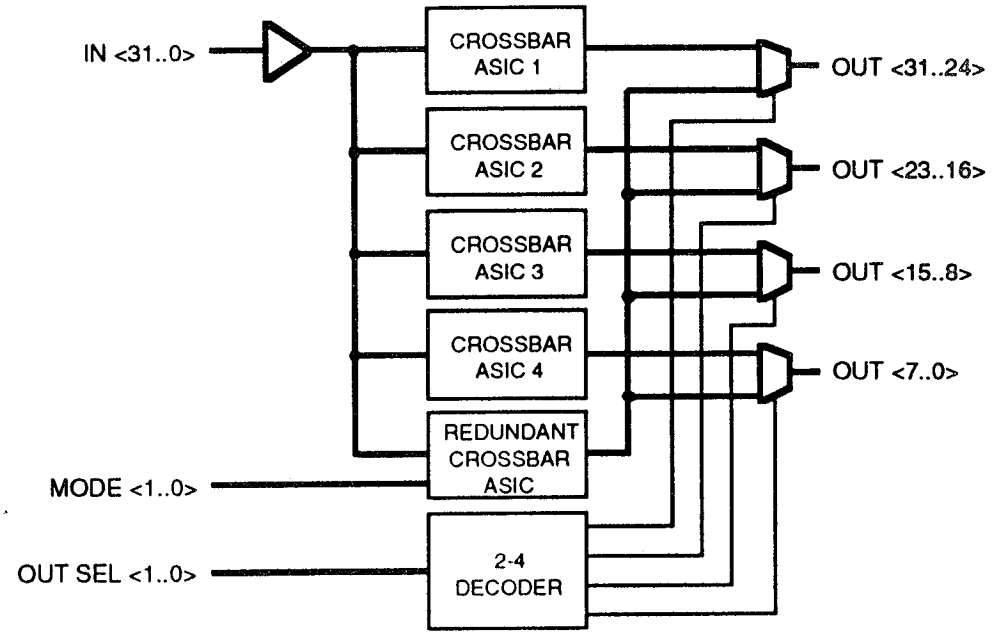

Figure 7-105. 1-For-N Switch Fabric ASICs Redundancy Block Diagram

A redundant Crossbar ASIC is added to the Switch Fabric. It accepts the same inputs as all other ASICs. The output of the redundant ASIC can be multiplexed to replace the outputs of any of the operational ASICs. When one of the operational ASICs have failed, the redundant ASIC will be set to its mode by MODE <1..0> and the associated multiplexor will select the redundant ASIC as its input. In this manner, the redundant ASIC has completely taken over the function of the failed ASIC.

7.3.4.1 Fault Detection

The redundant ASIC will be used to perform the fault detection in the Switch Fabric. The test will be performed on all online ASICs, one at a time, without affecting online traffic. The redundant ASIC will be programmed with the same mode input as the ASIC under test. The output of these two ASICs should now be identical. The corresponding multiplexor will perform a comparison on the two outputs; if they do not compare, a failure has been detected.

7.3.4.2 Fault Isolation

Once a failure has been detected, it must be isolated to the online ASIC or the redundant ASIC. The redundant ASIC will be checked first by sending a known packet through it and verifying this packet at the output multiplexors. If this ASIC is determined to failure free, the problem must exist in the online ASIC.

7.3.4.3 Redundancy Switching

When the fault has been isolated to the online ASIC, the redundant ASIC must be placed online to replace the failed part. The mode bits are set to the values for the failed ASIC. Since there is only a fixed delay in the Crossbar ASIC, the redundant ASIC can be put online immediately on the next packet boundary. Once the failed ASIC is offline, it may be further tested to verify its condition.

7.3.4.4 Summary

Advantages for the Switch Fabric ASIC redundancy are listed below.

- Only 1 ASIC required for any switch size.
- Fault detection possible in background mode.
- Fault detection and isolation easily implemented.
- Whole ASIC is replaced on a failure. It is likely that multiple paths have been destroyed rather than just one path.
- Redundancy circuitry limited to Switch Fabric module.

Disadvantages to this approach are as follows.

- For 8x8 Switch Fabric, N is equal to 1. This is a significant overhead for this size switch.
- Can only be used in Crossbar approach. This method is not practical for any of the Banyan/Batcher Networks.

7.3.5 Output Port Redundancy (1-for-N)

A complete 1-for-1 Switch Network redundancy is very costly in space and power. An alternative is to provide a 1-for-N redundant scheme for modules within the switch. The Output Ports in the Switch Network can be configured in this manner. A block diagram of this approach is shown in Figure 7-106.

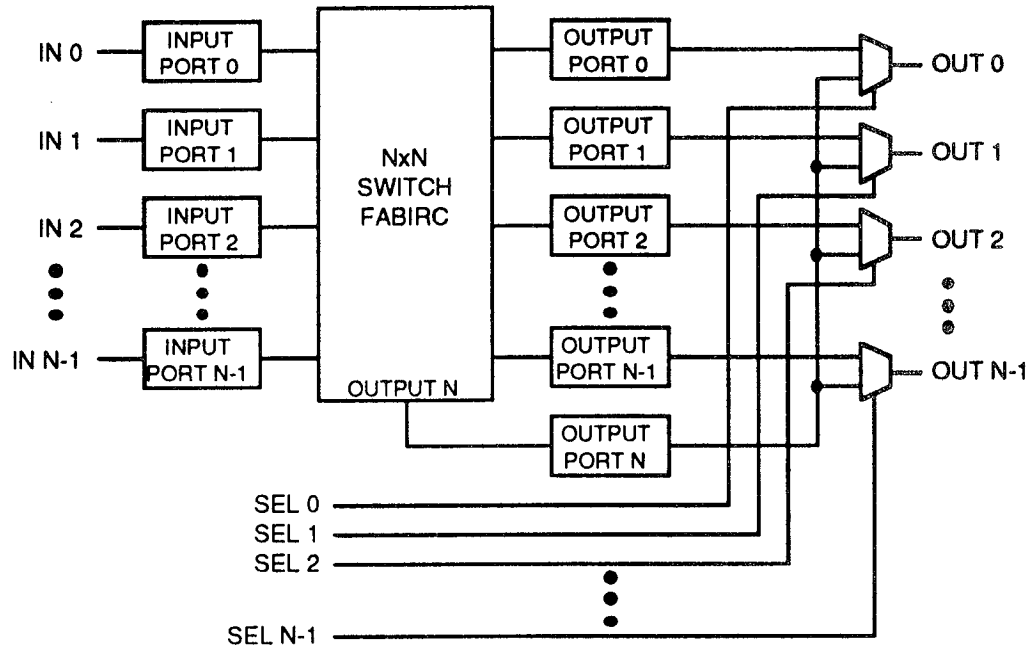

Figure 7-106. 1-For-N Output Port Redundancy Block Diagram

A redundant Output Port is added to the Switch Network. The output of this port can replace the output of any other Output Port by the output multiplexors. If a given Output Port is determined to contain a fault, its destination address in the Input Port mappings will be replace with Output N, the redundant port address. The selected multiplexor will enable the redundant port causing the failed Output Port to be bypassed.

7.3.5.1 Fault Detection

The fault detection scheme used will check the Output Ports on a non-interference basis to the online user traffic. The Input Ports will selectively send packets to the redundant Output Port. The output of this redundant port will be compared with the other Output Ports by the output multiplexors. This information is given to the signal processor.

Since the signal processor knows the destination Output Port(s) for a given Input Port, it will be able to determine which multiplexor should show valid compariosns. If a valid compariosn is not found from an expected multiplexor, a failure has occurred in either the Input Port, the Switch Fabric, the associated Output Port, or the redundant Output Port.

7.3.5.2 Fault Isolation

When a fault has been detected, the Input Port and Switch Fabric will be checked first by one of the methods described in the previous sections. If both of these are found to be fault free, either the associated Output Port or the redundant port has failed. The redundant Output Port can be verified by sending a known packet from one of the Input Ports to Output N. This packet can then be checked by the output multiplexors to determine if the redundant port is good. If it is determined that the redundant port is failure free, the affected Output Port is now suspect and will be switched offline.

7.3.5.3 Redundancy Switching

An Output Port will be replaced by the redundant Output Port by the following procedure.

- The signal processor will cause the affected output multiplexor to select the redundant Output Port.
- The map entries in all the Input Ports will be changed to reflect the new destination for packets for the given Switch Network output.
- At this time, the affected Output Port is now offline and can be checked to verify its status.

7.3.5.4 Summary

Advantages for a 1-for-N Output Port redundancy are listed below.

- Only one redundant Output Port required for any size switch.
- Fault detection possible in background mode.
- Fault detection and isolation easily implemented.

Disadvantages for this approach are listed below.

- More complex fault isolation algorithm than in the 1-for-1 Switch Network redundancy approach.
- Operational size of the Switch Network is reduced by each redundant Output Port added.
- More complex interconnect required.

Section 8

Conclusion

Through this study, it is reaffirmed that satellites can play a vital role in future broadband telecommunications services. Potential satellite applications include high-speed video distribution, B-ISDN trunking, private/business networks, thin-route and emergency communications, and NASA science networking. The major drivers for achieving the effectiveness of satellite communications are flexibility, high performance (i.e., a low bit error rate and high availability), and low user service cost. On-board baseband processing in conjunction with multiple spot beam operation potentially provides all these features. There are, however, a number of critical technology areas that will require further study and development. The following paragraphs summarize the study results and identify some of the system design issues and critical technologies for future research and development consideration.

The satellite transmission system can support high-speed transmission at a bit rate between 155 Mbit/s and 1.24 Gbit/s using an antenna size of between 2.4 m (155 Mbit/s at Ka-band) and 12 m (1.24 Gbit/s at Ku-band). These high-bit rates are primarily intended for use in trunking and high-speed circuit switched applications, and efficient bandwidth utilization, low bit error rates (e.g. 10^{-11}), and high availability are often demanded by this type of services. A simpler on-board switching payload, such as SS-TDMA and SS-FDMA (in RF or baseband), may be used for spot beam interconnection. The critical technologies required for high-speed transmission include bandwidth efficient modulation (e.g. 8-PSK and 16 QAM), concatenated coding (e.g. a combination of a high-rate convolutional code and Reed-Solomon code), satellite beam forming networks, and Ka-band earth station technology (antenna, HPA, and baseband processor). A bit rate of up to 1.24 Gbit/s must be supported by these subsystems.

The majority of future B-ISDN users will likely require bit rates between 1.544 Mbit/s and 155 Mbit/s and flexible interconnection among a large number of user terminals. In this environment, user earth station cost becomes of prime importance in designing a satellite system. Flexibility and efficient bandwidth utilization can be achieved with an on-board fast packet switch to provide an integrated circuit and packet switched service. Uplink access and transmission can be optimized to a group of users according to their traffic requirements, and downlink may use a single carrier TDM transmission to maximize the use of spacecraft power. The fast packet switch can also provide integrated operation of B-ISDN and non-B-ISDN traffic, including N-ISDN, frame relay, X.25/X.75, LANs, MANs, and SMDS. Critical technologies for low/medium speed

applications include an on-board fast packet switch, low-power multicarrier demodulators (MCDs), and other various multicarrier processing subsystems (FEC decoder/encoder, scrambler/descrambler, TDMA receive processor, etc.). The Ka-band earth station technology, such as a 10 watt SSPA, a low-cost antenna, and baseband processor, is extremely critical for a wide spread use of low-cost VSAT earth stations with an antenna size of 0.8 m to 1.2 m.

There are several architectural options for implementing an on-board fast packet switch. For low to medium capacity applications (1.5 - 2 Gbit/s), shard memory or shard medium switch structures (e.g. common memory, distributed memory with a parallel bus, and a fiber optic ring) seem most attractive. For a high capacity system (above 2 Gbit/s), multistage switching networks and a fiber optic ring can be used. Based on the detailed designs of multicast switches presented in this report, the per-port power consumption is about 9 watts using the current technology and is likely to be reduced to around 3 - 4 watts in the near future. This will translate into a total switching subsystem power requirement of 48 - 64 watts for a 10-Gbit system with each port operating at 622 Mbit/s. This estimate, however, does not include additional on-board hardware required for demodulation/remodulation, TDMA or TDM synchronization, FEC decoding/encoding, additional data memories (the above value assumes a storage of 128 512-bit packets), bit interleaving processing (if needed), and an on-board network controller. If all the necessary subsystems are included, the on-board baseband processor may require several hundred watts of power. This illustrates that a power requirement for a baseband switching function is relatively minor compared with the total processor power.

The single most important issue in the design of a fast packet switched system is probably flow and/or congestion control to prevent frequent buffer overflow at the satellite. This issue is also related to the estimation of a satellite buffer size to absorb traffic fluctuations for different downlink beams. The difficulty of developing an efficient control mechanism arises from the problem of characterizing user traffic and implementing a reactive control scheme due to a long satellite propagation delay. An attempt to estimate a proper buffer size has been made in this report for statistically well defined traffic models. The analytical and simulation results indicate that a buffer size of 50 to 100 packets may be needed for each link operating at 155 Mbit/s. The report also suggests several control schemes that may be effectively used for various network scenarios. Some of these techniques have been successfully tested in the past. Nevertheless, further study is needed to completely understand the problem and to develop efficient control procedures.

Appendix A

RF Transmission System Design

1 Satellite B-ISDN Transmission Issues

The support of B-ISDN in a satellite based system places severe requirements on the satellite transmission portion of the link due to the high bit-rates involved. Bandwidth availability becomes a very important issue, coupled with the inherently power limited nature of satellite links. In many applications, such as private networking, the size of the earth station has to be kept as small as possible, approaching current VSAT sizes. Further, B-ISDN performance requirements as they relate to BER and availability, has to be met. This section explores the transmission link design for satellite B-ISDN.

1.1 Transmission Bit Rates

A variety of B-ISDN services can be supported by a satellite network, with widely varying information rates and degree of burstiness. The B-ISDN interface rate is at 155.52 Mbit/s but the actual information rates may be as low as 64 kbit/s. Hence, the actual transmission rates over the satellite link may be lower than the B-ISDN interface rate. Typical transmission bit rates could range from 10 Mbit/s on the low end to 155 Mbit/s for user data. Users with lower traffic requirements would share a low rate carrier either by multiplexing their traffic at the earth station or by the use of time division multiple access. On the other hand, higher bit rates for such services as trunking, super computer networking, and certain science data distribution applications could reach 1.24 Gbit/s and may go higher.

1.2 Frequency Bands and General Satellite Characteristics

The large bandwidths associated with the transmission of B-ISDN carriers necessitates the utilization of wider transponder bandwidths than what is currently being used. They also necessitate the utilization of the higher frequency bands such as Ka-band (30/20 GHz) and possibly 50/40 GHz with their wider bandwidth allocations. An added benefit of going to higher frequencies is that the size of the earth terminal may be made smaller, however it is not clear whether this advantage can be realized in the presence of severe rain fade attenuations. The use of high power and high gain satellite repeaters becomes essential in order to meet the performance requirements at the B-

ISDN transmission rates. Spot beam coverage and on-board processing which includes on-board FEC add to the achievable improvements in link design. Another advantage of on-board processing is the separate optimization of the uplink and downlink which results in the most efficient use of earth station and satellite power.

1.3 Transmission Link Design Objective and Methodology

The transmission link design presented in this section attempts to determine typical bit rates that can be supported by various earth station sizes (antenna and HPA) primarily at Ku-band and Ka-band. The use of the 50/40 GHz band which is currently not being utilized for commercial communications is also considered since it may become necessary to accommodate high traffic volume. The approach used in the link design is to first characterize typical rain attenuation impairments at the transmission bands of interest. This is done using the Global Model [A-1] to determine the percent of time specific attenuations are exceeded in different rain climate regions. The second step is to use the representative data obtained from the Global Model in deriving uplink and downlink link budgets using representative bit rates, earth station sizes, and availability requirements.

1.4 Link Budget Limitations

Due to the large number of variations that could be examined in such an analysis, only a few representative cases are considered. Hence, the transmission link design is intended only to a give a feel for what can be achieved rather than being a comprehensive link design that covers each and every possible combination of satellite and earth station parameters. For the same reason, only one representative set of satellite parameters is used for each band, and the link budgets are simplified by omitting many implementation specific entries such as feed losses, pointing losses, inter-system and intra-system interference, modem losses, and the like. This is not to say that these losses are nonexistent, and indeed they are accounted for by the addition of 3 dB link degradation on each of the uplink and the downlink.

2 Global Model

Rain attenuation at Ku band and higher frequencies presents a severe impairment to the transmission link. An attempt is made to quantify the magnitude and probability of occurrence of rain attenuation for the three frequency bands of interest using the Global Rain Attenuation Model. The Global Model is a prediction tool to predict rain attenuation associated with a given exceedance time percentage for earth station location and certain link parameters. It uses the rain rate statistics for the earth station location, the earth station elevation angle, height above see level, transmission polarization and frequency to determine the rain attenuation for a given time exceedance percentage. If measured rain rate statistics at the earth station location are not available, the model uses representative rain rate statistics for the climate region in

which the earth station is located. The representative US climate regions are shown in Figure A-1.

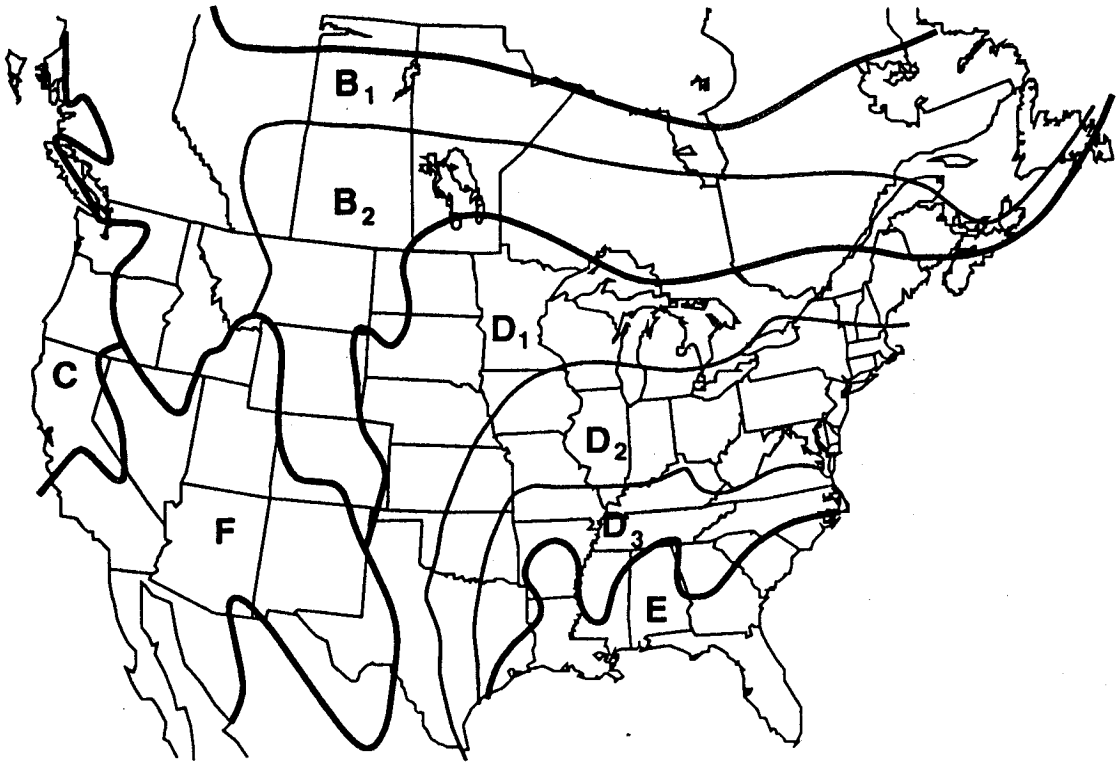

Figure A-1. Representative US Climate Regions [A-1]

2.1 Global Model Parameters

An implementation of the Global Model in spread sheet format is used to get rain attenuation for three of the US climate regions, region B, region D, and region E. Region D is more representative of average conditions in the US while region B represents dryer conditions and region E represents heavier precipitation conditions. The earth station elevation angle is assumed to be 30 degrees and the earth station height above sea level to be 100 meters. The polarization is assumed to be horizontal since this gives a more conservative estimates. Since rain attenuation on the downlink is also accompanied by an increase in earth station noise temperature which depends mainly on the LNA noise temperature, feed loss, and rain rate, the Global model implementation on the spread sheet also includes a calculation of the increase in noise temperature. The noise temperature degradation calculation assumes typical receiver parameters which are 0.4 dB feed loss and 250° K LNA. The increase in noise temperature degradation is added to the downlink rain attenuation to give the resultant fade margin. An example of the global model implementation in spread sheet format is given in Table A-1.

Table A-1. Global Model (Point Rain Rate Distributionin mm/hr)

PERCENT OF YEAR	RAIN CLIMATE REGION												HOURS PER YEAR
	A	B1	B	B2	C	D1	D= D2	D3	E	F	G	H	
0.001	28.5	45.0	57.5	70.0	78.0	90.0	108.0	126.0	165.0	66.0	185.0	253.0	0.09
0.002	21.0	34.0	44.0	54.0	62.0	72.0	89.0	106.0	144.0	51.0	157.0	220.5	0.18
0.005	13.5	22.0	28.5	35.0	41.0	50.0	64.5	80.5	118.0	34.0	120.5	178.0	0.44
0.010	10.0	15.5	19.5	23.5	28.0	35.5	49.0	63.0	98.0	23.0	94.0	147.0	0.88
0.020	7.0	11.0	13.5	16.0	18.0	24.0	35.0	48.0	78.0	15.0	72.0	119.0	1.75
0.050	4.0	6.4	8.0	9.5	11.0	14.5	22.0	32.0	52.0	8.3	47.0	86.5	4.38
0.100	2.5	4.2	5.2	6.1	7.2	9.8	14.5	22.0	35.0	5.2	32.0	64.0	8.76
0.200	1.5	2.8	3.4	4.0	4.8	6.4	9.5	14.5	21.0	3.1	21.8	43.5	17.52
0.500	0.7	1.5	1.9	2.3	2.7	3.6	5.2	7.8	10.6	1.4	12.2	22.5	43.80
1.000	0.4	1.0	1.3	1.5	1.8	2.2	3.0	4.7	6.0	0.7	8.0	12.0	87.60
2.000	0.1	0.5	0.7	0.8	1.1	1.2	1.5	1.9	2.9	0.2	5.0	5.2	175.20
5.000	0.0	0.2	0.3	0.3	0.5	0.0	0.0	0.0	0.5	0.0	1.8	1.2	438.00

2.2 Rain Degradation

Using the Global Model, Figures A-2 through A-4 show the expected rain attenuations plotted against the exceedance time percentage. As mentioned above, the downlink plots include the degradation due to the increase in noise temperature. The plots are shown for 1.0 % to 0.001 % exceedance time percentage which translates to availabilities of 99 % to 99.999 %.

3 Transmission Link Analysis

The transmission link analysis is performed to determine transmission bit rates that can be supported by various size earth stations and the link margins available for rain attenuation and other impairments. Understandably, the results of the analysis are highly dependent on the assumptions made with regard to the spacecraft parameters, specifically the spacecraft G/T and EIRP. For this reason, the link analysis uses the parameters of high EIRP domestic spacecraft at Ku-Band with a G/T of 5 dB/K and EIRP of 50 dBW, and those of ACTS spot beam parameters at Ka-band with a G/T of 20 dB/K and EIRP of 60 dBW. Both of these parameter sets are considered representative of what is currently achievable with an advanced spacecraft design, and further improvements are likely to be made in future generation spacecrafts. The parameters used for 50/40 GHz bands are scaled from the Ka-band parameters by the addition of a partial increase in spacecraft antenna gains. A satellite EIRP of 65 dBW and a satellite receiver G/T of 23 dB/K are thus assumed at 40/50 GHz.

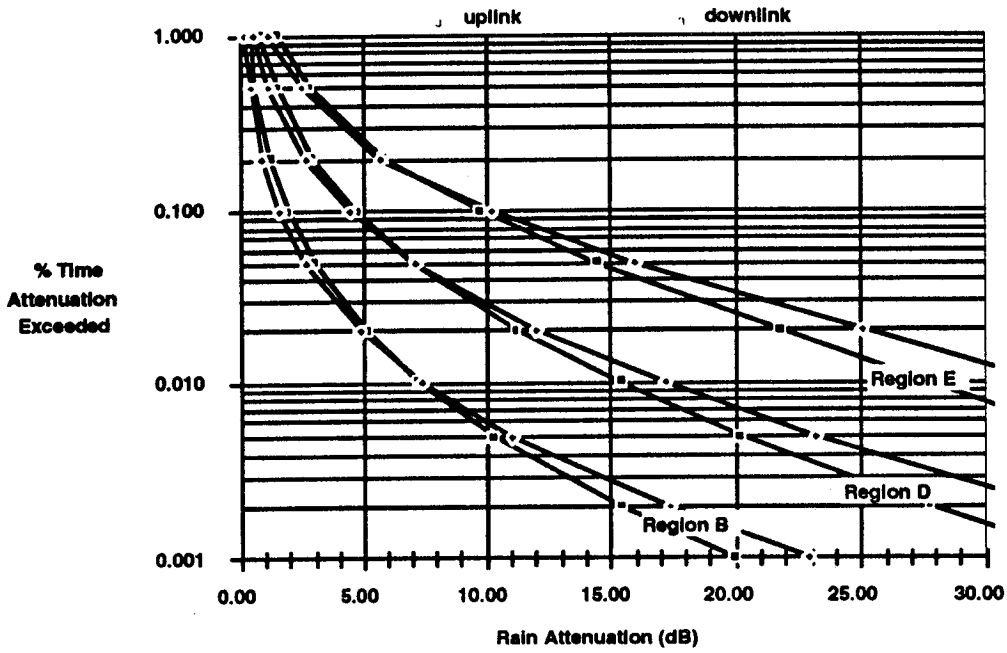

Figure A-2. Rain Attenuation at at 12/14 GHz

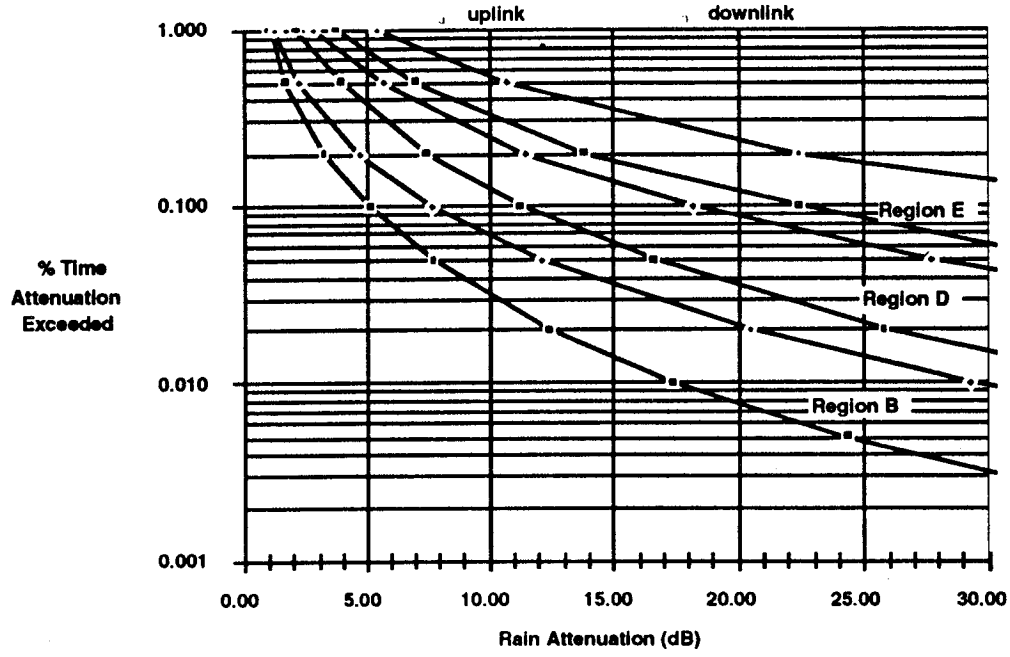

Figure A-3. Rain Attenuation at 20/30 GHz

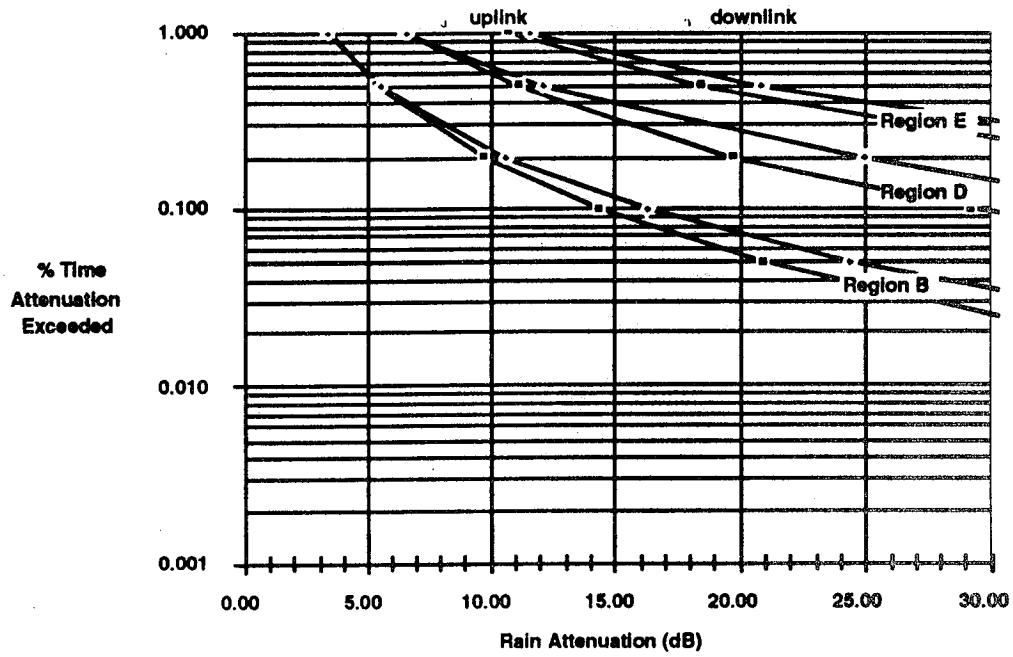

Figure A-4. Rain Attenuation at 40/50 GHz

3.1 Link Budget Specifics

Sample link budget calculations are shown in Tables A-2 through A-11 for the three frequency bands of interest. The link budgets assume separation of the uplink and downlink by on-board regeneration. The EIRP values which are shown in the table are given per carrier. For all the transmission rates of interest, the per carrier EIRP values are identical to the transponder EIRP values, implying single carrier operation. The only exception to this is when the transmission rate is 10 Mbit/s, in which case the satellite EIRP per carrier is scaled down by 7 dB (5 carriers per transponder). The modulation format used is coded QPSK and the bit error rate is assumed to be 10^{-8}. Code rates varying from 1/2 to 1/1 (no coding) are used. The theoretical Eb/No of the convolutional code rate is used in the link budgets, however additional losses of 3 dB are also included to account for modem implementation loss, feed loss, pointing loss, and possible loss due to interference. The link margins thus obtained are over and beyond the 3 dB losses included in the link budgets, and represent additional margins at a BER of 10^{-8}. Single carrier operation is assumed at the transmit earth station and on the satellite downlink so that no entry is made for intermodulation noise. The earth station elevation angle is assumed to be 30 degrees.

Table A-2. Representative Link Budget (Ku-Band, 620 Mbit/s)

UPLINK			DOWNLINK		
Frequency	14.2	GHz	Frequency	11.2	GHz
Antenna Diameter	9.0	meter	Antenna Diameter	9.0	meter
HPA Size	100.0	W	System Noise Temp.	214	K
Antenna Gain	59.9	dBi	Antenna Gain	57.9	dBi
E/S EIRP	79.9	dBW	E/S G/T	34.6	dB/K
Satellite G/T	5.0	dB/K	Satellite EIRP	50.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	207.3	dB	Free Space Loss	205.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	103.3	dBHz	C/No	105.0	dBHz
Required Eb/No	7.1	dB	Required Eb/No	7.1	dB
Carrier Info. Rate	620.0	Mbit/s	Carrier Info. Rate	620.0	Mbit/s
Requierd C/No	95.0	dBHz	Requierd C/No	95.0	dBHz
FEC Code Rate	0.75		FEC Code Rate	0.75	
Channel Bit Rate	826.7	Mbit/s	Channel Bit Rate	826.7	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	413.3	Msym/s	Symbol Rate	413.3	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	496.0	MHz	Occupied Bandwidth	496.0	MHz
Margin	8.3	dB	Margin	10.0	dB

Table A-3. Representative Link Budget (Ku-Band, 155 Mbit/s)

UPLINK			DOWNLINK		
Frequency	14.2	GHz	Frequency	11.2	GHz
Antenna Diameter	5.0	meter	Antenna Diameter	5.0	meter
HPA Size	50.0	W	System Noise Temp.	214	K
Antenna Gain	54.8	dBi	Antenna Gain	52.8	dBi
E/S EIRP	71.8	dBW	E/S G/T	29.5	dB/K
Satellite G/T	5.0	dB/K	Satellite EIRP	50.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	207.3	dB	Free Space Loss	205.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	95.2	dBHz	C/No	99.9	dBHz
Required Eb/No	6	dB	Required Eb/No	6	dB
Carrier Info. Rate	155.0	Mbit/s	Carrier Info. Rate	155.0	Mbit/s
Requierd C/No	87.9	dBHz	Requierd C/No	87.9	dBHz
FEC Code Rate	0.50		FEC Code Rate	0.50	
Channel Bit Rate	310.0	Mbit/s	Channel Bit Rate	310.0	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	155.0	Msym/s	Symbol Rate	155.0	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	186.0	MHz	Occupied Bandwidth	186.0	MHz
Margin	7.3	dB	Margin	12.0	dB

Table A-4. Representative Link Budget (Ku-Band, 50 Mbit/s)

UPLINK			DOWNLINK		
Frequency	14.200	GHz	Frequency	11.200	GHz
Antenna Diameter	2.4	meter	Antenna Diameter	2.4	meter
HPA Size	50.0	W	System Noise Temp.	214	K
Antenna Gain	48.5	dBi	Antenna Gain	46.4	dBi
E/S EIRP	65.5	dBW	E/S G/T	23.1	dB/K
Satellite G/T	5.0	dB/K	Satellite EIRP	50.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	207.3	dB	Free Space Loss	205.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	88.8	dBHz	C/No	93.5	dBHz
Required Eb/No	6	dB	Required Eb/No	6	dB
Carrier Info. Rate	50.0	Mbit/s	Carrier Info. Rate	50.0	Mbit/s
Requierd C/No	83.0	dBHz	Requierd C/No	83.0	dBHz
FEC Code Rate	0.50		FEC Code Rate	0.50	
Channel Bit Rate	100.0	Mbit/s	Channel Bit Rate	100.0	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	50.0	Msym/s	Symbol Rate	50.0	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	60.0	MHz	Occupied Bandwidth	60.0	MHz
Margin	5.8	dB	Margin	10.5	dB

Table A-5. Representative Link Budget (Ku-Band, 10 Mbit/s)

UPLINK			DOWNLINK		
Frequency	14.2	GHz	Frequency	11.2	GHz
Antenna Diameter	1.2	meter	Antenna Diameter	1.2	meter
HPA Size	20.0	W	System Noise Temp.	214	K
Antenna Gain	42.4	dBi	Antenna Gain	40.4	dBi
E/S EIRP	55.5	dBW	E/S G/T	17.1	dB/K
Satellite G/T	5.0	dB/K	Satellite EIRP	43.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	207.3	dB	Free Space Loss	205.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	78.8	dBHz	C/No	80.5	dBHz
Required Eb/No	6	dB	Required Eb/No	6	dB
Carrier Info. Rate	10.0	Mbit/s	Carrier Info. Rate	10.0	Mbit/s
Requierd C/No	76.0	dBHz	Requierd C/No	76.0	dBHz
FEC Code Rate	0.50		FEC Code Rate	0.50	
Channel Bit Rate	20.0	Mbit/s	Channel Bit Rate	20.0	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	10.0	Msym/s	Symbol Rate	10.0	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	12.0	MHz	Occupied Bandwidth	12.0	MHz
Margin	2.8	dB	Margin	4.5	dB

Table A-6. Representative Link Budget (Ka-Band, 1240 Mbit/s)

UPLINK			DOWNLINK		
Frequency	30.0	GHz	Frequency	20.0	GHz
Antenna Diameter	7.5	meter	Antenna Diameter	7.5	meter
HPA Size	100.0	W	System Noise Temp.	500	K
Antenna Gain	64.9	dBi	Antenna Gain	61.3	dBi
E/S EIRP	84.9	dBW	E/S G/T	34.3	dB/K
Satellite G/T	20.0	dB/K	Satellite EIRP	60.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	213.8	dB	Free Space Loss	210.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	116.7	dBHz	C/No	109.7	dBHz
Required Eb/No	7.0	dB	Required Eb/No	7.0	dB
Carrier Info. Rate	1240.0	Mbit/s	Carrier Info. Rate	1240.0	Mbit/s
Requierd C/No	97.9	dBHz	Requierd C/No	97.9	dBHz
FEC Code Rate	0.75		FEC Code Rate	0.75	
Channel Bit Rate	1653.3	Mbit/s	Channel Bit Rate	1653.3	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	826.7	Msym/s	Symbol Rate	826.7	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	992.0	MHz	Occupied Bandwidth	992.0	MHz
Margin	18.8	dB	Margin	11.8	dB

Table A-7. Representative Link Budget (Ka-Band, 620 Mbit/s)

UPLINK			DOWNLINK		
Frequency	30.0	GHz	Frequency	20.0	GHz
Antenna Diameter	5.0	meter	Antenna Diameter	5.0	meter
HPA Size	50.0	W	System Noise Temp.	500	K
Antenna Gain	61.3	dBi	Antenna Gain	57.8	dBi
E/S EIRP	78.3	dBW	E/S G/T	30.8	dB/K
Satellite G/T	20.0	dB/K	Satellite EIRP	60.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	213.8	dB	Free Space Loss	210.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	110.2	dBHz	C/No	106.2	dBHz
Required Eb/No	6.0	dB	Required Eb/No	6.0	dB
Carrier Info. Rate	620.0	Mbit/s	Carrier Info. Rate	620.0	Mbit/s
Requierd C/No	93.9	dBHz	Requierd C/No	93.9	dBHz
FEC Code Rate	0.50		FEC Code Rate	0.50	
Channel Bit Rate	1240.0	Mbit/s	Channel Bit Rate	1240.0	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	620.0	Msym/s	Symbol Rate	620.0	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	744.0	MHz	Occupied Bandwidth	744.0	MHz
Margin	16.3	dB	Margin	12.3	dB

Table A-8. Representative Link Budget (Ka-Band, 155 Mbit/s)

UPLINK			DOWNLINK		
Frequency	30.0	GHz	Frequency	20.0	GHz
Antenna Diameter	2.4	meter	Antenna Diameter	2.4	meter
HPA Size	50.0	W	System Noise Temp.	500	K
Antenna Gain	55.0	dBi	Antenna Gain	51.4	dBi
E/S EIRP	72.0	dBW	E/S G/T	24.5	dB/K
Satellite G/T	20.0	dB/K	Satellite EIRP	60.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	213.8	dB	Free Space Loss	210.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	103.8	dBHz	C/No	99.9	dBHz
Required Eb/No	6.0	dB	Required Eb/No	6.0	dB
Carrier Info. Rate	155.0	Mbit/s	Carrier Info. Rate	155.0	Mbit/s
Requierd C/No	87.9	dBHz	Requierd C/No	87.9	dBHz
FEC Code Rate	0.50		FEC Code Rate	0.50	
Channel Bit Rate	310.0	Mbit/s	Channel Bit Rate	310.0	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	155.0	Msym/s	Symbol Rate	155.0	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	186.0	MHz	Occupied Bandwidth	186.0	MHz
Margin	15.9	dB	Margin	11.9	dB

Table A-9. Representative Link Budget (Ka-Band, 50 Mbit/s)

UPLINK			DOWNLINK		
Frequency	30.0	GHz	Frequency	20.0	GHz
Antenna Diameter	1.2	meter	Antenna Diameter	1.2	meter
HPA Size	20.0	W	System Noise Temp.	500	K
Antenna Gain	48.9	dBi	Antenna Gain	45.4	dBi
E/S EIRP	62.0	dBW	E/S G/T	18.4	dB/K
Satellite G/T	20.0	dB/K	Satellite EIRP	60.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	213.8	dB	Free Space Loss	210.2	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	93.8	dBHz	C/No	93.8	dBHz
Required Eb/No	6.0	dB	Required Eb/No	6.0	dB
Carrier Info. Rate	50.0	Mbit/s	Carrier Info. Rate	50.0	Mbit/s
Requierd C/No	83.0	dBHz	Requierd C/No	83.0	dBHz
FEC Code Rate	0.50		FEC Code Rate	0.50	
Channel Bit Rate	100.0	Mbit/s	Channel Bit Rate	100.0	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	50.0	Msym/s	Symbol Rate	50.0	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	60.0	MHz	Occupied Bandwidth	60.0	MHz
Margin	10.8	dB	Margin	10.8	dB

Table A-10. Representative Link Budget (40/50-GHz, 620 Mbit/s)

UPLINK			DOWNLINK		
Frequency	50.0	GHz	Frequency	40.0	GHz
Antenna Diameter	7.5	meter	Antenna Diameter	7.5	meter
HPA Size	100.0	W	System Noise Temp.	500	K
Antenna Gain	69.3	dBi	Antenna Gain	67.4	dBi
E/S EIRP	89.3	dBW	E/S G/T	40.4	dB/K
Satellite G/T	23.0	dB/K	Satellite EIRP	65.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	218.2	dB	Free Space Loss	216.3	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	119.7	dBHz	C/No	114.7	dBHz
Required Eb/No	7.1	dB	Required Eb/No	7.1	dB
Carrier Info. Rate	620.0	Mbit/s	Carrier Info. Rate	620.0	Mbit/s
Requierd C/No	95.0	dBHz	Requierd C/No	95.0	dBHz
FEC Code Rate	0.75		FEC Code Rate	0.75	
Channel Bit Rate	826.7	Mbit/s	Channel Bit Rate	826.7	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	413.3	Msym/s	Symbol Rate	413.3	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	496.0	MHz	Occupied Bandwidth	496.0	MHz
Margin	24.7	dB	Margin	19.7	dB

Table A-11. Representative Link Budget (40/50-GHz, 155 Mbit/s)

UPLINK			DOWNLINK		
Frequency	50.0	GHz	Frequency	40.0	GHz
Antenna Diameter	2.4	meter	Antenna Diameter	2.4	meter
HPA Size	50.0	W	System Noise Temp.	500	K
Antenna Gain	59.4	dBi	Antenna Gain	57.5	dBi
E/S EIRP	76.4	dBW	E/S G/T	30.5	dB/K
Satellite G/T	23.0	dB/K	Satellite EIRP	65.0	dBW
Elevation Angle	30.00	deg	Elevation Angle	30.00	deg
Free Space Loss	218.2	dB	Free Space Loss	216.3	dB
Other Losses	3.0	dB	Other Losses	3.0	dB
C/No	106.8	dBHz	C/No	104.9	dBHz
Required Eb/No	6.0	dB	Required Eb/No	6.0	dB
Carrier Info. Rate	155.0	Mbit/s	Carrier Info. Rate	155.0	Mbit/s
Requierd C/No	87.9	dBHz	Requierd C/No	87.9	dBHz
FEC Code Rate	0.50		FEC Code Rate	0.50	
Channel Bit Rate	310.0	Mbit/s	Channel Bit Rate	310.0	Mbit/s
Modulation Format	QPSK		Modulation Format	QPSK	
Bits/Symbol	2	Bits	Bits/Symbol	2	Bits
Symbol Rate	155.0	Msym/s	Symbol Rate	155.0	Msym/s
BT Product	1.2		BT Product	1.2	
Occupied Bandwidth	186.0	MHz	Occupied Bandwidth	186.0	MHz
Margin	18.9	dB	Margin	16.9	dB

3.2 Link Budget Results

Although the parameters used in the link budgets are considered representative and are by no means optimal, the link budget analysis illustrates typical bit rates that can be supported by various earth station sizes. A summary of the results is given in Tables A-12 through A-14, which shows the uplink and downlink margins obtained for a given earth station size (antenna and HPA), transmission bit rate, FEC code rate, and associated transmission bandwidth. In addition, these tables shows the resultant uplink and downlink percent time availability for the margins shown. These availabilities are directly obtained from the plots in Figure A-2 for US climate region D.

A close examination of the results shows that, as expected, Ku-band availabilities are higher than Ka-band, which in turn are higher than 50/40 GHz. Ku-band availabilities are generally on the order of 99.95 percent to 99.99 percent, except for the smallest size stations. The majority of Ka-band availabilities are in the range from 99.90 to 99.95 percent, but a few are in the range from 99.7 to 99.90 percent. For the 50/40 GHz band, availabilities mostly range from 99.7 to 99.90 percent.

4 Conclusions

The results indicate that an information rate of over 1 Gbit/s can be supported by large Ka-band earth stations, and small VSAT class terminals can be used for bit rates of up to 155 Mbit/s. A relaxed threshold BER requirement, e.g. 10^{-6}, will result in greater margin and thus availability, a smaller antenna or HPA size, or a higher bit rate. If additional margin or a lower threshold BER e.g. 10^{-11}, is needed, a high rate outer code, such as a Reed-Solomon code, may be used with a slight increase in the required bandwidth. On the other hand, if more bandwidth efficiency is required, bandwidth and power efficient modulation/coding formats may be used (e.g. high-rate trellis-coded 8-PSK). This, combined with a high rate outer code, can result in significant improvements in both bandwidth occupancy and power requirements at the expense of added receiver complexity. Due to the on-board processing nature of the system, this added complexity can be accommodated either on-board the satellite and/or at the earth station depending on the system requirements. A comparison between FEC-coded QPSK and rate-8/9 8-PSK is shown in Table A-15 for three earth station sizes at Ka-band.

5 References

[A-1] Propagation Effects Handbook for Satellite Systems Design - A Summary of Propagation Impairments on 10 to 100 GHz Satellite Link with Techniques for System Design, NASA Reference Publication 1082(04), U. S. Department of Commerce, 1989.

Table A-12. Link Analysis Results (Ku-Band)

ANTENNA DIAMETER (m)	HPA SIZE (w)	BIT RATE (Mbit/s)	FEC CODE RATE	CARR. BW (MHz)	UPLINK MARGIN (dB)	UPLINK AVAIL. (%)	DOWNLINK MARGIN (dB)	DOWNLINK AVAIL. (%)
12	100	620	1	372	5.9	99.938	7.6	99.955
12	100	620	0.875	425	9.9	99.972	11.6	99.978
12	100	620	0.75	496	10.8	99.975	12.5	99.981
12	100	310	1	186	8.9	99.967	10.6	99.976
12	100	310	0.875	213	12.9	99.982	13.6	99.984
12	100	310	0.75	248	13.8	99.984	15.5	99.990
12	100	310	0.5	372	14.9	99.985	16.6	99.992
9	100	620	0.75	496	8.3	99.962	10.0	99.973
9	100	310	0.75	496	11.3	99.978	13.0	99.983
7.5	100	310	0.5	372	10.8	99.975	12.5	99.981
7.5	100	155	0.75	144	12.7	99.981	14.4	99.987
7.5	100	155	0.5	186	13.8	99.984	15.5	99.990
5	50	155	0.5	186	7.3	99.953	12.0	99.980
2.4	50	50	0.5	60	5.8	99.938	10.5	99.976
2.4	20	50	0.5	60	1.9	99.71	10.5	99.976
2.4	20	10	0.5	12	8.9	99.967	10.5	99.976
1.2	20	10	0.5	12	2.8	99.81	4.5	99.89

Table A-13. Link Analysis Results (Ka-Band)

ANTENNA DIAMETER (m)	HPA SIZE (w)	BIT RATE (Mbit/s)	FEC CODE RATE	CARR. BW (MHz)	UPLINK MARGIN (dB)	UPLINK AVAIL. (%)	DOWNLINK MARGIN (dB)	DOWNLINK AVAIL. (%)
12	100	1240	1	744	17.9	99.89	10.9	99.89
12	100	1240	0.875	850	21.9	99.924	14.9	99.938
12	100	1240	0.75	992	22.8	99.928	15.8	99.943
12	100	620	1	372	20.9	99.920	13.9	99.930
12	100	620	0.875	496	25.8	99.941	18.8	99.960
9	100	1240	1	744	15.4	99.86	8.4	99.83
9	100	1240	0.875	850	19.4	99.908	12.4	99.911
9	100	1240	0.75	992	20.3	99.913	13.3	99.921
9	100	620	1	372	18.4	99.90	11.4	99.90
9	100	620	0.875	425	22.4	99.925	15.4	99.941
9	100	620	0.75	496	23.3	99.930	16.3	99.948
7.5	100	1240	0.75	992	18.8	99.903	11.8	99.908
7.5	100	1240	0.5	1488	19.8	99.910	12.8	99.918
7.5	100	620	1	372	16.8	99.88	9.8	99.86
7.5	100	620	0.875	425	20.8	99.920	13.8	99.930
7.5	100	620	0.75	496	21.7	99.923	14.7	99.937
5	50	620	0.875	425	14.3	99.84	10.3	99.87
5	50	620	0.75	496	15.2	99.86	11.2	99.90
5	50	620	0.5	744	16.3	99.87	12.3	99 911
5	50	310	1	186	13.3	99.83	9.3	99.85
5	50	310	0.875	213	17.3	99.88	13.3	99.921
2.4	50	155	0.75	124	14.8	99.85	10.8	99.89
2.4	50	155	0.5	186	15.9	99.86	11.9	99.908
2.4	50	50	0.75	40	19.7	99.910	15.8	99.943
2.4	20	50	0.5	60	16.9	99.88	16.9	99.952
1.2	20	155	0.5	186	5.9	99.50	5.9	99.70
1.2	20	50	0.5	60	10.8	99.78	10.8	99.89
1.2	20	10	0.5	12	17.8	99.89	10.8	99.89

Table A-14. Link Analysis Results (40/50-GHz Band)

ANTENNA DIAMETER (m)	HPA SIZE (w)	BIT RATE (Mbit/s)	FEC CODE RATE	CARR. BW (MHz)	UPLINK MARGIN (dB)	UPLINK AVAIL. (%)	DOWNLINK MARGIN (dB)	DOWNLINK AVAIL. (%)
12	100	1240	1	744	20.9	99.73	15.9	99.70
12	100	1240	0.875	850	24.9	99.80	19.9	99.80
12	100	1240	0.75	992	25.8	99.81	20.8	99.81
12	100	1240	0.5	1488	26.9	99.82	21.9	99.82
9	100	1240	0.5	1488	24.4	99.79	19.4	99.79
9	100	620	0.75	496	26.3	99.82	21.3	99.82
7.5	100	620	0.75	496	24.7	99.80	19.7	99.80
5	50	155	0.5	186	25.3	99.80	23.3	99.84
2.4	50	155	0.5	186	18.9	99.69	16.9	99.73
2.4	50	50	0.5	60	23.8	99.78	21.9	99.82
2.4	20	10	0.5	12	26.9	99.82	21.9	99.82
1.2	20	10	0.5	12	20.8	99.73	15.8	99.70

Table A-15. Comparison Between FEC-Coded QPSK and Rate-8/9 Trellis-Coded 8-PSK Modulation with RS(223, 255) Outer Code at BER = 10^{-8}

MODULATION FORMAT	ANTENNA DIAMETER	HPA SIZE(watts)	BIT RATE (Mbit/s)	FEC CODING RATE	CARRIER BANDWIDTH (MHz)	UPLINK MARGIN (dB)	DOWNLINK MARGIN (dB)
QPSK	12	100	1240	0.75	992	22.8	15.8
	9	100	1240	0.75	992	20.3	13.3
	5	50	620	0.5	744	16.3	12.3
8-PSK	12	100	1240	0.78	638	24.5	17.4
	9	100	1240	0.78	638	22.0	14.9
	5	50	620	0.78	319	16.9	12.8

Appendix B

Optical Transmission System Design

1 Introduction

The objective of this part of the study is to perform tradeoff analysis of the optical power budget required for ground to geosynchronous (GEO) satellite links and to identify critical technology/lead-time items that need further development and testing before its deployment. In this report, the results of detailed optical link budget calculations are presented as functions of various link parameters such as: data rates, antenna sizes, modulation/detection schemes, selected optical wavelengths and back-ground radiation.

Ground to GEO-satellite optical link power budget is critically dependent on the chosen optical technology and its feasibility in the specified time frame. For this study, selection of the appropriate technology and configuration/system has been based on COMSAT's previous experience with optical intersatellite links (ISL) which involved both system feasibility as well as critical hardware development [B-1] - [B-5]. The scope of these projects included evaluation of potential optical technologies (e.g. CO_2, Nd-YAG, InGaAsP/Nd-YAG, GaAlAs lasers) and optical modu-lation/detection techniques (e.g.. direct and heterodyne, analog and digital modulation of single-carrier and subcarrier multiplexed signals). For the present study, link performances at B-ISDN rates of 620 Mbit/s (OC-12), 1.24 Gbit/s (OC-24) and 2.48 Gbit/s (OC-48) were evaluated. Improvement of link performance using FEC and block codes (BCH) was also considered. Finally, the degradation of the link under various background radiation conditions was assessed.

At high transmission rates, inherent advantages of heterodyne vs. direct detection technology become more pronounced. However, stringent selection criteria of components, drive circuit complexity of optical trans-mitter and local oscillator and also phase matching of optical signals at the mixer/receiver are strong deterrents for implementing heterodyne technique. Although considerably improved in the last couple of years, optical heterodyne technology is still thought to be very complex and expensive for deployment in mid-1990s. Thus, a semiconductor diode laser based optical technology with direct detection (On-Off Keying or Quaternary Pulse Position Modulation) is recommended. For the optical links requiring a high data rate throughput, use of wavelength division multiplexing (WDM) technique could relax both the transmitted power and receiver bandwidth requirements of direct detection optical

transceivers. Adverse environmental effects on sharp-cutoff, narrow bandwidth multilayer optical filters, multiplexers/demultiplexers, diode lasers and photodetectors will require further examination.

2 Tradeoffs for the Optical Satellite Link

Optical power budget of ground to GEO-satellite links are discussed in this section. The design considerations for this link are described in detail as are the various tradeoffs.

2.1 Link Configuration and Component Selection

For the specifications of this task (see Table B-1), we consider a baseline configuration shown in Figure B-1. The Earth to geosynchronous satellite link assumes the use of a directly modulated semiconductor diode laser (AlGaAs) transmitter emitting at 0.85 μm wavelength, a direct detection optical receiver with an avalanche photodetector (APD) and suitable modulation/demodulation schemes (on-off keying/QPPM).

Table B-1. Ground to Geosynchronous Satellite Optical Link Parameters

PARAMETERS	DIRECT DETECTION	HETERODYNE DETECTION
Average Transmitter Power	500 mW, 1W	250mw, 500 mW
Wavelength	850 nm	850 nm
Range (Earth-Geosynchronous)	36,0000 km	36,000 km
Noise Equivalent Angle (NEA)	0.3 μradians	0.3 μradians
Tx Antenna Efficiency Factor (Off-axis Cassegrain)	-0.9 dB	-0.9 dB
Defocussing Loss	2 dB	2 dB
Total Absorption Loss	2.5 dB	2.5 dB
Point-Ahead Loss	1 dB	1 dB
Modem Implementation Margin	1 dB	1 dB
System Margin	6 dB	6 dB
Detector Responsitivity	0.6 A/W	0.6 A/W
Detector Dark Current	1 nA	1 nA
Detector (APD/PIN) Gain	125	1
Detector Shunt Capacitance	2 pF	2 pF
Preamplifier Noise Temperature	500° K	
Heterodyne Mixing Loss		1 dB
Quantum Noise Factor		1 dB
Extinction Ratio	0.05	0.05
Accumulated Bit Error Rate	10^{-8}	10^{-8}

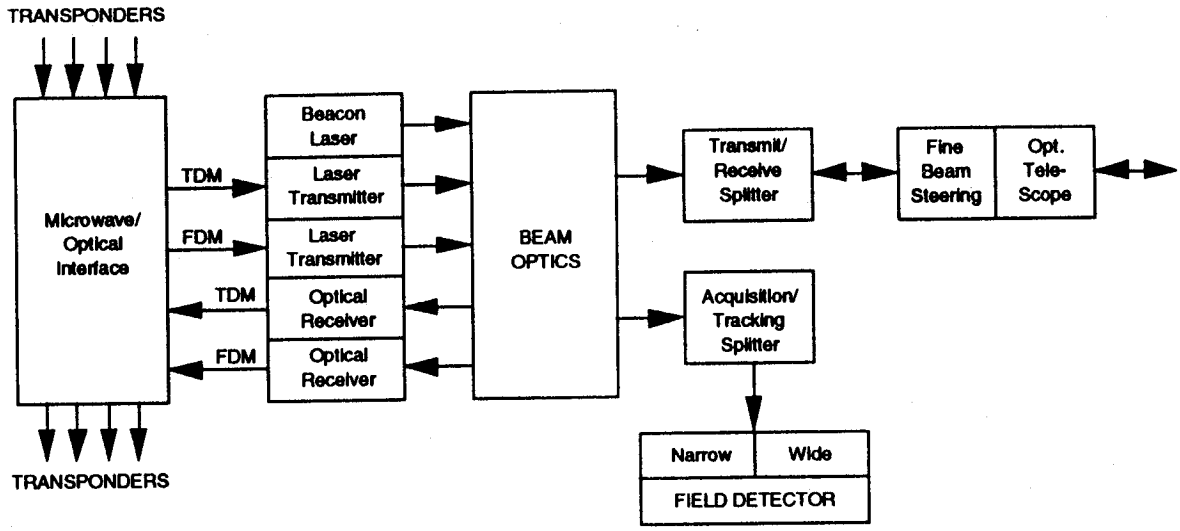

Figure B-1. Baseline Configuration of Optical Satellite Link Payload

The optical transmitter output power is chosen in the range of 50 mW to 1W cw. Commercially available single-element diode lasers do not emit more than 100 mW cw. High output power (100 mW and above) laboratory devices have poor lifetime (100 to 2,000 hours). High power transmitters will require optical power combining of several lasers. However, practical problems of implementing such combining schemes, e.g. stability of beam optics, synchronization of modulation drive to the lasers, etc., limit the number of lasers and thus transmitter output power to 1 kW. Four important reasons for selecting short wavelength (GaAlAs) vis-a-vis long wavelength (InGaAsP) diode lasers are: (a) lower atmospheric absorption losses, (b) maturity of technology, (c) higher optical antenna gain at shorter wavelength due to diffraction-limited optics and (d) better thermal properties of III-V ternaries.

Transmission using M-ary PPM with M>2 is more efficient than OOK. However, M-ary PPM systems requires the lasers to be driven at a higher rate. The selected lasers must be compatible with the effective modulation bandwidth, and the high-speed electronics capable of deliver-ing large currents required to drive the high-power diode lasers. Another important consideration stems from the fact that the relative advantage of M-ary PPM for M>2 is based on the assumption that the peak output power of the laser is increased M-fold compared to OOK scheme. Diode lasers operated at large drive currents and/or high output powers usually degrade at the facets due to defects caused by heat dissipation. If the M-ary PPM test performances, which are very limited at this stage, are true indicators, it is most likely that lasers may not be reliably operated at the high peak power levels as appropriate for a large M-ary PPM system.

2.2 Optical Power Budget Analysis

The received optical power of a ground to GEO-satellite link fluctuates with time because the pointing and tracking subsystem has a finite response time. An accurate

assessment of the resulting BER needs to take into account the pointing error distribution with time and the far-field pattern of the optical antenna. The power level in the beam center and the relative widths of the two functions determine the BER of the optical link. Assuming that the power level in the center is fixed, a wider beam or a smaller tracking error clearly results in a smaller time percentage with low receive power or high relative error probability. The overall effect of tracking error on the ground to GEO-satellite optical link system per-formance is parameterized by the noise equivalent angle (NEA) of the link.

Thus, to make the power budget tradeoff analysis useful, a large number of system parameters must be judiciously selected at the outset. Among these are the laser power output, the NEA and the optical receiver characteristics (e.g. quantum efficiency, dark current, preamplifier noise temperature, etc.). Other parameters such as load resistor and avalanche gain of the the photodetector are optimized for each specific case.

The relevant system parameters for the OOK modulated/ direct detection optical ISL are given in Table B-2. The loss parameters shown in this table control the link budget through the range equation. Several factors such as selection of suitable beam optics, point-ahead considerations for satellite drift, effi-ciency and obscuration of transmit/receive antenna, extinction ratio and link distance determine the link budget. These factors require to be carefully estimated for they are dependent on the link requirements and limitations of the available technology.

The power budget calculation includes the effect of tracking errors by computing the BER in small increments using an Airy far-field pattern and a Gaussian distribution for the tracking NEA. The incremental contri-butions to the BER are determined stepwise by increasing the tracking error and calculating the optical receive power and the signal-to-noise ratio at the detector output. For each increment of the tracking error, a time percentage is obtained. The accumulated BER of the optical link is then the sum of all the incremental contributions.

To aid in the calculation, the range equation is programmed to include the loss factors such as: atmospheric absorption, absorption and scattering in the optics, beam defocusing, transmit/receive beam splitting, optical multiplexing and demultiplexing, beam truncations and point-ahead errors as variable parameters. The sensitivity of the antenna dia-meter and the link budget to variation in the loss factors places practical upper and lower bounds on these estimates.

An overall system margin of 6 dB is assumed for the GEO to ground link budget presented here. Given these link parameters and the laser transmitter power, the link BER performance at the required data rates can then be predicted for specified noise equivalent angle and receiver sensi-tivity.

Table B-2. The Range Equation and Receiver Sensitivity of an Optical Intersatellite Link

A. Range Equation

Detection analysis results in the determination of the optical power required at the receiver. The range equation relates the received power at the detector to the power at the transmitter which has been reduced by various system losses. The range equation for the optical link is given below:

$$P_r = P_t \bullet E_m \bullet G_t \bullet G_r \bullet L_{sp} \bullet L_{abs} \bullet L_p \bullet L_m$$

where, P_r = Average received power in Watts,

P_t = Average optical transmitter power in watts

E_m = Modulation efficiency

G_t = Transmit antenna gain = $\left(\frac{\pi D_t}{\lambda}\right)^2 \bullet L_o \bullet L_d$

where; Dt = diameter of transmit antenna

λ = wavelength of optical transmission

L_o = antenna obscuration loss

$= 1 - \Psi^2$ (Ψ = obscuration factor= ratio of secondary to primary mirror radii)

L_d = defocussing loss in optics

G_r = Receive antenna gain= $\left(\frac{\pi D_r}{\lambda}\right)^2 \bullet L_o$

where; Dr = diameter of receive antenna

L_{sp} = Free Space Loss= $\left(\frac{\lambda}{4\pi R}\right)^2$

where; R = Range (Ground to satellite) in meters

L_{abs} = Absorption loss in transmit and receive optics/telescope

L_p = Pointing and tracking loss, including point-ahead loss

L_m = System Margin

Table B-2. The Range Equation and Receiver Sensitivity of an Optical Intersatellite Link (Cont'd)

B. Optical Receiver Sensitivities for Baseband Digital Formats

For a given value of Pt and other appropriate optical link system parameters, the range equation can be used to determine the antenna diameter requirements if the received power is known. The value of Pr is dependent upon the modulation format and the signal to noise ratio or the BER requirement at the receiver output. The receiver sensitivity at the specified data rate has been calculated for the OOK modulation/direct detection system using the relationship below. Binary Pulse Position Modulation (M-ary PPM, M=2) is a variation of the OOK modulation format. The receiver sensitivity is given by:

$$Pr = \frac{h\nu}{2\eta}\sqrt{\frac{C}{N}}\left\{F(\overline{M})R_b\sqrt{\frac{C}{N}} + \frac{2}{q}\sqrt{\left[qI_dF(\overline{M}) + \frac{2kT}{\overline{M}^2R_e}\right]R}\right\}$$

where h = Planck's constant (J.s)

ν = optical frequency (Hz)

C/N = carrier-to-noise ratio

$F(\overline{M})$ = excess noise factor

R = bit rate (bit/s)

q = electron charge (Coulombs)

I_d = photodetector thermally generated dark current (Amps)

k = Boltzmann's constant (Joules/°K)

T = temperature (°K)

$\overline{M}$ = average detector multiplication factor

R_e = thermal noise equivalent resistance (Ω)

When the optical source is modulated by a QPSK carrier, the average optical power required achieve a given signal to noise ratio is:

$$Pr = \frac{h\nu}{\eta}F(\overline{M})R\frac{C}{N}\left\{1 + \sqrt{1 + 2\left[\frac{qI_dF(\overline{M}) + 2kT/\overline{M}^2Re}{q^2F^2(\overline{M})R\ C/N}\right]}\right\}$$

2.3 Tradeoff Study Results

The system parameters of the OOK modulation/direct detection scheme with adaptive threshold, given in Table B-2, were used to compute the power budgets for the Earth to GEO optical links. Figure B-2 shows plots of optical antenna diameter vs. data rate for accumulated BERs of 10^{-7}, 10^{-8} and 10^{-9}; assuming 500 mW of transmitted power, OOK modulation and 0.3 μradian NEA. It is seen, under these conditions, that a transmitter aperture of 24 cm is needed to achieve a BER of 10^{-8} at 620 Mbit/s and that much larger apertures are needed at the higher data rates and for higher BERs. Figure B-3 demonstrates how the accumulated BER varies as a function of data rate for transmitter powers of 500 mW and 1 W. By increasing the optical transmitter power from 500 mW to 1W the optical link can support data rates of upto 1.3 Gbit/s as opposed to only 620 Mbit/s (assuming accumulated BER of 10^{-8}). This is further exemplified in Figure B-4 which shows how transmitter aperture varies as a function of transmit power, assuming a fixed accumulated BER of 10^{-8} and OOK modulation format. The trans-mitter aperture has been plotted for 3 data rates (620 Mbit/s, 1.24 Gbit/s and 2.48 Gbit/s) as well as for 3 sets of receiver apertures (i.e., S, the ratio of the receiver aperture to transmitter aperture, = 1, 2 and 3). This assumes the possibility of having larger receive or transmit apertures at the ground station than on the satellite. Assuming a transmit power of 1 W and a data rate of 620 Mbit/s, the transmit aperture can be reduced from 24 cm to about 12 cm by increasing S from 1 to 3.

Figure B-5 is a plot of accumulated BER as a function of aperture diameter for 500 mW and 1 W transmit powers. The effect of a constant background optical power of 1.56 μW (due to Sun and various stars) is also included here. It is seen here that the aperture diameter would have to be increased from 24 cm. to 35 cm. in order to maintain a BER of 10^{-8} in the presence of background radiation (assuming transmit power of 500 mW and data rate of 620 Mbit/s). To maintain 10^{-8} BER at 2.48 Gbit/s with back-ground radiation (1.56 μW) the required antenna diameter becomes im-practical. A similar effect is shown in Figure B-6, where the transmit power is the variable and the transmit aperture is kept constant. To maintain a BER of 10^{-8} in the presence of background radiation, the transmit power has to be increased from 500 mW to approximately 4 W. However, unlike the case in Figure B-5, it is possible, at 2.48 Gbit/s, to maintain a BER of 10^{-8} in the presence of background radiation, provided that the transmit power is in-creased from about 1W to about 10W (assuming transmitter aperture of 30.1 cm.).

An optimum NEA of 0.3 μradians is assumed in the calculation of BER as functions of data rate, transmit power and transmit antenna size. A tight NEA (i.e. $\leq$ 0.3 μrad) can enhance the BER performance; a useful tradeoff consideration in reducing optical power and/or antenna diameter requirements for the specified BER. Practical technology limitations of pointing and tracking subsystems, however, restrict this approach to moderate NEAs. A much relaxed NEA, on the other hand, may not meet the specified BER even with large antenna and/or high power transmitter.

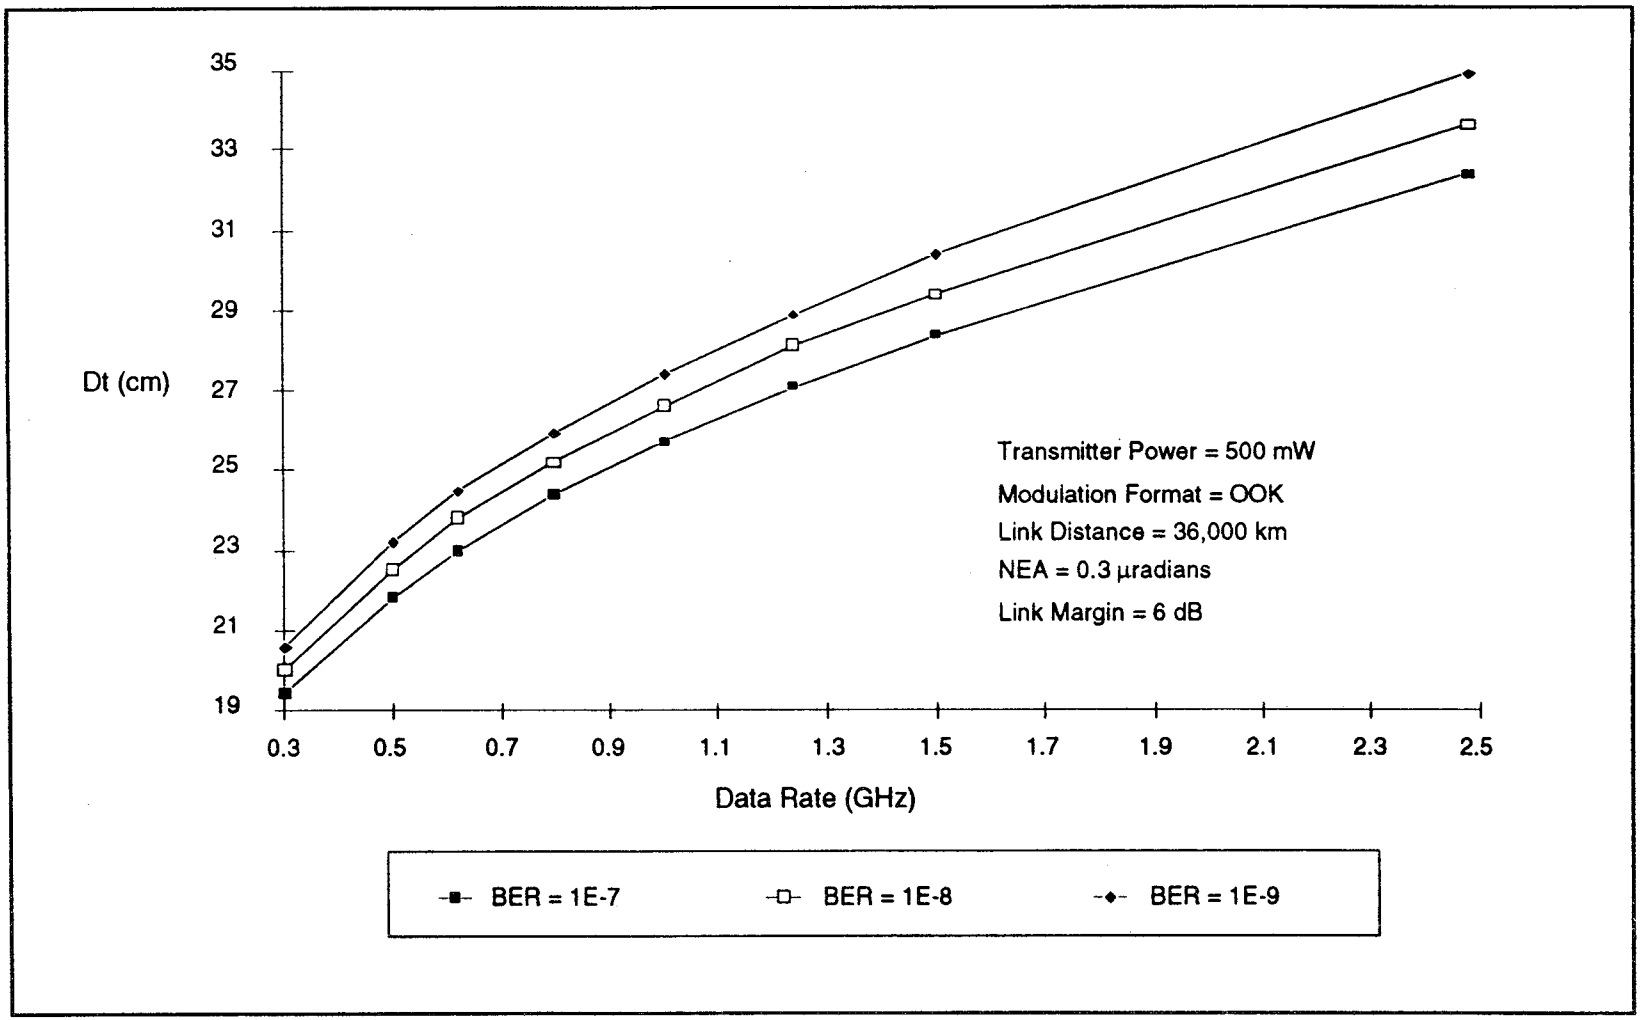

Figure B-2. Antenna Aperture vs. Data Rate for various B.E.Rs.

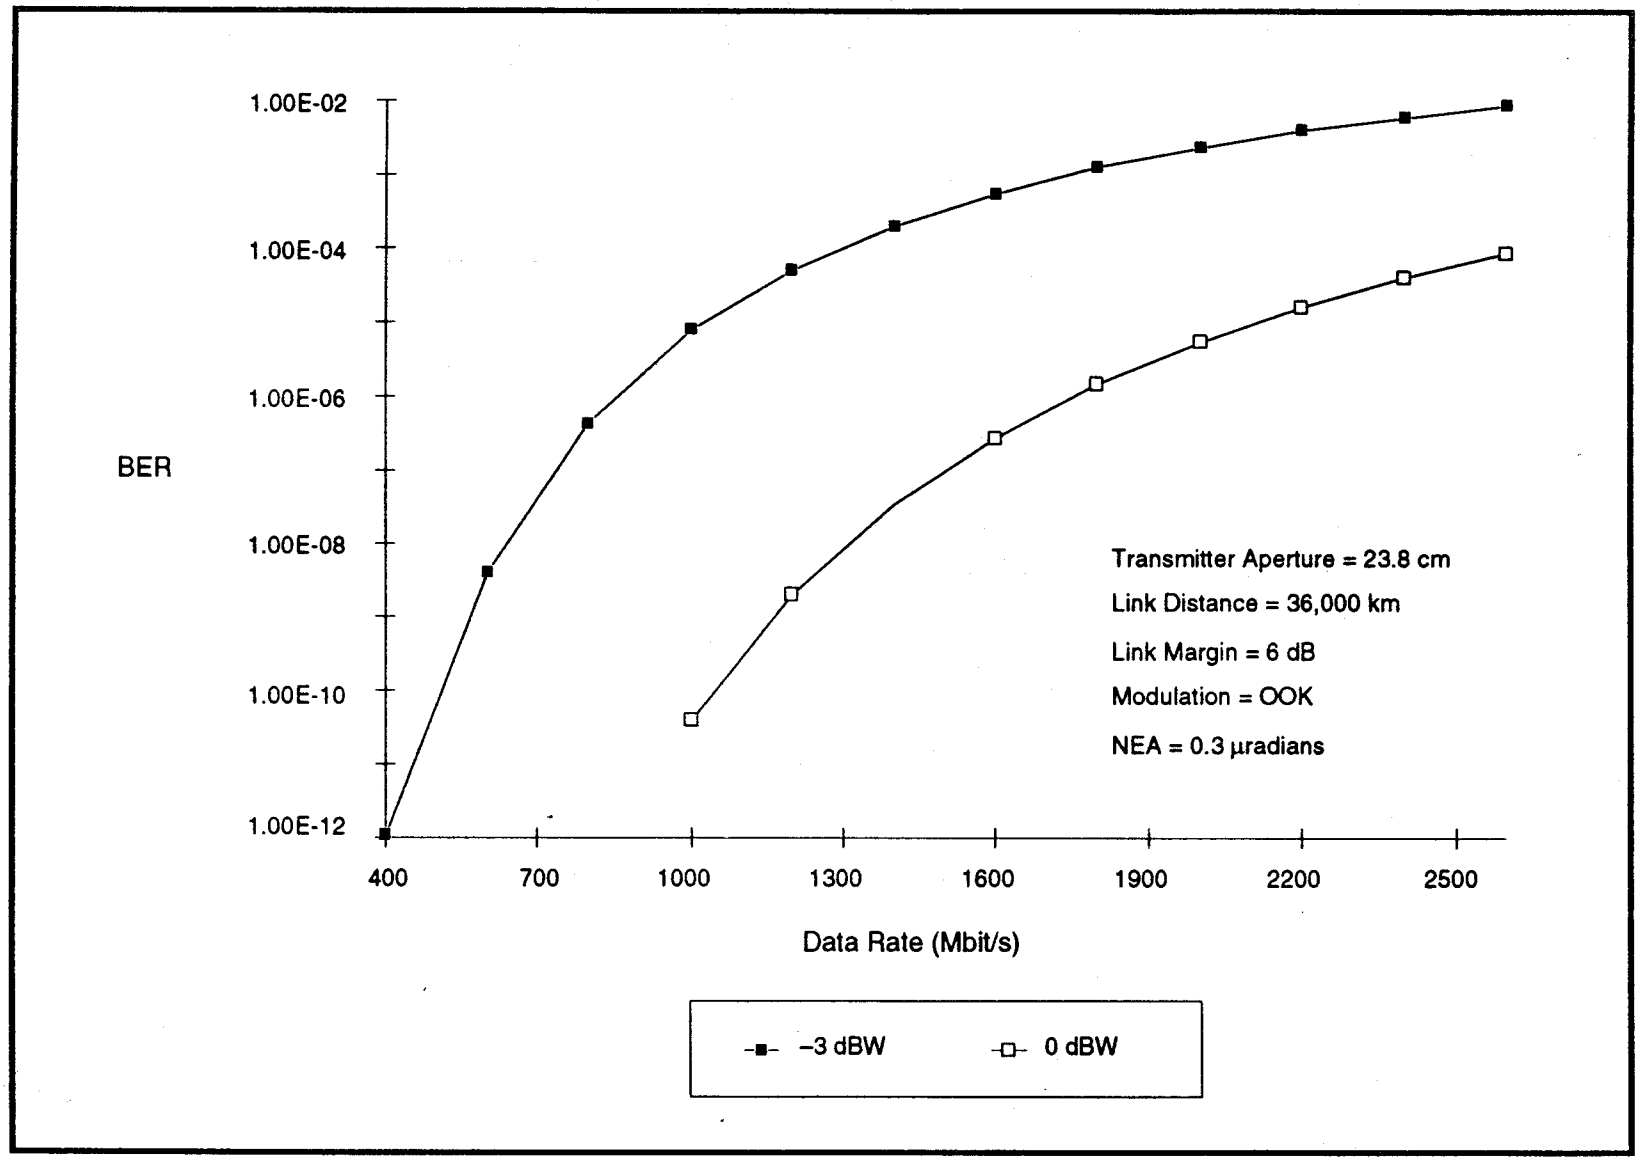

Figure B-3. BER vs. Data Rate for tranmitter power of 500 mW and 1 Watt.

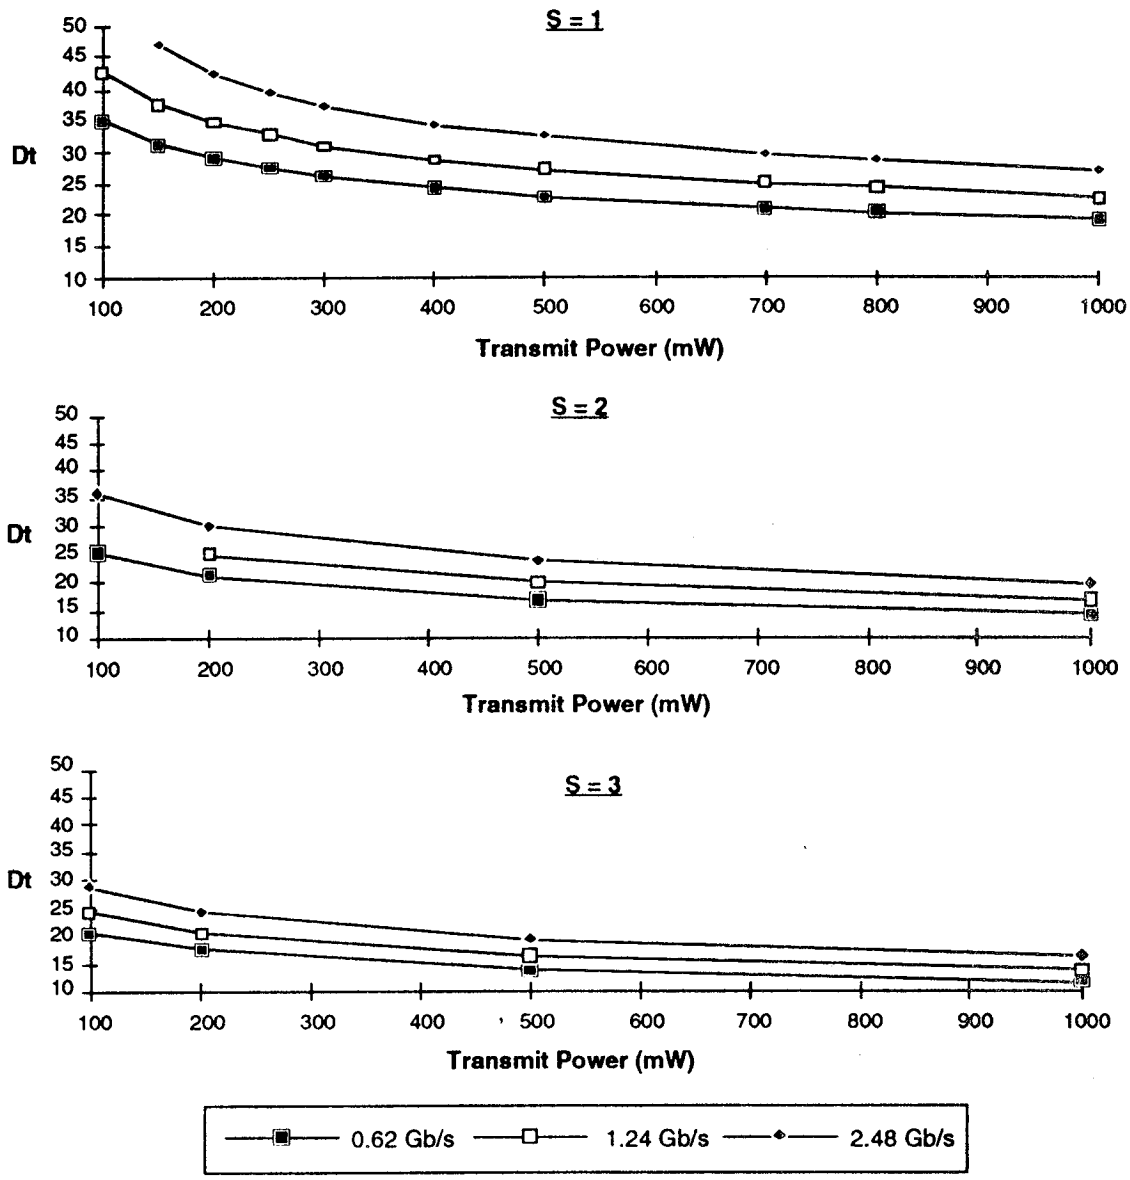

Figure B-4. Transmitter Aperture as a function of Transmit Power for Various Data Rates

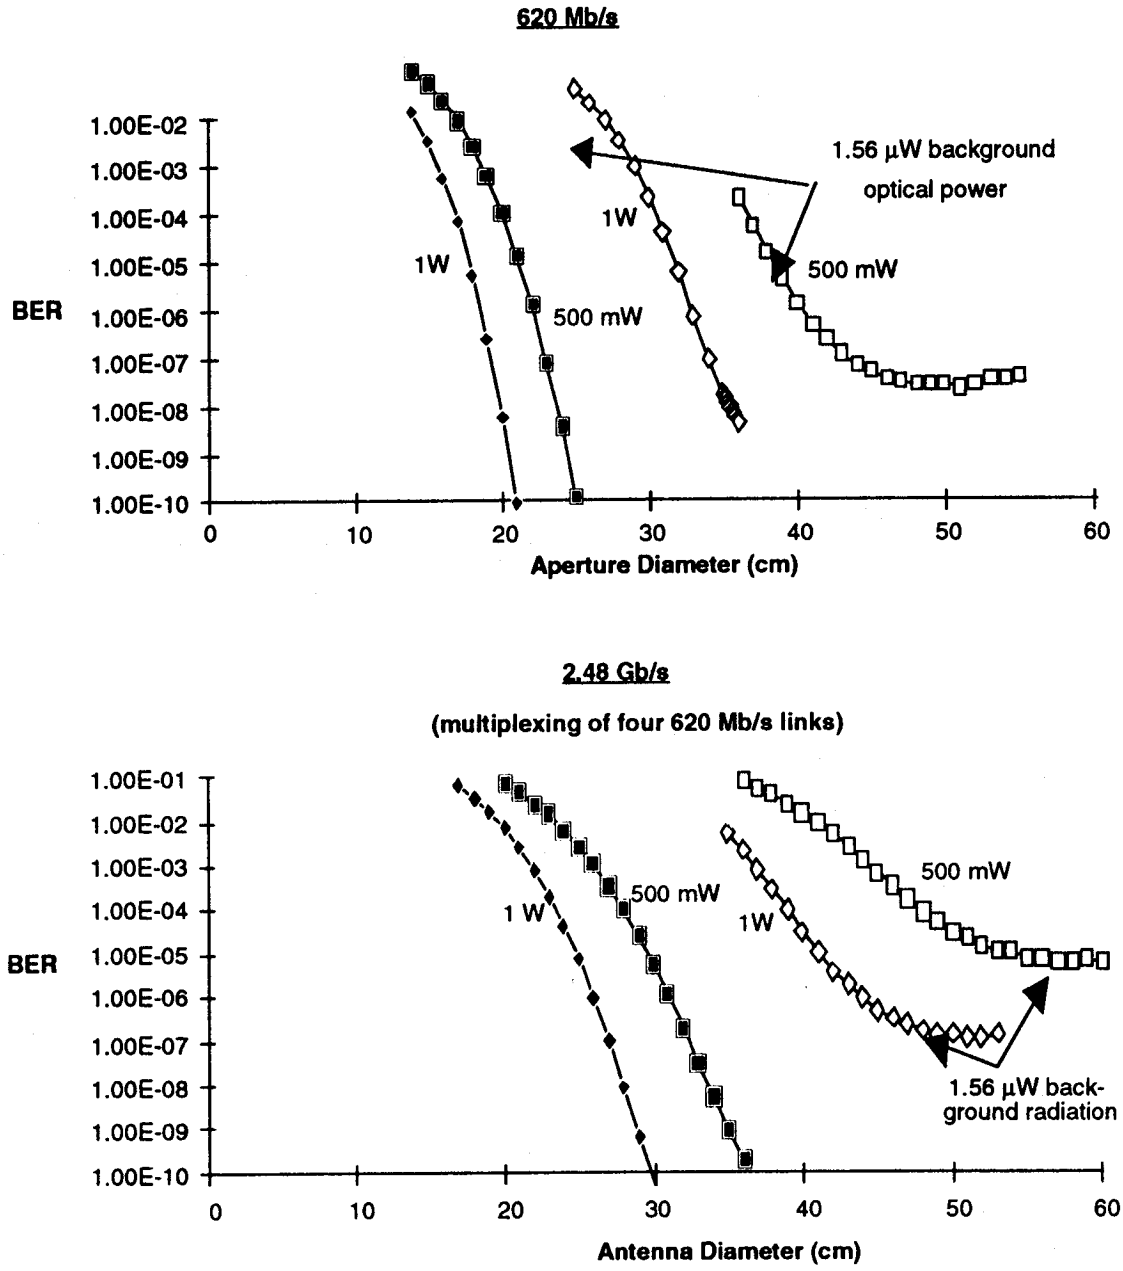

Figure B-5. BER as a function of antenna aperture with & without background radiation

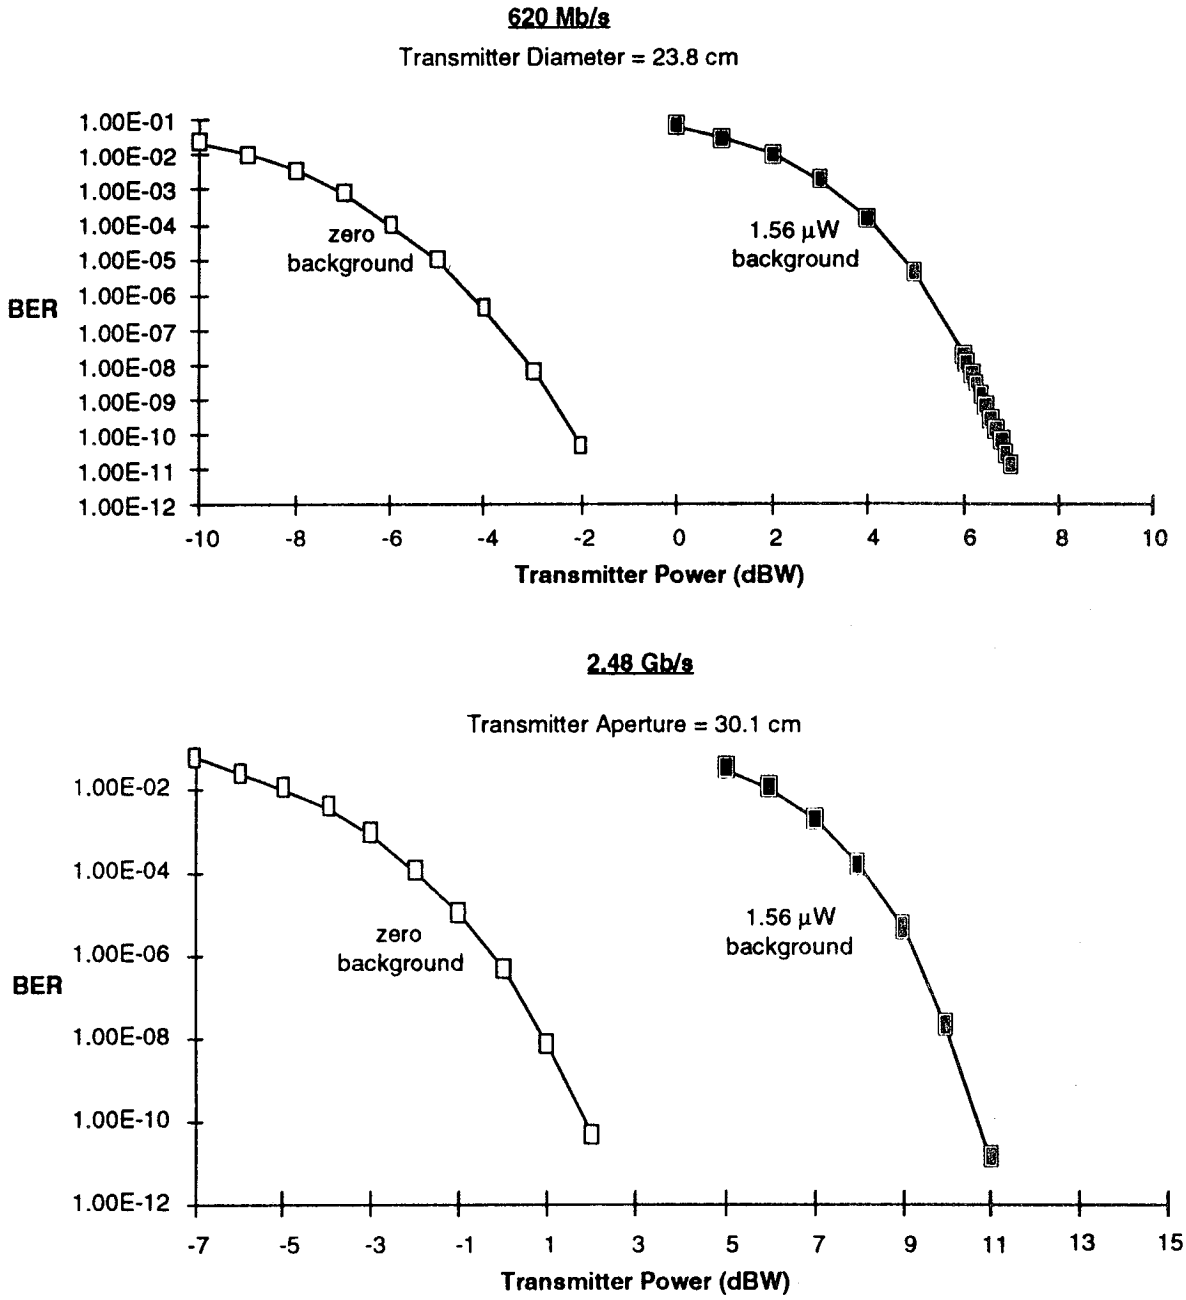

Figure B-6. BER as a function of transmitter power with & without background radiation

The design parameters for the OOK modulation/ direct detection ground to GEO-satellite optical link, for the specified BERs, can be obtained from the data given in Tables B-3 through B-6, where the required antenna diameter is traded with the optical transmitter power, loss factors and the NEA capability of the pointing and tracking device. The following obser-vations on the OOK direct detection and QFSK heterodyne detection link budget tradeoffs are worth noting:

- For the systems designed for 620 Mbit/s optical throughput, the input data architecture need not use optical multiplexing. Time-division multiplexing (TDM) method used to process the RF/microwave signals at the transmit and receive terminals on-board satellites is adequate for the recommended optical power levels and antenna diameters.

- The ground to GEO-satellite link calculations for the different data rates (620 Mbit/s, 1.24 Gbit/s and 2.48 Gbit/s) show reasonable antenna diameters (between 24 and 34 cm) for a 500 mW cw optical trans-mitter. Operation at 1 W provides a nominal saving in antenna diameter size (by 2 to 5 cm). At 100 mW these requirements are increased by 10 to 15 cm. Single diode lasers can deliver upto 100 mW output power. However, the antenna sizes necessary to maintain the high data rates will result in increased mass and volume of the satellite payload. A 500 mW power level appears to be reasonable from mass/volume considerations. Such high-power optical transmitters are designed to combine (dichroic filter/ incoherent) the output power of several lasers.

- For a reasonable antenna diameter (approx 24 cm.), the system BER degrades extremely rapidly as the NEA goes beyond 0.3 μradians. Table B-3 shows the power penalty from increasing NEA from 0.3 to 0.4 and 0.5 μradians. At 620 Mbit/s and for BER of 10^{-8}, the tracking penalty increases from 0.3 dB to 1.6 dB as the NEA increases from 0.3 to 0.5 μradians. At 1.24 and 2.48 Gbit/s it is impossible to maintain BER of 10^{-8} when the NEA is 0.5 μradians. At 1.24 Gbit/s and BER of 10^{-7}, the tracking penalty goes up from 0.4 to 2.6 dB as the NEA goes up from 0.3 to 0.5 μradians. Consequently, a tight control on the value of NEA to ≤0.3 μradians is considered of paramount importance.

- Increasing the value of S (i.e., ratio of receiver antenna diameter to transmitter diameter) allows reduction of the transmitter antenna size. For example; assuming OOK modulation at 2.48 Gbit/s rate with a transmitter power of 500 mW, the transmitter diameter required to maintain a BER of 10^{-8} is 33.6 cm for S=1, whereas for S=3 it is just 19 cm. For the QFSK system the corresponding values at S=1 and S=3 are 23.3 cm. and 13.4 cm., respectively. Both the transmitter power and antenna diameter can be reduced by using a larger S. This, however, requires a larger receive antenna. Such a design is accept-able in the Earth Station terminal, because the mass/volume limitations on-board the satellite are no longer valid.

Table B-3. Tracking Penalty with 500 mW Transmitter Power.

Modulation Format = OOK

Link Distance = 36,000 km

System Margin = 6 dB

S = 1

Data Rate (Gb/s)	B.E.R.	Tracking Penalty (dB) for following NEAs		
		0.3 μrad	0.4 μrad	0.5 μrad
0.62	1.00E-07	0.345	0.618	0.959
	1.00E-08	0.328	0.637	1.606
	1.00E-09	0.704	0.881	∞
1.24	1.00E-07	0.358	0.563	2.576
	1.00E-08	0.499	1.006	∞
	1.00E-09	0.414	1.482	∞
2.48	1.00E-07	0.339	1.784	∞
	1.00E-08	0.728	∞	∞
	1.00E-09	0.733	∞	∞

Table B-4. Transmitter Apertures as a function of Data Rate.

Transmit Power = 500 mW

Noise Equivalent Angle (NEA) = 0.3 μradians

Link Distance = 36,000 km

System Margin = 6 dB

BER = 1.0 E-8

Data Rate	OOK			QFSK		
(Gb/s)	S=1	S=2	S=3	S=1	S=2	S=3
0.62	23.8	16.6	13.6	16.4	11.6	9.4
1.24	28.1	19.7	16.0	19.6	13.8	11.2
2.48	33.6	23.4	19.0	23.3	16.4	13.4

Table B-5. Transmitter Apertures for different Multiplexing Schemes.

Transmit Power = 500 mW

Noise Equivalent Angle (NEA) = 0.3 μradians

Link Distance = 36,000 km

System Margin = 6 dB

S = 1

Data Rate (Gb/s)	Multiplexing Scheme	Mux/Dmux Loss (dB)	Required Transmitter Aperture (cm)		
			BER = 1E-7	BER = 1E-8	BER = 1E-9
0.62	Single	0	23	23.8	24.5
	2 units @ 310 Mb/s	2	22	22.7	23.3
1.24	Single	0	27.1	28.1	28.9
	2 units @ 620 Mb/s	2	25.8	26.7	27.5
	4 units @ 310 Mb/s	4	24.6	25.5	26.3
2.48	Single	0	32.4	33.6	34.9
	2 units @ 1.24 Gb/s	2	30.6	31.7	32.7
	4 units @ 620 Mb/s	4	29.1	30.1	31.1
	8 units @ 310 Mb/s	8	31.4	32.6	33.6

Table B-6. Transmitter Apertures for different Modulation Formats.

Transmit Power = 500 mW

Noise Equivalent Angle (NEA) = 0.3 μradians

Link Distance = 36,000 km

System Margin = 6 dB

S = 1

Data Rate (Gb/s)	Modulation Format	Required Transmitter Aperture (cm)		
		BER = 1E-7	BER = 1E-8	BER = 1E-9
0.62	OOK	23	23.8	24.5
	QPPM	22.6	23.5	24.2
	QFSK	15.8	16.4	16.9
1.24	OOK	27.1	28.1	28.9
	QPPM	27	27.9	28.75
	QFSK	18.9	19.6	20.3
2.48	OOK	32.4	33.6	34.9
	QPPM	32.3	33.5	34.7
	QFSK	22.5	23.3	24.1

- Another possible, albeit less effective, method of reducing antenna diameters is to use optical multiplexing of data. The trade-off here is between the reduced data rate for each link to the added multiplexer/ demultiplexer losses and the need for more transmitters and re-ceivers. For example, a 2.48 Gbit/s rate can be achieved by using 2 links (at different wavelengths) each at 1.24 Gbit/s; or by using 4 links at 620 Mbit/s each or by using 8 links at 310 Mbit/s each. It is seen from the results that using two links at 1.24 Gbit/s instead of one at 2.48 Gbit/s reduces the antenna diameter requirement (for 10^{-8} BER) from 34 to 32 cm. Using 4 links at 620 Mbit/s reduces the antenna diameter requirement to 30 cm. A further increase in the number of multi-plexed units, however, sets the antenna diameter requirement back up to 33 cm.

- Judicious use of coding (3 - 5 dB gain, typically) can reduce the optical power budget and/or antenna size for the link. For example, at 1.24 Gbit/s data rate and 500 mW transmit power, an antenna dia-meter of 27 cm is required to maintain 10^{-8} BER. A 3-db coding gain would reduce this to 22.7 cm., i.e. 16% reduction. This reduction is even more pronounced at lower transmitter powers.

- Optical heterodyne detection has significant advantages over direct detection; optical power budget and/or transmit aperture size requirements can be relaxed, specially at high data rates. At 2.48 Gbit/s both OOK and QPPM modulation schemes require an antenna aperture of ≈34 cm (assuming 500 mW transmit power). Using the QFSK scheme the aperture is reduced to 23 cm, resulting in a considerable savings in the payload mass and size. The technology of the coherent receiver subsytem is, however, very complex.

3 Focal Plan and RF/Optical Interface

To arrive at an optimal focal plane design and layout, numerous aspects of the ground to GEO-satellite optical link specifications must be considered. The major aspects of the focal plane design relevant to this study are identified in Figure B-7.

It is apparent that the ground to GEO-satellite optical link payload can be grouped into two subsystems: the communications and the pointing, acquisition and tracking subsystems. However, to successfully achieve the specified mission, these two subsystems should be interfaced in terms of a few common system parameters, such as antenna diameter and NEA. Special beam optics and adaptive feedback control units are needed to provide the required compatibility of the two beams. Also, the compact payload, for use on-board satellite, is likely to affect the details of its focal plane layout. The details at the subsystem and component levels are largely determined by the specific goals of the ground to GEO-satellite link and the optical power budget tradeoff considerations.

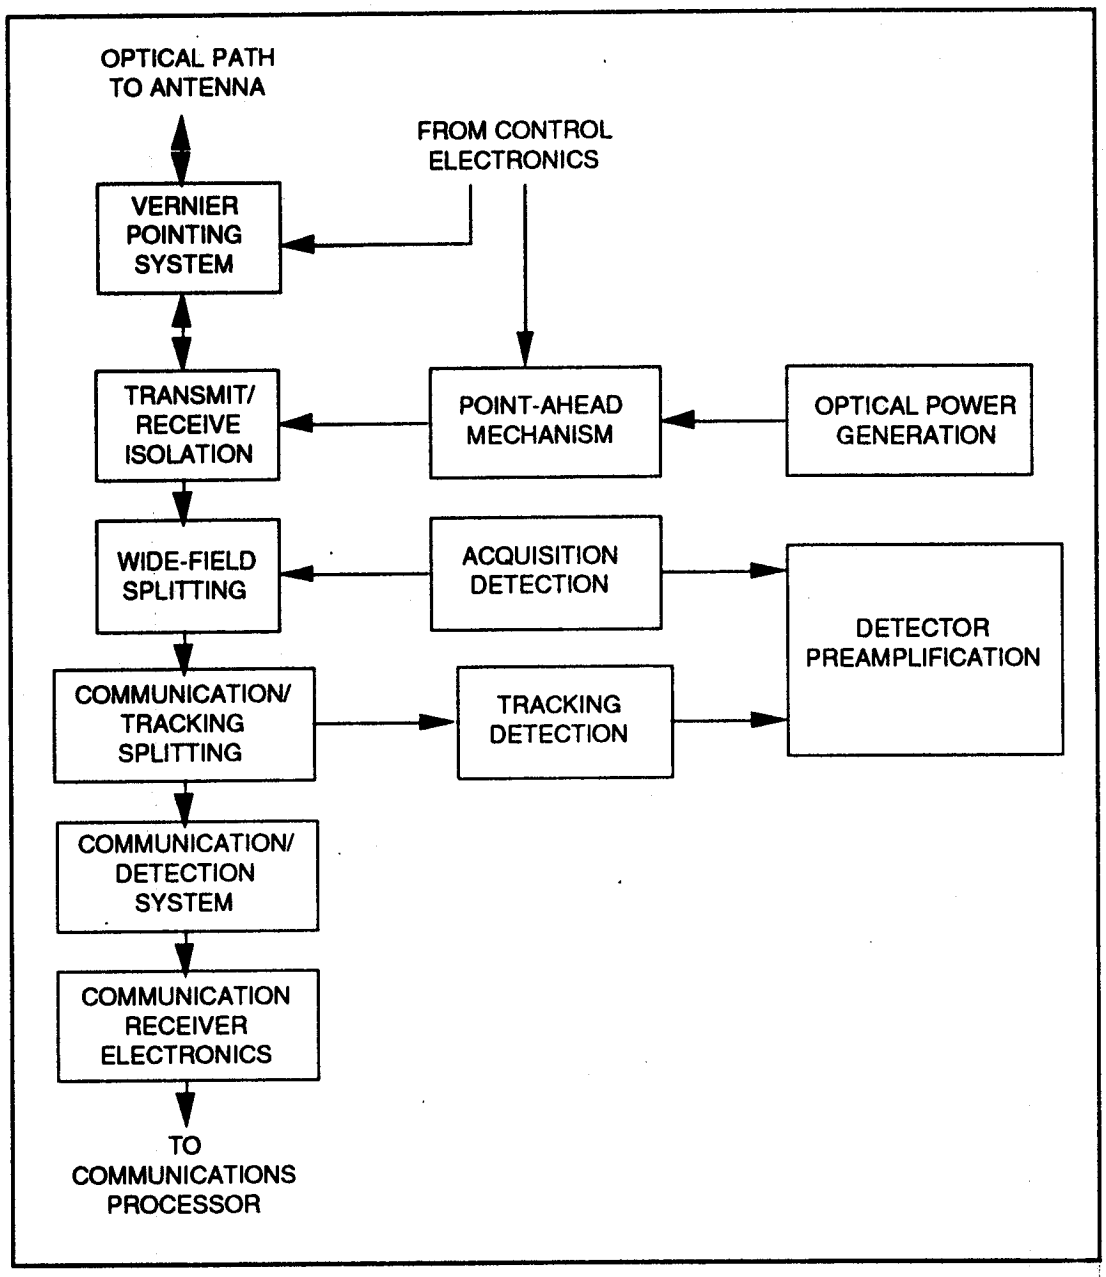

Figure B-7. Major Aspects of Focal Plane Assembly

4 RF/Optical Interface

The central issue revolving around RF/Optical interface is, how best to provide the required interconnectivity between various network elements on the ground stations. Two possibilities exist: (a) the satellite functioning only as a relay element, in which case the RF/Optical interface is straight-forward or (b) the satellite performs on-board processing, in which case the RF/Optical interface is considerably more complicated.

The primary design issues are network synchronization, switching and routing requirements, TDM frame formats and the data buffer require-ments. Combining multiple data sources, which is one of the primary chores of the RF/optical interface, requires special timing and control tech-niques. Network timing and synchronization can be accomplished via a variety of methods, such as using independent timing sources, a central timing reference, or a combination of both approaches. In addition to providing the configuration control information on command from the network control station located on earth, the control unit may also be re-quired to select appropriate RF transmitters or receivers supporting more than one satellite/ground stations.

The goals to select the modulation formats and coding schemes for the ground to GEO-satellite optical link mission must focus on flexibility and ease of implementation of the interface configuration. The choice of multi-plexing techniques offers the possibility of reducing the interface hardware requirements and complexity. Thus, major tradeoffs are made in above areas to gain in size, weight and prime power requirements for the optimal architecture.

5 Conclusion

In this particular ground to GEO-satellite link application, short wavelength GaAlAs technology appears to be promising for a variety of reasons:

- small payload size,
- high efficiency optical transmitter,
- direct modulation capability,
- wavelength selectability and tunability,
- high reliability,
- high level of research and development support by fiber optic industries and
- reasonable optical antenna/telescope diameter for the chosen transmitter power level and modulation/detection schemes.

However, a few critical technology areas and long lead time items which will require further development and performance/lifetesting to make the direct detection/OOK optical satellite links flightworthy include:

- optical power combining for high-power laser transmitters,
- static and dynamic stabilization of laser transmitters,
- device (high power laser; photoreceiver; narrow-bandwidth, sharp-cutoff optical filters; ICs, etc.) reliability in adverse and benign environments,

- high-speed ICs/logic chips for laser drivers and wide bandwidth optical receivers (for M-ary PPM transceivers), and on board signal processing (RF/optical interface) and

- focal plane design optimization and RF/optical interface.

It should be noted that the success of the ground-to-GEO satellite optical link depends on locating lasercom ground stations with minimal atmos-pheric degradation such that the specified link performance can be main-tained. This necessarily restricts the selection of ground terminal sites and operating wavelengths. Useful ground stations are few and far between, which in some cases may be distant from the data sources and receptors and require considerable terrestrial feed/distribution networking.

6 References

[B-1] "Microwave-Optical Interface Study", Final Report under contract with NASA GSFC, November 1984.

[B-2] "The Study of Optical Intersatellite Links", Final Report under contract with INTELSAT, August 1985.

[B-3] "TDAS Laser Intersatellite Communications Study", Final Report under subcontract with Ball Aerospace Systems Division/NASA GSFC, August 1986.

[B-4] "Optical Intersatellite Link Technology", Final Report under project support authorization STA-5 with COMSAT International Satellite Services Engineering, December 1986.

[B-5] "Architecture Study for Processing Payloads", Final Report under contract with INTELSAT, March 1989.

Appendix C

Satellite B-ISDN Traffic Analysis

1 Introduction

In this section, the impact of ATM traffic on the satellite B-ISDN is discussed. The input traffic characteristics, the statistical multiplexer buffer size, and the output link transmission rate are the major factors that affect cell transfer performance in the satellite B-ISDN system. Queueing models are built to analyze the cell transfer performance through statistical multiplexing at the earth station and space segment. The statistical multiplexer at the earth station and the on-board fast packet switch are assumed to be operated in slotted mode. Therefore, all the cells coming from different input lines are synchronous. Since the rates of the network node interfaces (NNIs) and user network interfaces (UNIs) were standardized for the B-ISDN, it is assumed that all the input lines at the earth station will have the same transmission rate. Two rates have been defined for NNIs and UNIs: one is 155.52 Mbit/s and the other 622.08 Mbit/s. It is assumed that a statistical multiplexer at the earth station interfaces at the 155.52 Mbit/s rate. In order to support a higher rate interface, a demultiplexer is needed to bring the 622.08 Mbit/s rate down to the 155.52 Mbit/s rate. Although the satellite transmission capacity is not constrained by terrestrial interfaces, for easy modeling the output transmission rate of the earth station is assumed to be an integer multiple of 155.52 Mbit/s. In the queueing model, data rates are normalized to 155.52 Mbit/s. The output transmission rate of the earth station will be an integer r, where $r \geq 1$. The slot time for rate 1 is equivalent to one cell transmission time at the input line, which is termed as unit time. Since the cell length is 53 bytes, one unit time is equal to 2.726 μs. At the earth station, cells coming to the multiplexer will be statistically multiplexed and stored in an output cell buffer. At the space segment, cells coming to the fast packet switch will be routed to the proper output ports and temporarily stored in the output buffers.

The ATM traffic can be characterized by the following parameters: peak bit rate, mean bit rate, mean peak duration, mean inactive duration, and burstiness. In the queueing models developed, a two-state Markovian chain associated with these parameters is used to describe the input traffic characteristics.

Intuitively, a cell loss ratio will be decreased by increasing the buffer size; however, cells will experience a longer transfer delay. Cell loss ratio and cell transfer delay will be

decreased if the output speed of the multiplexer is increased. However, if the output speed is set to too high, then the output link utilization will be lower, and the transmission capacity will not be utilized efficiently. There are trade-offs among these parameters to achieve the best performance of the satellite B-ISDN system. Queueing models are built to exploit the relationship among these parameters. The queueing models are assumed to be a stationary queue, i.e., the transient behavior of the queue is not considered. The satellite B-ISDN queueing model is depicted in Figure C-1. Queueing equations for the statistical multiplexer performance measurements, including buffer occupancy, cell loss ratio, cell transfer delay, cell delay jitter, and output link utilization, will be derived. This section will focus on the earth station queueing model, since the space segment model can be derived directly from the earth station model.

Statistical multiplexing and buffering (traffic smoothing) alone cannot achieve optimal efficiency of the satellite B-ISDN resource usage. Segregation of traffic based on priority is another important issue. To better exploit the efficiency of the network resources, different priorities can be assigned to different connections carrying different types of service. Within the ATM cell header, a one-bit field cell loss priority (CLP) can be used for this purpose. Whenever CLP is set (CLP Bit = 1) in a cell, this cell will be dropped first in case of network congestion. Utilizing the CLP field, the performance of the high-priority cells will be improved. If more than two priorities are used, the one-bit reserved field (RES) can also be utilized. It is anticipated that the improvement of cell loss ratio (CLR) will exceed the cell delay using the priority concept, because the performance of delay is dominated mostly by the buffer size. Although priority can help in guaranteeing the QOS for high-priority cells, it is hard to maintain the QOS for the low-priority cells. Nevertheless, this is the penalty of using cell priority.

The resource management functions can be divided into three categories: network control level (VCI/VPI routing and connection management), call control level (admission control and overload congestion control), and cell control level (usage parameter control and buffer management). The admission control, which is a decision to accept a new call based on the available network resources, will be discussed in this section. Two performance issues necessitate consideration when discussing the admission control. One is the blocking probability and the other is the cell level congestion related parameters: CLR, cell transfer delay (CTD), and cell delay jitter (CDJ). Only the cell level congestion related parameters will be considered in this section.

Although most analysis focuses on the earth station model, one performance measurement will be examined using the space segment model: the sensitivity of the on-board buffer with respect to the output transmission link. The object is to minimize the on-board buffer size while maintaining a certain channel efficiency. The overall CLR for the satellite system will also be discussed.

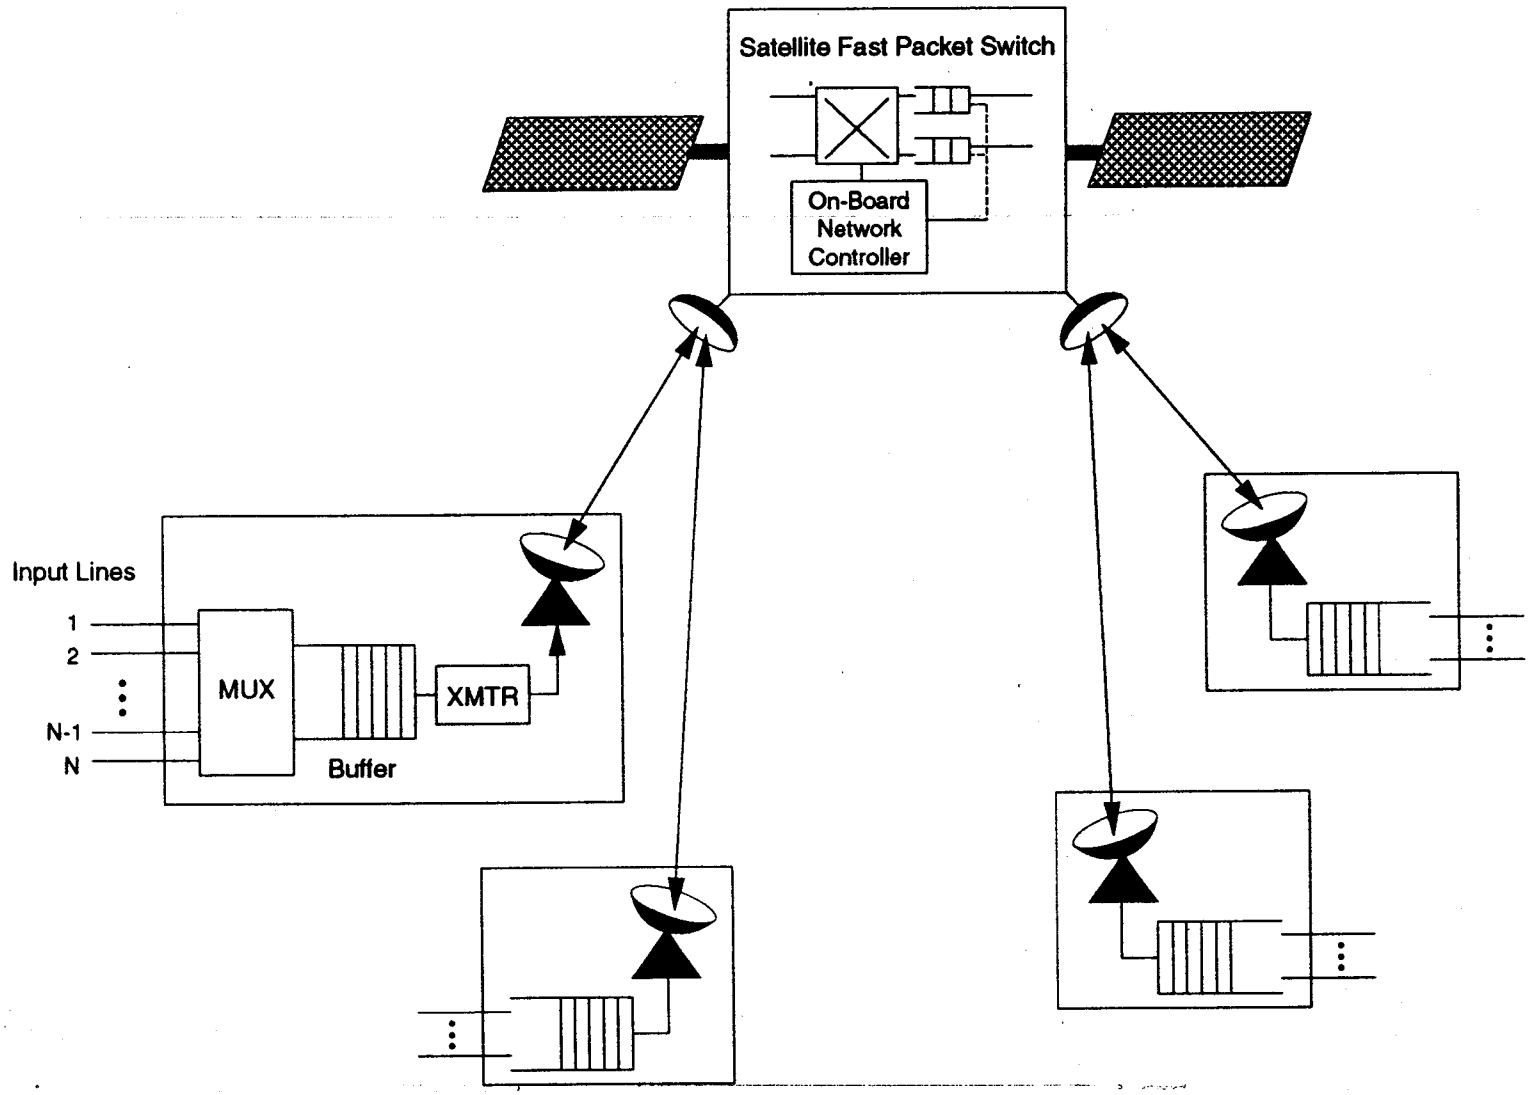

Figure C-1. Satellite B-ISDN Queueing Model

Four queueing models, as shown in Table C-1, will be developed for different priorities and buffer allocation schemes. Model A assumes that the incoming cells have no priority and the output queue is using the complete sharing scheme. Model B assumes that there are two priorities in the incoming cells. The high-priority cells will be put into the buffer first, followed by the low-priority cell. Model C is the same as the second model except there are two queues for two priority cells. Hence, the buffer storage is divided into two areas, which is essentially the complete partitioning (CP) scheme. Model D is also almost the same as the second model except the queue will always scan through the buffer and transmit the high-priority cells first. The buffer still uses the complete sharing scheme (CS), but the server does not use the first come first serve (FCFS) discipline. Numerical programs will be developed to solve these queueing equations. Results for different models, traffic scenarios, and design factors will be obtained and discussed.

Table C-1. Earth Station Model Alternatives

MODELS	BUFFER ALLOCATION	PRIORITY	QUEUE DISCIPLINE
Model A	Complete Sharing	No	FCFS
Model B	Complete Sharing	2	FCFS
Model C	Complete Partitioning	2	High Priority First
Model D	Complete Sharing	2	High Priority First

2 Notations for the Queueing Model

The statistical multiplexer at the earth station is modeled as a queueing system with a finite amount of storage of size b. Cells coming to the earth station will be stored in the buffer. This queue is assumed to be located at the output of the statistical multiplexer to simplify modeling. The customers of the queueing system are the cells arriving from different input lines. The server of the queue is the output line (output transmitter). For Model A, the queue is an FCFS system. This queueing model uses a synchronous cell arrival model, in which cells will arrive at the statistical multiplexer simultaneously based on slot timing.

The notation used in this section is summarized as follows.

- N The number of input lines interfaced with the earth station.
- P_A The probability that one input line is active with a mean value of T_A unit times.
- ρ_A The probability that one active input line will receive one information cell during one unit time.
- P_I The probability that one input line is inactive with a mean value T_I unit times.

ρ_I The probability that one inactive input line will receive one information cell during one unit time.

Q_1 The probability that one input line sends one information cell during one unit time.

Q_0 The probability that one input line sends one idle cell during one unit time.

$Q_{j(k)}$ The probability that input line k sends j information cells during one unit time.

e_i The probability that there are i information cells from input lines entering the statistical multiplexer during one unit time.

q_j The probability that the buffer queue will stay in state j (the buffer will have j cells).

B The burstiness of the input line.

b The buffer size at the output of the statistical multiplexer.

R The average number of cells entering into the system in one unit time.

r The transmission rate of the statistical multiplexer at the earth station, i.e., the output line of the statistical multiplexer at the earth station can transmit up to r cells in one unit time.

U The output link utilization.

3 Input Line Traffic Characterization

During one unit time, an input line can transmit either one empty cell or one information cell. For the bursty traffic, a two-state Markovian chain is used to describe its burstiness (see Figure C-2). At each unit time, each input line is either in an active state (A) or an inactive state (I). In the active state, the traffic source will generate information cells with rate ρ_A, and in the inactive state, the traffic source will generate information cells with rate ρ_I, where ρ_A is much greater than ρ_I. Burstiness is indicated by the ratio of the peak rate to the average rate and by the active state period. With the assumption of the two-state Markovian chain model, the periods of states A and I will be geometrically distributed (which is a memoryless process), with respective mean values T_A and T_I unit times. Let P_A (or P_I) be the probability that the input line is active (or inactive), then

$$P_A = \frac{T_A}{T_A + T_I} \quad \text{and} \quad P_I = \frac{T_I}{T_A + T_I} .$$

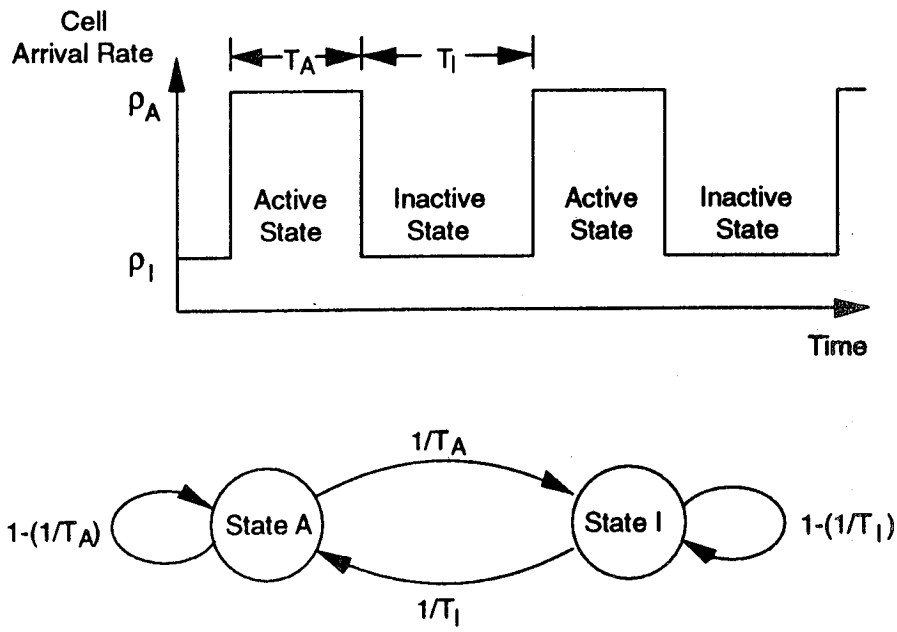

Figure C-2. Two-State Markovian Cell Arrival Model for Bursty Traffic

If Q_1 (or Q_0) denotes the probability that one input line sends one information cell (or one idle cell) during a unit time, then

$$\begin{cases} Q_1 = P_A\rho_A + P_I\,\rho_I \\ Q_0 = P_A\,(1 - \rho_A) + P_I\,(1 - \rho_I) \end{cases}$$

Burstiness of an input line is defined as the ratio of the peak rate to the average rate. The peak rate is ρ_A and the average rate is $P_A\,\rho_A + P_I\,\rho_I$. Hence the burstiness, B, is defined as

$$B = \frac{\rho_A}{P_A\rho_A + P_I\,\rho_I}.$$

If $\rho_I = 0$, then

$$B = \frac{\rho_A}{P_A\rho_A}.$$

It can be derived that

$$\frac{1}{T_I} = \frac{1}{T_A}\,\frac{1}{(B - 1)}$$

This equation depicts the relationship between the inactive/active periods and the burstiness, B, in Figure C-2.

Assume e_i is the probability that i cells arrive at the earth station in one unit time. Since there are N incoming lines to the earth station, N is the maximum number of cells that can arrive at the earth station in a unit time. Hence, the maximum value of i for e_i is N.

Assume $Q_{j(k)}$ is the probability that input line k sends j information cells during a unit time, where j is 1 or 0. The distribution of $Q_{j(k)}$ represents the characteristics of input line k. Then e_i can be derived as:

$$e_i = \sum_{j(1)+..+j(N)= i} \prod_{k=1}^{N} Q_{j(k)}$$

If these N lines all have the same characteristic, then $Q_{1(k)} = Q_1$ and $Q_{0(k)} = Q_0$. Therefore, e_i can be simplified as

$$e_i = \binom{N}{i} Q_1^{\,i} Q_0^{\,N-i}, \quad \text{where } i \le N.$$

4 Performance Measurements for Model A

Three cell transfer performance measurements will be discussed: CLR, CTD and CDJ. The bandwidth efficiency indication, output utilization, will also be derived.

4.1 Buffer Occupancy

Assume b is the buffer size of the output queue, and q_j is the probability that there are j cells waiting in the queue. It is assumed that when there is no cell in the buffer, the cell can flow through the buffer without any delay.

For $r \le b$, the balance equations for q_j can be written as

$$\begin{cases} \sum_{i=0}^{r} q_i \left[\sum_{k=0}^{r-i} e_k \right] = q_0 \\ \sum_{i=0}^{\mathrm{Min}(r+j,b)} q_i \, e_{(r+j-i)} = q_j \qquad \text{for } 1 \le j < b \text{ and } r+j-i \le N \\ \sum_{i=0}^{b} q_i \left[\sum_{k=r+b-i}^{N} e_k \right] = q_b \end{cases}$$

where one of the above equations is redundant.

The discrete-time Markovian chain state transition diagram for the output queue is shown in Figure C-3.

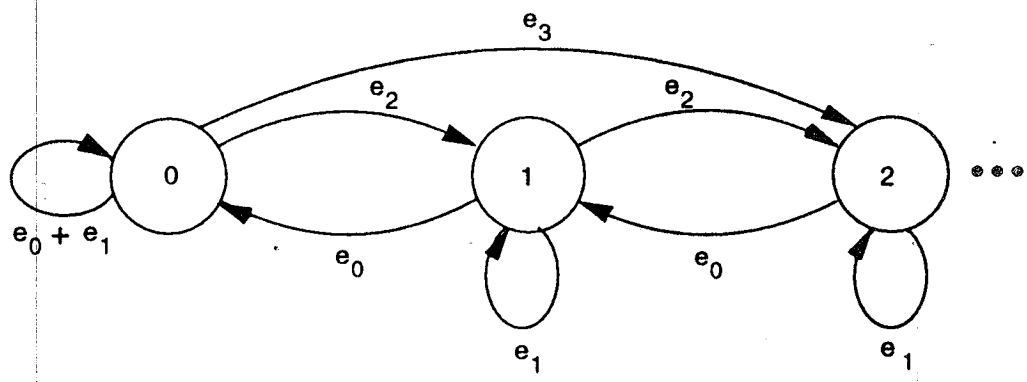

Figure C-3. Discrete-Time Markovian Chain State Transition for the Buffer Occupancy

For $r > b$, the balance equations for q_j can be written as

$$\begin{cases} \sum_{i=0}^{b} q_i \left[\sum_{k=0}^{r-i} e_k \right] = q_0 \\ \sum_{i=0}^{b} q_i \, e_{(r+j-i)} = q_j \qquad \text{for } 1 \le j < b \text{ and } r+j-i \le N \\ \sum_{i=0}^{b} q_i \left[\sum_{k=r+b-i}^{N} e_k \right] = q_b \end{cases}$$

where one of the above equations is redundant.

And by conservation of probability,

$$\sum_{j=0}^{b} q_j = 1$$

The above system comprises b+1 linear equations and b+1 unknowns. A numerical method can be used to solve the above equations.

The mean buffer length E[q] can be derived as

$$E[q] = \sum_{j=0}^{b} q_j j$$

The buffer length standard deviation σ_q^2 can be derived as

$$\sigma_q^2 = \sum_{j=0}^{b} q_j j^2 - E^2[q]$$

4.2 Cell Loss Ratio

The CLR is defined as the number of lost cells to the sum of the number of lost and successfully delivered cells. The cells will not be lost unless the buffer overflows. The buffer will overflow if the number of cells waiting in the buffer plus the number of arriving cells exceeds the sum of the transmission rate and the size of the buffer. That is, the buffer will overflow if $j + i > r + b$. Given a particular queue state j, the conditional probability that the buffer will overflow can be derived as

$$\sum_{i=r+b-j+1}^{N} e_i$$

Hence, the average number of lost cells, given that the queue is in state j, is given as

$$\sum_{i=r+b-j+1}^{N} e_i \ (j + i - r - b)$$

Finally, the number of lost cells during one unit time can be written as

$$\sum_{j=0}^{b} q_j \left[\sum_{i=r+b-j+1}^{N} e_i \, (j + i - r - b) \right]$$

The cell loss ratio, CLR, is defined as the ratio of the average number of lost cells during a unit time to the average number of cells entering the system in a unit time. The average number of cells entering the system in a unit time, R, is given as

$$R = \sum_{i=0}^{N} i \, e_i$$

Thus, the cell loss ratio can be expressed as

$$CLR = \frac{\displaystyle\sum_{j=0}^{b} q_j \left[\sum_{i=r+b-j+1}^{N} e_i \, (j+i-r-b) \right]}{R}$$

4.3 Cell Transfer Delay and Cell Delay Jitter

The cell transfer delay consists of cell waiting time in the buffer and cell service time (cell transmission time). The cell transfer delay will be computed first. Assume that the number of cells in the buffer is j. Whenever there are i cells coming from the input line, these i cells will be put at the back of these j cells. To simplify the queueing model, the multiplexer randomly puts the incoming cells into the buffer each time. Assume the probability that a cell-of-interest will wait for k cell times before being transmitted is Pr[K = k], where a cell time is defined as the amount of time needed for the output line to transmit one cell. According to this definition, one unit time equals r cell times. After the arrival of i cells, the total number of cells in the buffer is i+j. This cell-of-interest will be in position k+1 of all the i+j cells in the buffer (see Figure C-4).

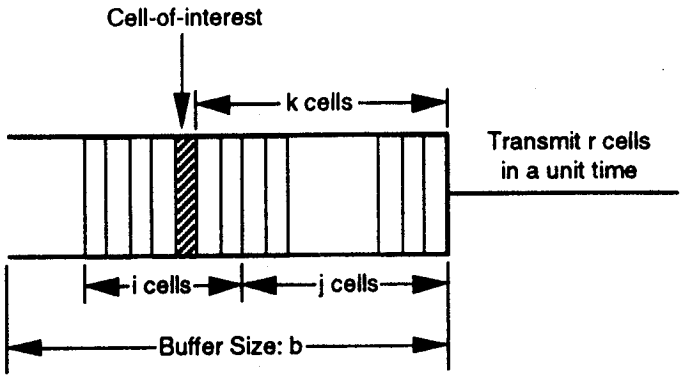

Figure C-4. Cell Transfer Delay for Cell k+1

$Pr[K=k \mid q_j]$ is the probability that the cell is in position k+1 of the buffer, given that there are j cells in the buffer, where k ranges from 0 to r+b-1 and $k \geq j$. The ranges of k are derived as follows.

By definition, the position of the cell-of-interest is greater than j. Hence, when k is less than j, $Pr[K=k \mid q_j]$ is equal to 0. Cell waiting time only applies to the non-lost cells; hence, k is always less than or equal to r+b-1. When k is greater than r+b-1, $Pr[K=k \mid q_j]$ is equal to 0.

Now, consider that k is between r + b - 1 and j, i.e., $r + b - 1 \geq k \geq j$. Define $P_{batch(i \mid j)}$ to be the probability that a batch of cells with size i has come to the cell buffer when the buffer is in state j. Define $P_{pos(i \mid j)}$ to be the probability that the cell-of-interest is in one particular position in this batch, given the buffer is in state j. Then, $Pr[K=k \mid q_j]$ can be expressed as

$$Pr[K=k \mid q_j] = \sum_{i=k-j+1}^{N} P_{batch(i \mid j)} P_{pos(i \mid j)}, \text{ where } r + b - 1 \geq k \geq j.$$

In order to calculate $P_{batch(i \mid j)}$, the average size of the batch given that the buffer is in state j, $E[e_i \mid q_j]$, has to be known first. With infinite buffer size, there is no cell loss; hence, $E[e_i \mid q_j] = E[e_i]$. With a finite buffer size, the average batch size depends on the queue state. When the queue is in state j, the batch size will be i when $i \leq b + r - j$ and will be $b + r - j$ when $i > b + r - j$. Hence, $E[e_i \mid q_j]$ can be expressed as

$$E[e_i \mid q_j] = \sum_{i=0}^{b+r-j} i\, e_i + \sum_{i=b+r-j+1}^{N} (b+r-j)\, e_i$$

Probabilities $P_{batch(i \mid j)}$ and $P_{pos(i \mid j)}$ can be computed in two different cases based on the value of i.

a. If $i \le b + r - j$, then $P_{batch(i|j)} = \frac{i\, e_i}{E[e_i \mid q_j]}$ and $P_{pos(i|j)} = \frac{1}{i}$

b. If $i > b + r - j$, then $P_{batch(i|j)} = \frac{(b+r-j)\, e_i}{E[e_i \mid q_j]}$ and $P_{pos(i|j)} = \frac{1}{b+r-j}$

Therefore,

$$Pr[K=k \mid q_j] = \frac{\sum_{i=k-j+1}^{N} e_i}{\sum_{i=0}^{b+r-j} i\, e_i + \sum_{i=b+r-j+1}^{N} (b+r-j)\, e_i}$$

It can be verified that $\sum_{k=0}^{b+r-1} Pr[K=k \mid q_j] = 1$.

Then,

$$Pr[K = k] = \sum_{j=0}^{b} q_j\, Pr[K=k \mid q_j]$$

E[K] and σ_K^2 can be derived as follows:

$$E[K] = \sum_{k=0}^{r+b-1} Pr[K = k]\, k$$

$$\sigma_K^2 = \sum_{k=0}^{r+b-1} Pr[K = k]\, k^2 - E^2[K]$$

The cell transfer delay (CTD) in unit time is the sum of the cell waiting time and the cell transmission time. The cell transmission time is only one cell time. Hence,

$$CTD = (E[K] + 1) / r$$

Since the cell transmission time is a constant, the cell delay jitter (CDJ) in unit time is the variance of the waiting time σ_K. Hence,

$$CDJ = \sigma_K / r$$

4.4 Output Link Utilization

Let t_j be the probability that the output line transmits j cells during a unit time, where j ranges from 0 to r. By this definition, t_j can be written as

$$t_j = \sum_{k=Max(0,j-b)}^{j} e_k \, q_{j-k} \qquad \text{for } 0 \le j \le r-1$$

$$t_r = \sum_{k=0}^{r-1} q_k \left[\sum_{i=r-k}^{N} e_i \right] + \sum_{k=r}^{b} q_k \left[\sum_{i=0}^{N} e_i \right] \quad \text{if } b \ge r .$$

Since $\sum_{i=0}^{N} e_i = 1$, the above equation can be simplified as

$$t_r = \sum_{k=0}^{r-1} q_k \left[\sum_{i=r-k}^{N} e_i \right] + \sum_{k=r}^{b} q_k \quad \text{if } b \ge r .$$

$$t_r = \sum_{k=0}^{b} q_k \left[\sum_{i=r-k}^{N} e_i \right] \quad \text{if } b < r$$

The output average utilization of the statistical multiplexer at the earth station is defined as the ratio of the average number of cells transmitted by the output line to the output rate r. Hence, the output link utilization (U) is written as

$$U = \frac{\sum_{j=0}^{r} t_j \, j}{r}, \text{ where } 0 \le j \le r.$$

For r=1, the U can be simplified as follows.

$$U = (1 - e_0 \, q_0)$$

The utilization can also be expressed as the average arrival rate times the average service time. Therefore, another way of calculating U is

$$U = R(1\text{-}CLR)\frac{1}{r}$$

5 Performance Measurements for Model B

In Model B, the cell loss priority (CLP) in the ATM cell will be activated to improve the performance for delay-sensitive and loss-sensitive traffic. In ATM cells, there is 1 bit available for CLP. A cell with CLP=0 (not set) has a higher priority than a cell with a CLP=1. This cell loss priority for a given virtual channel connection (VCI) should remain constant for the duration of the virtual connection so that no missequencing will happen. The multiplexer will always put the incoming high-priority cells into the buffer first, followed by the low-priority cells. The sequence of the cells already in the buffer will not be changed. In case of buffer overflow, incoming low-priority cells are dropped first, followed by the high-priority cells. After the cell enters the buffer, all cells will be served in an FCFS manner. This queueing model is shown in Figure C-5. The queueing equations for buffer occupancy and output link utilization of this model are essentially the same as in Model A. The cell loss ratio, cell transfer delay and cell delay jitter for the cells with different priorities are the main focus. The improvement of CLR, CTD, and CDJ for the high-priority cells is compared with Model A to examine the effects of using the priority concept.

5.1 Notations

$e(L)_l$ The probability that there are l low-priority information cells entering the statistical multiplexer during one unit time.

$e(H)_h$ The probability that there are h low-priority information cells entering the statistical multiplexer during one unit time.

N(L) The maximum number of low-priority cells entering the statistical multiplexer during one unit time.

N(H) The maximum number of high-priority cells entering the statistical multiplexer during one unit time.

R(H) The average number of high-priority cells entering the system in one unit time.

R(L) The average number of low-priority cells entering the system in one unit time.

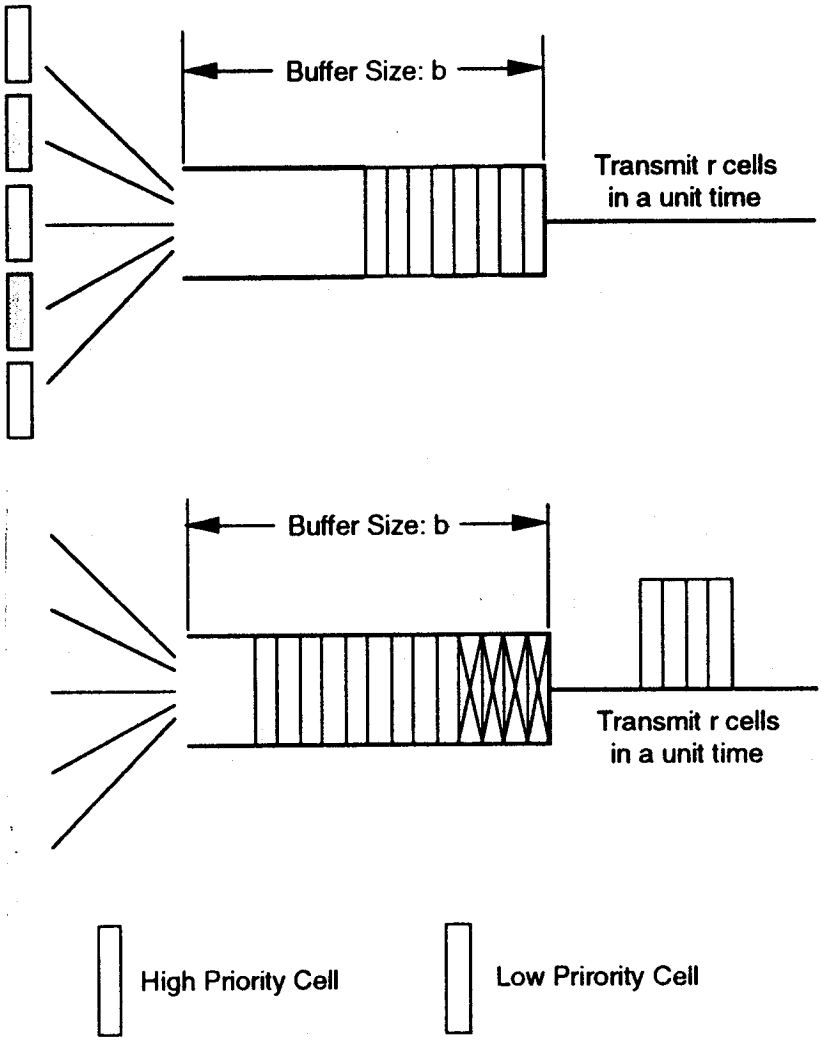

Figure C-5. Cell Priority Concept in Model B

5.2 Cell Loss Ratio

The cell loss ratio will be calculated for high-priority cells and low-priority cells separately.

5.2.1 High-Priority Cells

The high-priority cells will be lost if the number of high-priority cells plus the number cells waiting in the buffer (which contains high- and low-priority cells) exceed the sum of the transmission rate and the size of the buffer. Hence, the average number of lost high-priority cells in a unit time can be written as

$$\sum_{j=0}^{b} q_j \sum_{h=0}^{N(H)} e(H)_h \ \text{Max}((h + j - r - b) , 0)$$

The average number of high-priority cells entering the system in a unit time is

$$R(H) = \sum_{h=0}^{N(H)} e(H)_h \, h$$

Thus the cell loss ratio for high-priority cells (CLRH) can be expressed as

$$CLRH = \frac{\sum_{j=0}^{b} q_j \sum_{h=0}^{N(H)} e(H)_h \ \text{Max}((h + j - r - b) , 0)}{R(H)}$$

5.2.2 Low-Priority Cells

The low-priority cells will be lost if the number of low-priority cells plus the number of high-priority cells and the number of cells in the buffer exceed the sum of the transmission rate and the size of the buffer. The average number of lost low-priority cells in a unit time can be written as

$$\sum_{j=0}^{b} q_j \sum_{l=0}^{N(L)} \sum_{h=0}^{N(H)} e(L)_l \, e(H)_h \, \text{Max} (\text{Min}((l + h + j - r - b) , l), 0)$$

The average number of low-priority cells entering the system in a unit time is

$$R(L) = \sum_{l=0}^{N(L)} e(L)_l \, l$$

Thus, the cell loss ratio for the low-priority cells can be expressed as

$$CLRL = \frac{\sum_{j=0}^{b} q_j \sum_{l=0}^{N(L)} \sum_{h=0}^{N(H)} e(L)_l \, e(H)_h \, \text{Max}(\text{Min}((l + h + j - r - b) , l), 0)}{R(L)}$$

There is an easy way of calculating the CLRL. The basic concept is that the total number of lost cells is the same as the sum of the number of lost high-priority cells and the number of lost low-priority cells. Hence, the relationship among CLR, CLRH, and CLRL is expressed as follows.

$$R * CLR = RH * CLRH + RL * CLRL$$

Remember that R is the average number of cells entering the system in a unit time.

$$R = \sum_{i=0}^{N} i\, e_i$$

$$RH = \sum_{i=0}^{N(H)} h\, e(H)_h$$

$$RL = \sum_{i=0}^{N(L)} l\, e(L)_l$$

5.3 Cell Transfer Delay and Cell Delay Jitter

The cell transfer delay and cell delay jitter will be discussed separately for high-priority cells and low-priority cells.

5.3.1 High-Priority Cells

Assume the number of cells in the buffer is j. Whenever there are h high-priority cells coming from the input line, these h cells will be put at the back of these j cells. Hence, computation of $Pr[K=k \mid q_j]$, which is the probability that the high-priority cell will wait k cell times given there are j cells in the buffer, is essentially the same as in Model A. Again the cell waiting time will only apply to the non-lost high-priority cells; hence, k only can go up to r + b - 1.

Hence, when k is greater then r + b - 1, $Pr[K=k \mid q_j]$ is equal to 0.

By definition, when k is less than j, $Pr[K=k \mid q_j]$ is equal to 0.

Now consider $r + b - 1 \geq k \geq j$. Following the same arguments as in Model A, the probability that the cell-of-interest will be in position k + 1 among all the cells in the buffer, given the buffer is in state j, is derived as

$$Pr[K=k \mid q_j] = \frac{\sum_{h=k-j+1}^{N} e(H)_h}{\sum_{h=0}^{b+r-j} h\, e(H)_h + \sum_{h=b+r-j+1}^{N(H)} (b+r-j)\, e(H)_h}$$

After $Pr[K=k \mid q_j]$ is obtained, the corresponding $E[K]$ and σ_K^2 can also be computed.

5.3.2 Low-Priority Cells

Assume the number of cells in the buffer is j. Whenever there are h high-priority cells and l low-priority cells coming from the input lines, these h cells will be put at the back of the j cells followed by the l cells. The objective is to find the cell delay within these l cells. Hence, this puts a constraint on the values of k. The value of k has to be greater than or equal to j + h, and less than j + h + l. Again the cell waiting time will only apply to the non-lost low-priority cells; hence, k only can go up to r+b-1.

Define a new variable s = j + h, where j is the number of cells in the buffers and h is the number of incoming high-priority cells. Since only non-lost low-priority cells are considered, at least one cell buffer should exist for the low-priority cell-of-interest; hence, the values of s are limited between 0 and r+b-1. Define an event $\Psi = \{ 0 \le s \le r + b - 1\}$. Define a new probability $Pr[K=k \mid (s \mid \Psi)]$ which is the probability that the cell-of-interest will be in position k + 1 of all the possible r + b positions given $Pr[(s \mid \Psi)]$. Since only the non-lost low-priority cells are of interest, the variable, k+1, of the cell-of-interest can only go up to r + b.

Hence, when k is greater than r + b - 1, $Pr[K=k \mid (s \mid \Psi)]$ is equal to 0.

By definition, when k is less than s, $Pr[K=k \mid (s \mid \Psi)]$ is equal to 0.

Following the same principle as in Model A,

$$Pr[K=k \mid (s \mid \Psi)] = \sum_{l=k-s+1}^{N(L)} P_{batch(l \mid j;h)} \, P_{pos(l \mid j;h)}, \text{ where } r+b-1 \ge k \ge s.$$

In order to calculate $P_{batch(l \mid j;h)}$, the average size of the batch, $E[e(L)_l \mid \Psi]$, given that the buffer is in state j and the number of arriving high-priority cells is h, must be known first. When the queue is in state j and the number of arriving high-priority cells is h, the batch size will be l when $l \le b + r - s$ and will be $b + r - s$ when $l > b + r - s$. Hence, $E[e(L)_l \mid \Psi]$ can be expressed as

$$E[e(L)_l \mid \Psi] = \sum_{l=0}^{b+r-s} l \, e(L)_l + \sum_{l=b+r-s+1}^{N(L)} (b+r-s) \, e(L)_l$$

The $P_{batch(l \mid j;h)}$ and $P_{pos(l \mid j;h)}$ can be computed in two different cases based on the value of l.

a. If $l \le b+r-s$ then $P_{batch(l \mid j;h)} = \dfrac{l \, e(L)_l}{E[e(L)_l \mid \Psi]}$ and $P_{pos(l \mid j;h)} = \dfrac{1}{l}$

b. If $l > b+r-s$, then $P_{batch(l|j;h)} = \frac{(b+r-s)\, e(L)_l}{E[e(L)_l \mid \Psi]}$ and $P_{pos(l|j;h)} = \frac{1}{b+r-s}$

Therefore,

Following the same procedure as in Model A, $Pr[K=k \mid (s \mid \Psi)]$ can be derived as

$$Pr[K=k \mid (s \mid \Psi)] = \frac{\sum_{l=f-S+1}^{N(L)} e(L)_l}{\sum_{l=0}^{b+r-s} l\, e(L)_l + \sum_{l=b+r-s+1}^{N(L)} (b+r-s)\, e(L)_l}$$

The $Pr[s \mid \Psi]$ can be derived as follows.

$$Pr[s \mid \Psi] = \frac{\sum_{j+h=s} q_j e(H)_h}{\sum_{s=0}^{b+r-1} \sum_{j+h=s} q_j e(H)_h}$$

Now the probability that the cell-of-interest will be in position k + 1 of all the r + b possible positions is

$$Pr[K=k] = \sum_{s=0}^{b+r-1} Pr[s \mid \Psi]\, Pr[K=k \mid (s \mid \Psi)]$$

6 Performance Measurements for Model C

In Model C, the ATM cell loss priority will also be activated. However, there are two output queues in the statistical multiplexer instead of just one. One output queue is dedicated to the high-priority queue while the other is dedicated to the low-priority queue. The output server will always transmit the cells in the priority queue first. Whenever there are no cells to transmit in the high-priority queue, the low-priority queue will be served. This queueing model is depicted in Figure C-6.

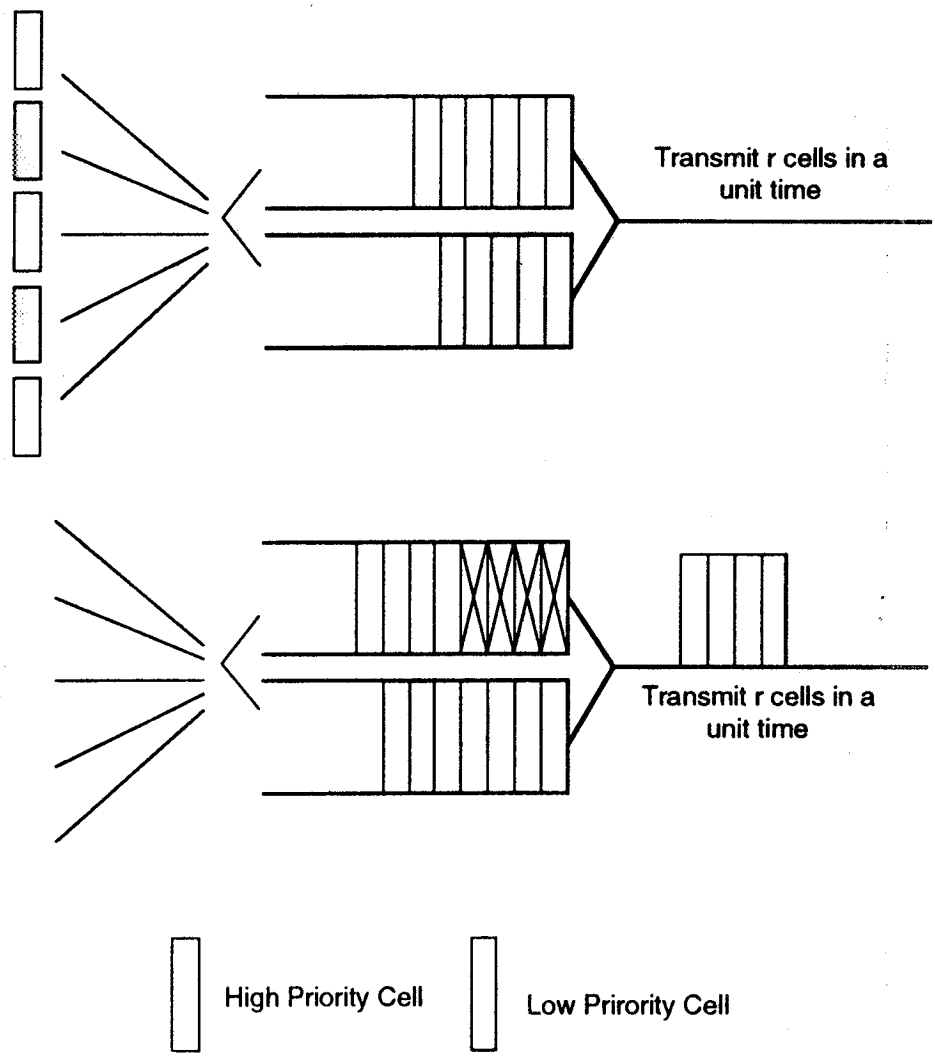

Figure C-6. Cell Priority Concept in Model C

6.1 Notations

The notation introduced in Model B will not be repeated here.

b(L) The buffer size for the low-priority cells.

b(H) The buffer size for the high-priority cells.

$q(L)_j$ The probability that the buffer queue assigned to low-priority cells will stay in state j (the high-priority buffer will have j cells).

$q(H)_j$ The probability that the buffer queue assigned to high-priority cells will stay in state j (the high-priority buffer will have j cells).

6.2 High-Priority Cells

Since the output transmitter will always serve the high-priority cells first, the processing of high-priority cells is independent of the processing of low-priority cells. Hence, the queueing equations for buffer occupancy, cell loss ratio, cell delay, and output link utilization are the same as those in Model A.

6.3 Low-Priority Cells

Processing of low-priority cells depends on the state of the high-priority queue and the high-priority cell arrival rate; hence, the queueing equations for buffer occupancy, cell loss ratio, cell delay, and output link utilization will be recalculated.

6.3.1 Buffer Occupancy

Assume b(L) is the buffer size assigned to the low-priority cells, and $q(L)_j$ is the probability that there are j low-priority cells waiting in the buffer. Remember $t(H)_j$ is the probability that the output line transmits j high-priority cells during one unit time.

For r < b(L), the balance equations for $q(L)_j$ can be written as

$$\begin{cases} \displaystyle\sum_{i=0}^{r} q(L)_i \sum_{z=i}^{r} t(H)_{r-z} \left[\sum_{k=0}^{\mathrm{Min}(z-i,N(L))} e(L)_k \right] = q(L)_0 \\ \displaystyle\sum_{i=0}^{\mathrm{Min}(r+j,b(L))} q(L)_i \sum_{z=\mathrm{Max}(0,i-j)}^{\mathrm{Min}(r,i-j+N(L))} t(H)_{r-z}\, e(L)_{z+j-i} = q(L)_j \quad \text{for } 1 \le j \le b(L)-1 \text{ and } i-j<r \end{cases}$$

For r > b(L), the balance equations for $q(L)_j$ can be written as

$$\begin{cases} \displaystyle\sum_{i=0}^{b(L)} q(L)_i \sum_{z=i}^{r} t(H)_{r-z} \left[\sum_{k=0}^{\mathrm{Min}(z-i,N(L))} e(L)_k \right] = q(L)_0 \\ \displaystyle\sum_{i=0}^{b(L)} q(L)_i \sum_{z=\mathrm{Max}(0,i-j)}^{\mathrm{Min}(r,i-j+N(L))} t(H)_{r-z}\, e(L)_{z+j-i} = q(L)_j \quad \text{for } 1 \le j \le b(L)-1 \text{ and } i-j<r \end{cases}$$

And by conservation of probability,

$$\sum_{j=0}^{b(L)} q(L)_j = 1$$

The above system consists of b(L)+1 equations and b(L)+1 unknowns.

6.3.2 Cell Loss Ratio

The cell loss ratio for low-priority cells is:

$$CLRL = \frac{\sum_{j=0}^{b(L)} q(L)_j \sum_{z=0}^{r} t(H)_{r-z} \{ \sum_{k=b(L)+z-j+1}^{N(L)} e(L)_k [k+j-z-b(L)] \}}{\sum_{i=0}^{N(L)} i\, e(L)_i}$$

6.3.3 Cell Transfer Delay and Cell Delay Jitter

Assume the number of cells in the low-priority buffer is j. Whenever there are l low-priority cells coming from the input lines, these l cells will be put at the back of the j cells. The objective is to find the cell delay within these l cells. The difference in the derivation of cell delay in this model compared to the previous two is that the number of low-priority cells which can be transmitted in one unit time is not a constant r. Instead, this number depends on the number of cells being transmitted in the high-priority queue. Assume the number of low-priority cells which can be transmitted in a unit time is z. This implies the number of high-priority cells being transmitted in a unit time is r - z. Define a new probability $Pr[F=f \mid q(L)_j ; z]$, which is the probability that the cell-of-interest will be in position f + 1 of all the possible z + b(L) positions. Since only the non-lost low-priority cells are interested, the variable, f + 1, of the cell-of-interest can only go up to z + b(L).

Hence, when f is greater than z+b(L)-1, $Pr[F=f \mid q(L)_j ; z]$ is equal to 0.

By definition, when f is less than j, $Pr[F=f \mid q(L)_j ; z]$ is equal to 0.

Following the same procedure as in Model A or Model B, $Pr[F=f \mid q(L)_j ; z]$ can be derived as

$$Pr[F=f \mid q(L)_j; z] = \frac{\sum_{l=f-j+1}^{N(L)} e(L)_l}{\sum_{l=0}^{b(L)+z-j} l\, e(L)_l + \sum_{l=b(L)+z-j+1}^{N(L)} (b(L)+z-j)\, e(L)_l}$$

Now the probability that the cell-of-interest will be in position f + 1 of all the b(L) + z possible positions is

$$Pr[F=f] = \sum_{j=0}^{b(L)} \sum_{z=0}^{r} t(H)_{r-z} \, Pr[F=f \mid q(L)_j ; z]$$

Define another probability Pr[K=k | F=f], which is the probability that the cell-of-interest will have to wait for k cell time before being transmitted, given that this cell is in position f + 1 at the current unit time. The computation of Pr[K=k | F=f] depends on the value of r - z, which is the total number of high-priority cells being transmitted in this unit time. If f + 1 ≤ z, then this cell-of-interest can be transmitted in this unit time; otherwise, this cell has to wait at least until the next unit time to be transmitted. Hence, there are two cases based on the value of z. Assume the value of z at this unit time is z(1).

a. f+1 ≤ z(1). This cell-of-interest will be transmitted during this unit time. The cell waiting time is f + r - z(1), where r - z(1) is the transmission time of the high-priority cells.

b. f+1 > z(1). This cell-of-interest will not be transmitted at this unit time. However, the position of this cell will be advanced by z(1). Hence, the cell-of-interest will be in position f+1-z(1) at the next unit time. The cell waiting time at this instant is the same as the unit time, which is r cell time.

Consider the next unit time:

b1. Assume f + 1 ≤ z(1) + z(2). This cell-of-interest will be transmitted during this unit time. The total cell waiting time is r + f - z(1) + r - z(2).

b2. Assume f + 1 < z(1) + z(2). This cell will not be transmitted at this unit time. This cell will be in position f + 1 -z(1) - z(2). The total waiting time up to this instant is 2r.

By recursively following this procedure, the cell waiting time can be calculated. It is possible that the cell waiting time will be infinite; however, in the computation, an assumption is made that the probability for this event to occur is very small. Hence, this event will be ignored.

6.3.4 Output Link Utilization

Let t_j be the probability that the output line transmits j cells during one unit time, where j ranges from 0 to r. These j cells consists of high- and low-priority cells.

$$t_j = \sum_{z=r-j}^{r} t(H)_{r-z}\, t(L)_{j+z-r}$$

where $t(L)_j$ can be obtained from the following equations.

$$t(L)_j = \sum_{k=Max(0,j-b(L))}^{j} e(L)_k\, q(L)_{j-k} \qquad \text{for } 0 \le j \le r-1$$

$$t(L)_r = \sum_{k=0}^{r-1} q(L)_k \left[\sum_{i=r-k}^{N(L)} e(L)_i \right] + \sum_{k=r}^{b(L)} q(L)_k \left[\sum_{i=0}^{N(L)} e(L)_i \right] \quad \text{if } b(L) \ge r$$

$$t(L)_r = \sum_{k=0}^{b(L)} q(L)_k \left[\sum_{i=r-k}^{N(L)} e(L)_i \right] \quad \text{if } b(L) < r$$

Hence, the output link utilization (U) is written as

$$U = \frac{\sum_{j=0}^{r} j\, t_j}{r}, \text{ where } 0 \le j \le r.$$

7 Performance Measurements for Model D

In Model D, the ATM cell loss priority will also be activated. There is one output queue in the statistical multiplexer. However, the queue discipline is not FCFS. Instead, the high-priority cells will always be inserted in front of the low-priority cells. A pointer concept is used to implement this queueing system so that insertion of cells can be done easily. In this sense, the queue has been divided into two logical separated areas, and the boundary between two queues is movable. The output server will serve the cells in the queue sequentially. A non-preemptive queue is assumed so that cells being served cannot be interrupted. This queueing model is depicted in Figure C-7.

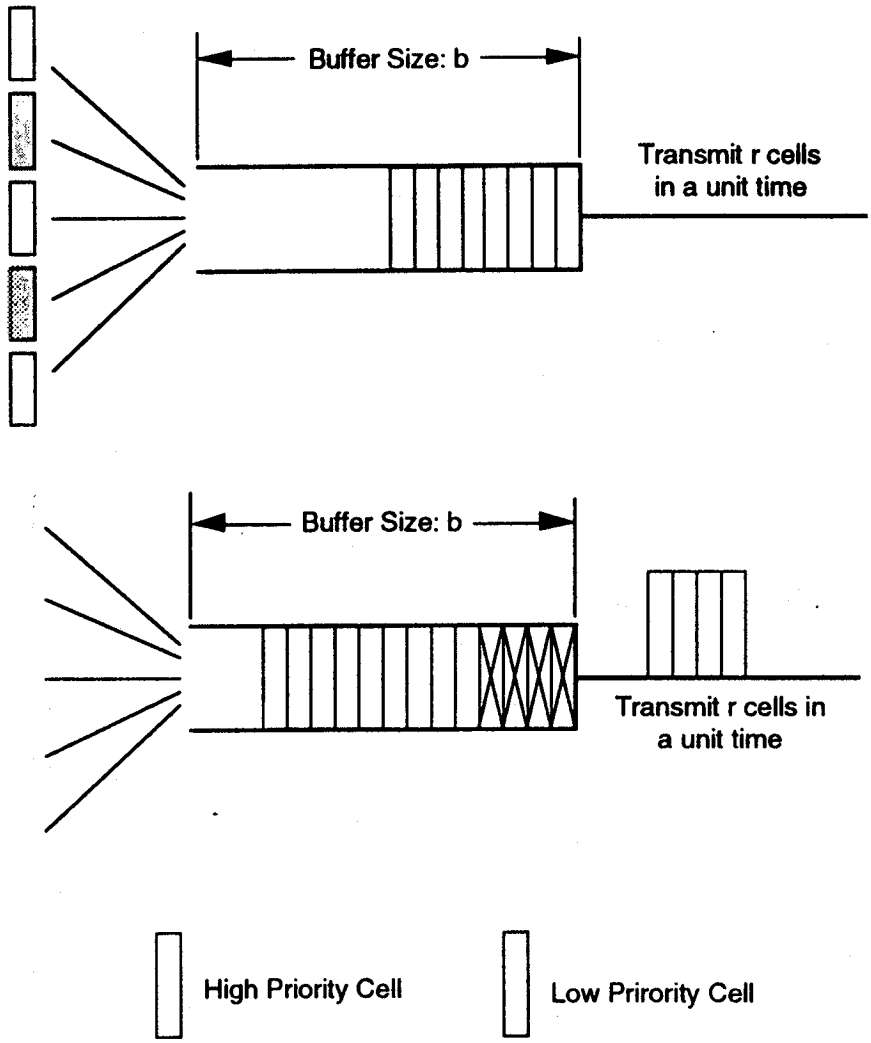

Figure C-7. Cell Priority Concept in Model D

7.1 Notations

q_j The probability that there are j cells in the buffer.

$q_{l'}$ The probability that there are l low-priority cells in the buffer.

q_h The probability that there are h high-priority cells in the buffer.

7.2 Buffer Occupancy

Let q_j, $q_{h'}$, and $q_{l'}$ be the corresponding probability distributions of buffer occupancy for the total number of cells, the number of high-priority cells, and the number of low-priority cells, respectively. The computation of $q_{h'}$ is exactly the same as in Model A

since the high-priority cells will not be affected by the low-priority cells. The computation of q_j is also the same as in Model A since priority will not affect buffer occupancy of the total number of cells. There are two ways of computing $q_{l'}$. The first option is to use the balance equations. The second option is to take advantage of the fact that the values of q_j and $q_{h'}$ are known; hence, $q_{l'}$ may be calculated from these values. The second option is adopted.

An example of buffer size 3 is used to illustrate the procedure of computing $q_{l'}$. Next, a general procedure will be introduced.

Use $q_{[0^X L^Y H^Z]}$ to represent the probability that there are X empty buffers, Y low-priority cells, and Z high-priority cells, where X+Y+Z = b. All possible combinations of X,Y, and Z for buffer size 3 are:

q[03], q[02H], q[02L], q[0H2], q[0LH], q[0L2], q[H3], q[LH2], q[L2H], and q[L3].

Writing down the equations for q_j and $q_{h'}$:

$$q_0 = q_{[0^3]}$$

$$q_1 = q_{[0^2H]} + q_{[0^2L]}$$

$$q_2 = q_{[0H^2]} + q_{[0LH]} + q_{[0L^2]}$$

$$q_3 = q_{[H^3]} + q_{[LH^2]} + q_{[L^2H]} + q_{[L^3]}$$

$$q_{0'} = q_{[0^3]} + q_{[0^2L]} + q_{[0L^2]} + q_{[L^3]}$$

$$q_{1'} = q_{[0^2H]} + q_{[0LH]} + q_{[L^2H]}$$

$$q_{2'} = q_{[0H^2]} + q_{[LH^2]}$$

$$q_{3'} = q_{[H^3]}$$

There are 10 unknown variables within these equations. However, the number of independent equations for q_j and $q_{h'}$ is only 7. Three additional equations are needed. These three equations can be found through the balance equations for the three

variables $q_{[0LH]}$, $q_{[L^2H]}$, and $q_{[LH^2]}$. After these 10 variables are known, the values for $q_{l'}$ can be found through the following equations.

$$q0' = q[03] + q[02H] + q[0H2] + q[H3]$$

$$q1' = q[02L] + q[0LH] + q[LH2]$$

$$q2' = q[0L2] + q[L2H]$$

$$q3' = q[L3]$$

The general procedure for a buffer with size b is introduced as follows: First the values for q_j and $q_{h'}$ are found. The total number of independent equations given by q_j and $q_{h'}$ is 2b + 1. The total number of unknown variables for $q_{[0^X L^Y H^Z]}$ is $\frac{1}{2}$ b (b + 1), where X + Y + Z = b. The other independent equations can be found from the balance equations for the variables $q_{[0^X L^Y H^Z]}$. For easy computation, the value of z can be limited to at least 1; hence, the number of balance equations in this case is $\frac{1}{2}$ b (b-1), which is sufficient to calculate the unknown variables.

7.3 Cell Loss Ratio

7.3.1 High-Priority Cells

The cell loss ratio for high-priority cells (CLRH) is the same as in Model B.

$$CLRH = \frac{\sum_{h'=0}^{b} q_{h'} \sum_{h=0}^{N(H)} e(H)_h \, Max\,((h+h'-r-b), 0)}{\sum_{h=0}^{N(H)} e(H)_h . h}$$

7.3.2 Low-Priority Cells

Cell loss in the low-priority cells can be distinguished into two situations. The first is the incoming low-priority cells cannot find space in the queue. The second is the low-priority cells already in the queue are pushed out of the buffer due to the excessive arrival of the high-priority cells. Therefore, the cell loss ratio of the low-priority cells depends on the buffer occupancy and the number of new arriving high-priority cells.

Two cases can be considered based on whether high-priority cells are lost or not.

First, assume that the high-priority cells are not lost. If the number of cells in the buffer plus the incoming cells exceeds the sum of the transmission rate and buffer size, then low-priority cells will be lost. In other words, if $h' + l' + h + l > r + b$, then low-priority cells are lost where $l > 0$. Since the assumption is that the high-priority cells should not be lost, $h + h' \leq r + b$.

Secondly, assume that the high-priority cells are lost being conditional upon the number of high-priority cells in the buffer plus the number of incoming cells being larger than the sum of the transmission rate and the buffer size. If high-priority cells are lost, then all the $l'+l$ low-priority cells will be lost. In other words, if $h' + h > r + b$, then $l' + l$ low-priority cells are lost. Based on these two cases, the CLRL can be derived as

$$\frac{\sum_{h'=0}^{b} \sum_{l'=0}^{N(L)} q_{h'}\, q_{l'} \left[\sum_{l=0}^{r+b-h'} e(L)_l\, e(H)_h\, \mathrm{Max}(h'+l'+h+l-r-b,\, 0) + \sum_{h=r+b-h'+1}^{N(H)} e(L)_l\, e(H)_h\, (l'+l) \right]}{\sum_{l=0}^{N(L)} e(L)_l . l}$$

There is an easy way of calculating the CLRL. The basic concept is that the total number of lost cells is the same as the sum of the number of lost high-priority cells and the number of lost low-priority cells. Hence, the relationship among CLR, CLRH, and CLRL is expressed as follows.

$$R * CLR = RH * CLRH + RL * CLRL$$

Remember that R is the average number of cells entering the system in a unit time.

$$R = \sum_{i=0}^{N} i\, e_i$$

$$RH = \sum_{i=0}^{N(H)} h\, e(H)_h$$

$$RL = \sum_{i=0}^{N(L)} l\, e(L)_l$$

7.4 Cell Transfer Delay and Cell Delay Jitter

7.4.1 High-Priority Cells

The calculation of CTDH and CDJH is the same as in Model A or Model B.

7.4.2 Low-Priority Cells

The calculation of CTDL is analogous to that in Model C. In Model C, there are two physically separate queues; in Model D, there are two logically separate queues. Define a new variable s = j + h, where j is the number of cells in the buffers and h is the number of incoming high-priority cells. Since only non-lost low-priority cells are considered, at least one buffer should exist for the low-priority cell-of-interest; hence, the values of s are limited between 0 and b + r - 1. Define an event $\Psi = \{ 0 \le s \le b+r-1\}$ and a new probability $Pr[F=f \mid (s \mid \Psi)]$, which is the probability that the cell-of-interest will be in position f + 1 of all the possible r + b positions given $(s \mid \Psi)$. Since only the non-lost low-priority cells are of interest, the variable, f + 1, of the cell-of-interest can only go up to r+b.

Hence, when f is greater than r + b - 1, $Pr[F=f \mid (s \mid \Psi)]$ is equal to 0.

By definition, when f is less than s, $Pr[F=f \mid (s \mid \Psi)]$ is equal to 0.

Following the same procedure as in Model B, $Pr[F=f \mid (s \mid \Psi)]$ can be derived as

$$Pr[F=f \mid (s \mid \Psi)] = \frac{\sum_{l=f-S+1}^{N(L)} e(L)_l}{\sum_{l=0}^{b+r-s} l\, e(L)_l + \sum_{l=b+r-s+1}^{N(L)} (b+r-s)\, e(L)_l}$$

The $Pr[s \mid \Psi]$ can be derived as follows.

$$Pr[s \mid \Psi] = \frac{\sum_{j+h=s} q_j e(H)_h}{\sum_{s=0}^{b+r-1} \sum_{j+h=s} q_j e(H)_h}$$

Now, the probability that the cell-of-interest will be in position f + 1 of all the b + r possible positions is

$$\Pr[F=f] = \sum_{s=0}^{b+r-1} \Pr[s \mid \Psi] \Pr[F=f \mid (s \mid \Psi)]$$

Define another probability Pr[K=k | F=f] which is the probability that the cell-of-interest will have to wait for k cell time before being transmitted, given that this cell is in position f+1. The values of k depend on the values of f and r. Two cases can be distinguished.

a. If $f + 1 \leq r$, this cell-of-interest will be transmitted during this unit time. Hence k is equal to f.

b. If $f + 1 > r$, this cell-of-interest will not be transmitted during this unit time. However, the position of this cell-of-interest will be advanced to $f + 1 - r$. The cell waiting time at this instant is r.

Now consider the next unit time. At this unit time, there might be new incoming high-priority cells. Hence, this low-priority cell-of-interest will not stay at $f + 1 - r$. The new position for this cell-of-interest is $f + 1 - r + h$, where h is the number of new incoming high-priority cells. Again two cases can be distinguished.

b1. If $f + 1 + h \leq 2r$, this cell-of-interest will be transmitted at this unit time. The total cell waiting time is $r + f - r + h$.

b2. If $f + 1 + h > 2r$, this cell-of-interest will not be transmitted at this unit time. This cell will be in the new position of $f + 1 + h - 2r$. The cell waiting time at this instant is 2r.

By recursively performing this procedure, the cell waiting time can be calculated. Two events will be ignored in this calculation. The first one is that the cell waiting time might be infinite. The second one is the low-priority cell might be lost due to the excessive arrival of high-priority cells.

After Pr[K=k | F=f] is calculated, the probability that a low-priority cell will wait for k cell time, Pr[K=k], can be easily obtained.

$$\Pr[K=k] = \sum_{f=0}^{b+r-1} \Pr[F=f] \Pr[K=k \mid F=f]$$

After Pr[K=k] is obtained, the calculation of CTD and CDJ is a direct matter.

8 Space Segment Queueing Model

The space segment is assumed to be a fast packet switch with output buffering (see Figure C-8). The interconnection network used in the fast packet switch is assumed to be nonblocking. The switch is operated in slotted mode. The switch will route the cells

8 Space Segment Queueing Model

The space segment is assumed to be a fast packet switch with output buffering (see Figure C-8). The interconnection network used in the fast packet switch is assumed to be nonblocking. The switch is operated in slotted mode. The switch will route the cells based on the VPI/VCI field in the header to the appropriate down-link beam. The ATM performance through the satellite node is measured by observing one output port. In this context, the queueing model developed for the statistical multiplexer can be extended to the space segment. The notation used in this subsection is introduced as follows.

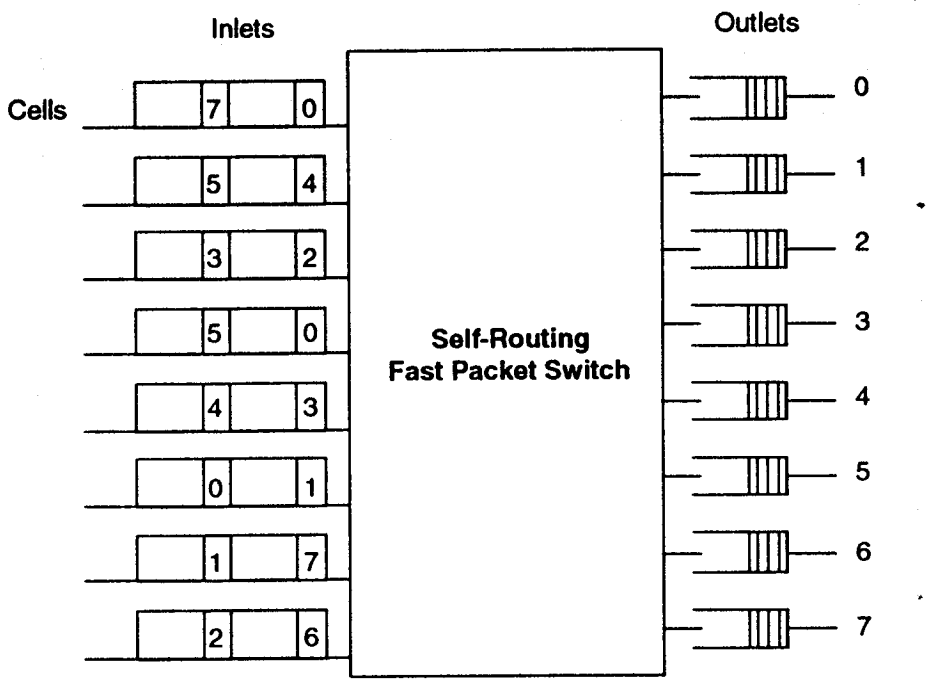

Figure C-8. An 8 x 8 Fast Packet Switch with Output Buffering

N_{beam} The number of spot beams in the satellite system.

s_i The probability that an earth station will transmit i cells destined to a particular down link beam during a unit time.

The probability that the earth station will transmit j cells during one unit time has been derived, which is t_j. Denote the probability that k of these j cells will be destined to a particular down-link beam as P(j,k). Assume all cells have equal probability to be destined to each down-link beam. Then P(j,k) can be derived as

$$P(j,k) = \binom{j}{k} \left(\frac{1}{N_{beam}}\right)^k \left(1 - \frac{1}{N_{beam}}\right)^{j-k}$$

where $\binom{j}{k} = \frac{j!}{k!\,(j-k)!}$, and s_i can be derived as

$$s_i = \sum_{j=i}^{r} t_j P(j,i)$$

The s_i is equivalent to Q_j at the earth station except that the maximum value of j is 1 while the maximum value for i is r. The derivation of performance measurements is the same as in Model A, which will not be repeated here. The main objective is to evaluate the sensitivity of the switch output buffer with respect to the output transmission rate.

9 Computation Results

Based on these queueing models, the ATM performance measurements will be computed and compared using different input traffic characteristics, different buffer sizes, and different output transmission rates. A conclusion will be made based on these comparisons.

9.1 CLR, CTD and CDJ

The performance parameters (CLR, CTD, and CDJ) will be computed as a function of different traffic characteristics, number of input lines, size of the buffer, output line rate, and different priority and buffer allocation schemes.

9.1.1 Burstiness of Input Traffic

The burstiness is defined as the peak bit rate divided by the mean bit rate. The effect of burstiness on the ATM performance is shown in Figures C-9, C-10, and C-11. The other parameter which varies in these figures is the output transmission rate. In these computations the number of lines at each earth station is 10, and all the lines share the same characteristics. It can be seen from Figure C-9 that to maintain CLR of 10^{-9}, the output transmission rate must be increased from 1 to 2 to achieve the QOS when burstiness is higher than 7.5. In other words, higher burstiness demands more capacity to achieve the same QOS. In Figures C-10 and C-11, when the output transmission rate is 1, the CTD and CDJ increases exponentially when the burstiness goes beyond 8. If this earth station also interfaces with the circuit emulation services (AAL Class 1 Service), then the exponential increase of CTD and CDJ might become a problem for timing recovery at the endpoint. This situation is depicted in Figures C-12 and C-13. In these computations, the number of lines are also 10; however, there are 2 trunk lines with burstiness equal to 1, and there are 8 bursty lines. The degradation of the circuit emulation service can be remedied using the priority concept, which will be discussed latter. To effectively support the B-ISDN traffic and to efficiently design the earth station, the information regarding the traffic characteristics is critical. In general, the higher the burstiness, the worst the ATM performance through the statistical multiplexer.

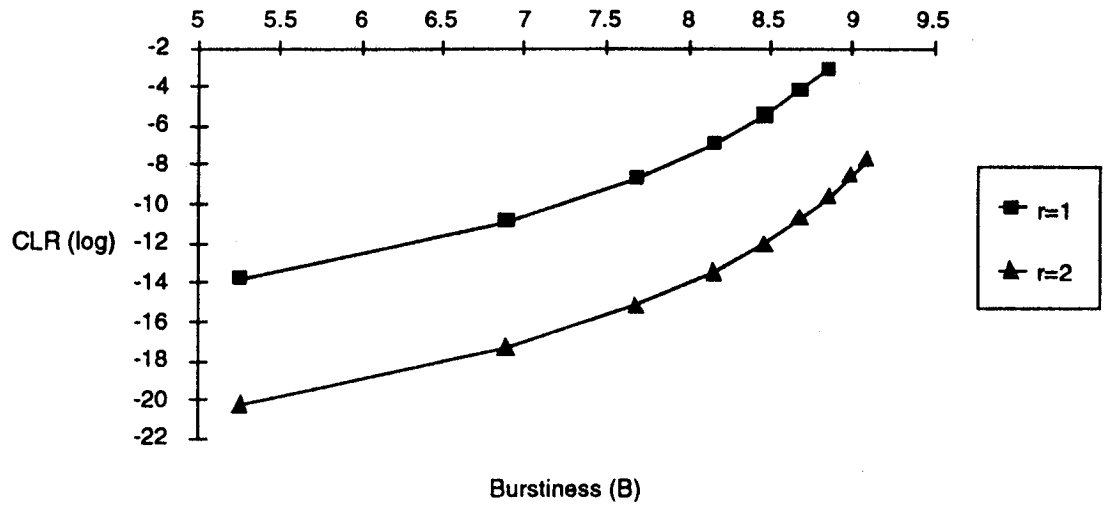

Figure C-9. Cell Loss Ratio vs. Burstiness for Different Output Transmission Rates

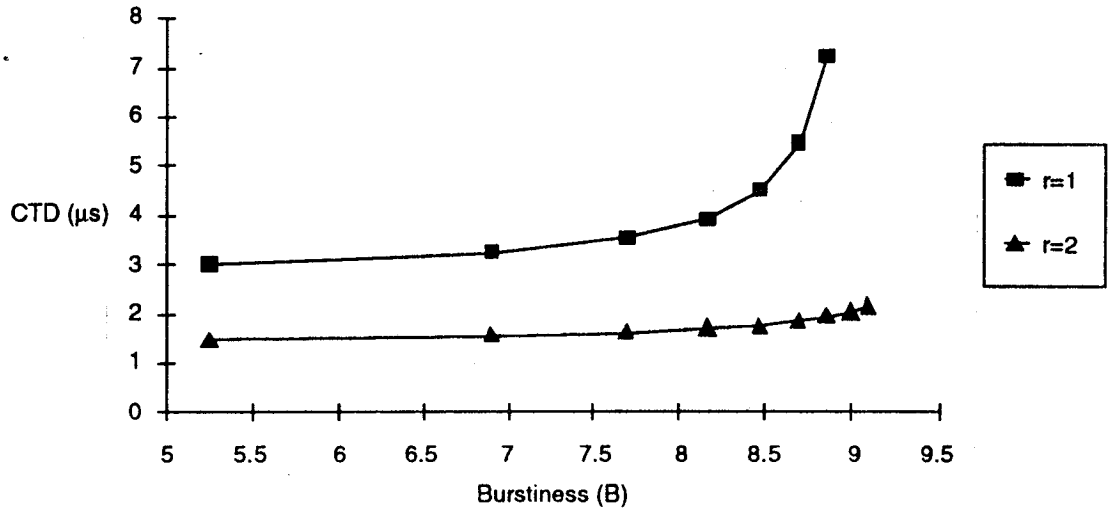

Figure C-10. Cell Transfer Delay vs. Burstiness for Different Output Transmission Rates

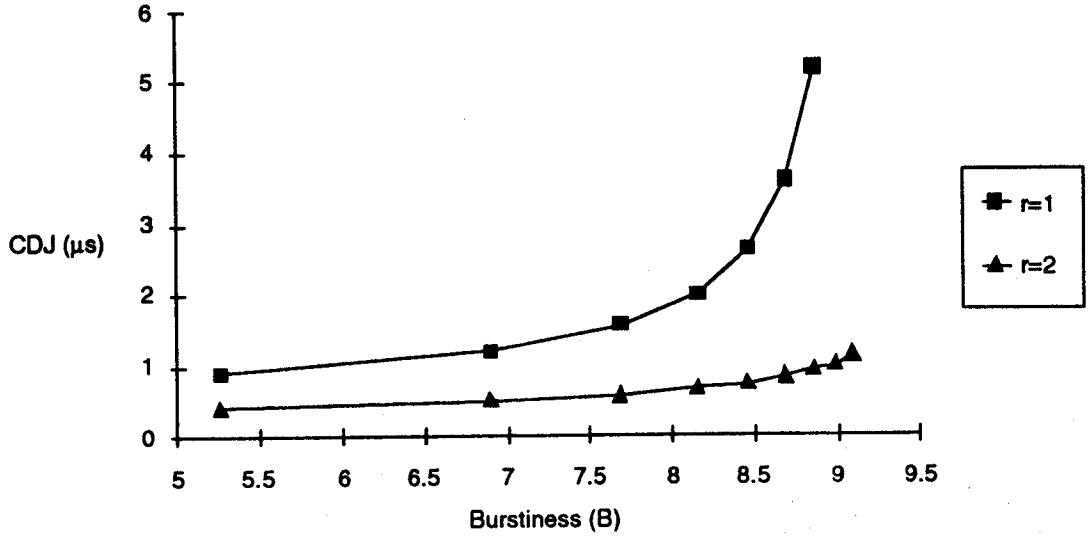

Figure C-11. Cell Delay Jitter vs. Burstiness for Different Output Transmission Rates

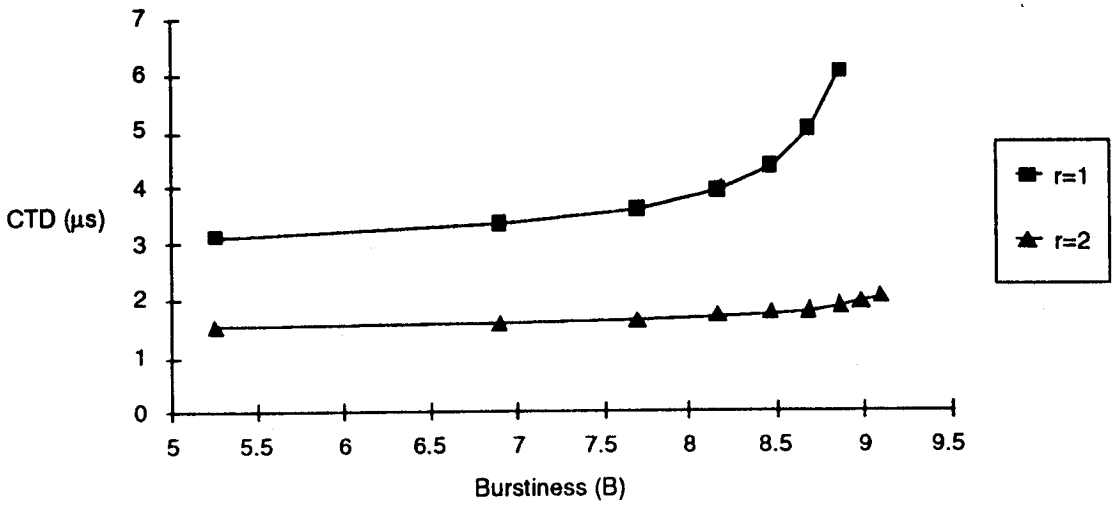

Figure C-12. Cell Transfer Delay vs. Burstiness for Mixed Traffic Scenario

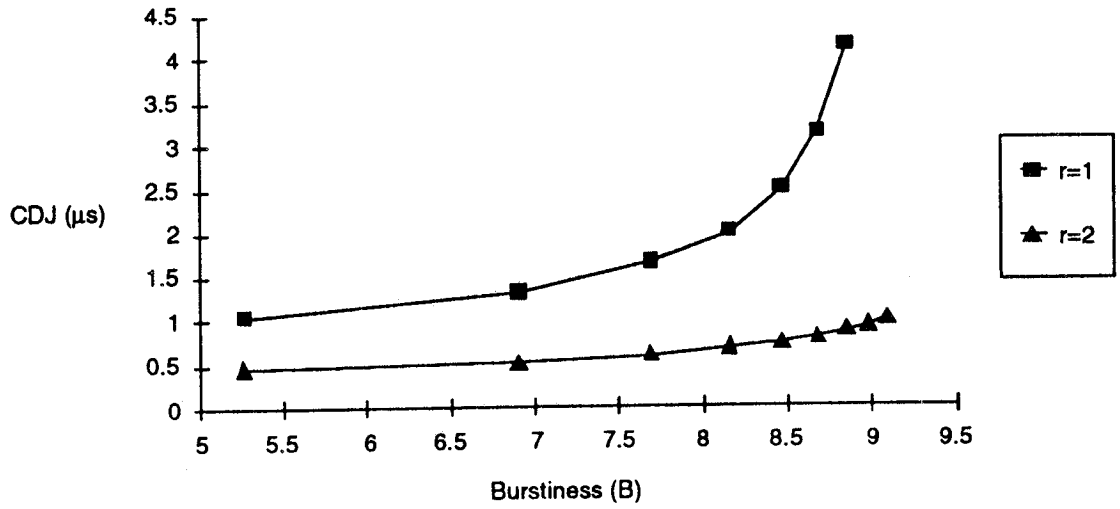

Figure C-13. Cell Delay Jitter vs. Burstiness for Mixed Traffic Scenario

9.1.2 Buffer Size

The impact of the buffer size at the earth station on the ATM performance is shown in Figures C-14, C-15, and C-16. When the average arrival rate (R) is small, the improvement of CLR due to the increase of the buffer size is large compared to that when the average arrival rate is large. The reason is that when the arrival rate is high, the output link is almost saturated; hence, adding buffers will slightly alleviate the output saturation problem. However, when the arrival rate is small, the buffer is not fully utilized; therefore, increasing the buffer size makes the buffer more empty; hence, CLR is improved significantly. As for CTD and CDJ, when the buffer size is small, increasing the buffer size will increase the delay; however, when the buffer size is large, increasing the buffer size has almost no effect on CTD and CDJ. The reason is even though the buffer size is increased, the buffer occupancy is still concentrating on the head portion instead of the tail portion; hence, delay will not be affected.

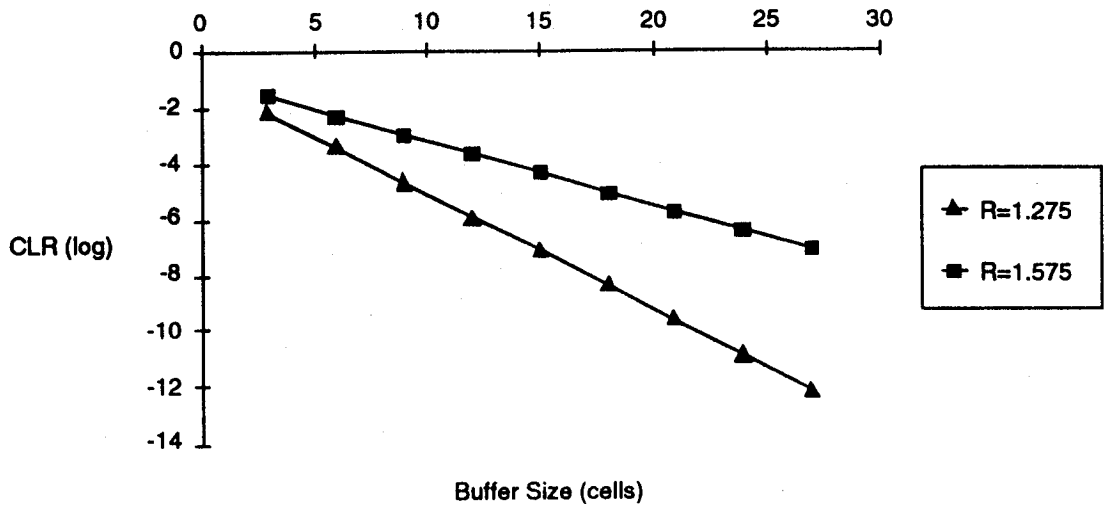

Figure C-14. Cell Loss Ratio vs. Buffer Size for Different Average Arrival Rates

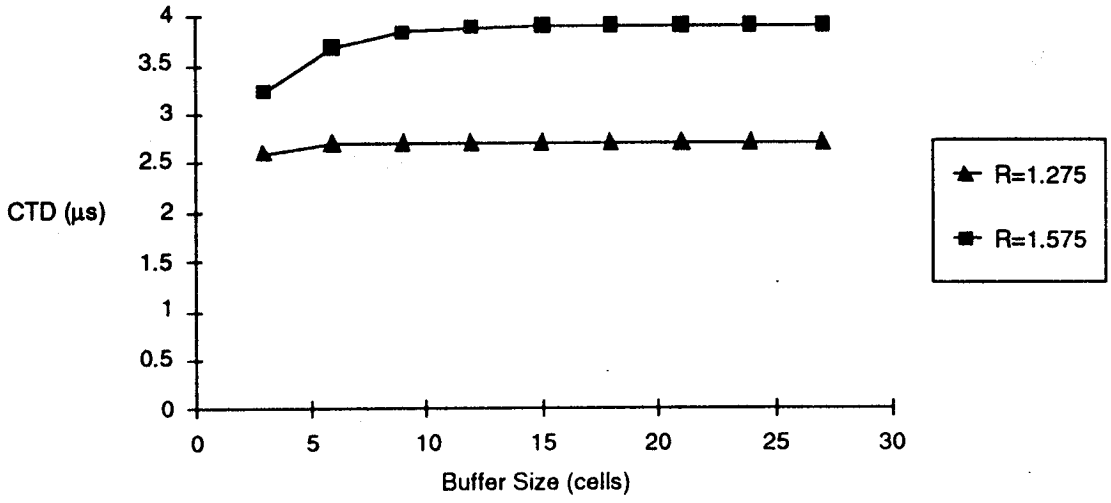

Figure C-15. Cell Transfer Delay vs. Buffer Size for Different Average Arrival Rates

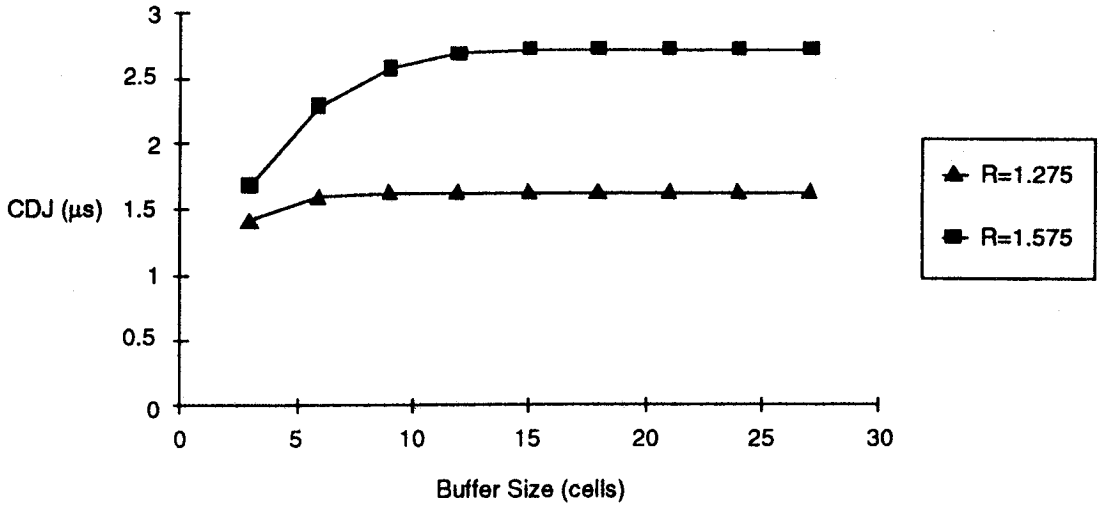

Figure C-16. Cell Delay Jitter vs. Buffer Size for Different Average Arrival Rates

In Figure C-17, CLR versus the buffer size for different output transmission rates is presented. To reach a required CLR at the earth station, for example 10^{-9}, using either a high output transmission rate and a small amount of buffer or a low transmission rate and a large amount of buffers. Hence by adding buffers at the earth station, the output transmission rate can be reduced, i.e., the transmission capacity requirement is reduced. The penalty of this capacity reduction approach using more buffers is that the delay will be increased.

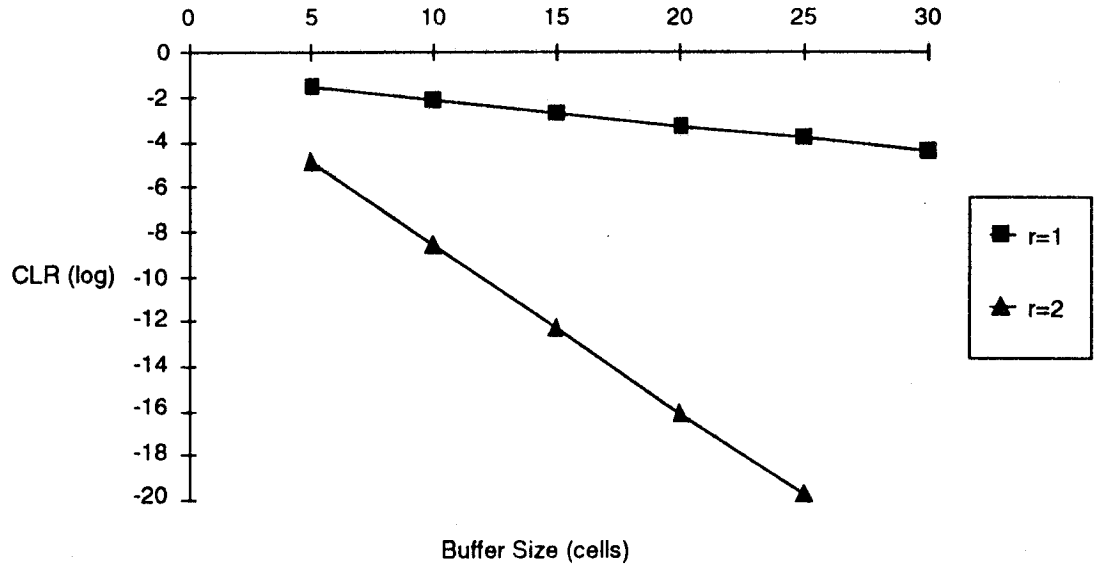

Figure C-17. Cell Loss Ratio vs. Buffer Size for Different Output Transmission Rates

9.1.3 Output Line Rate

The impact of the output line rate to the ATM performance is depicted in Figures C-18, C-19, and C-20. As expected, when the line rate is large, the improvement in performance is noticeable. However, since the bandwidth is a limited resource, the output utilization efficiency has also to be considered.

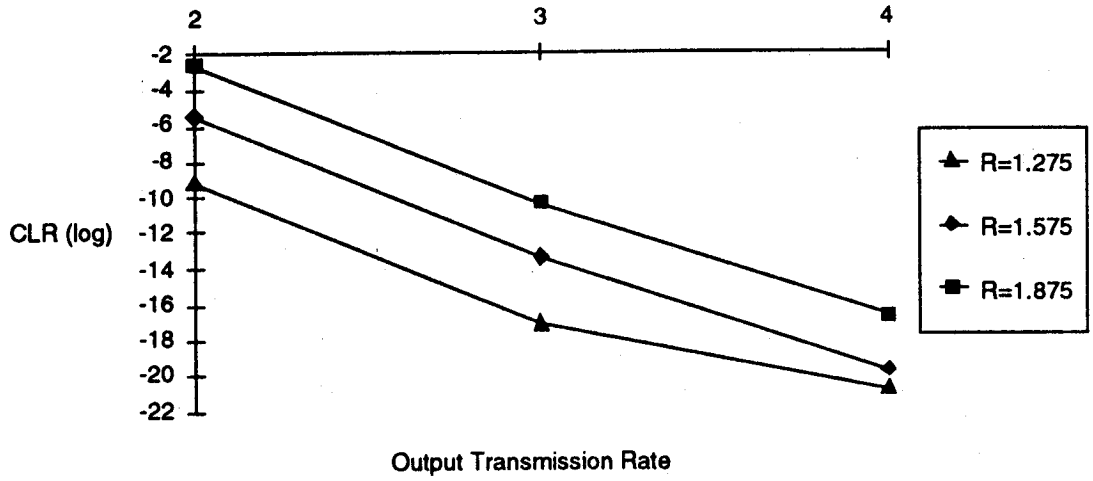

Figure C-18. Cell Loss Ratio vs. Output Transmission Rate for Different Average Arrival Rates

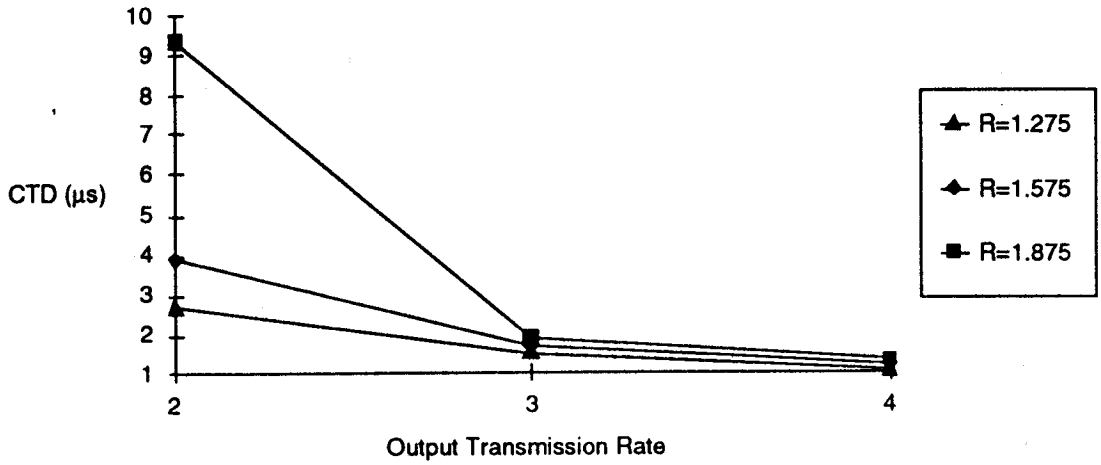

Figure C-19. Cell Transfer Delay vs. Output Transmission Rate for Different Average Arrival Rates

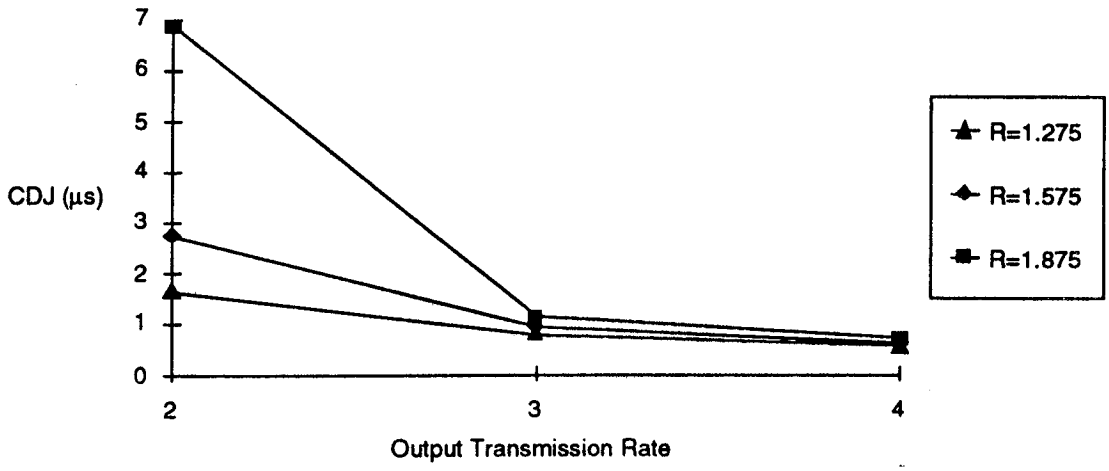

Figure C-20. Cell Delay Jitter vs. Output Transmission Rate for Different Average Arrival Rates

9.1.4 Number of Input Lines

Figures C-21, C-22, and C-23 show the effect of the number of lines on the ATM performance. For each curve, the load (ratio of average arrival rate and output transmission rate) of the earth station is kept constant. The only variable is the number of input lines (N). Although the load at the earth stations is the same, the larger the number of input lines the better the performance. Therefore the gain of statistical multiplexing is larger when the number of input lines is large. This is easy to understand since if there is only one input line, the statistical multiplexer becomes a deterministic multiplexer. The greater the number of the input lines, the smaller the possibility that bursts from different input lines will come together. Hence, the sum of the average bit rates of incoming lines indicates a good approximation for the output transmission rate if the number of input lines is large.

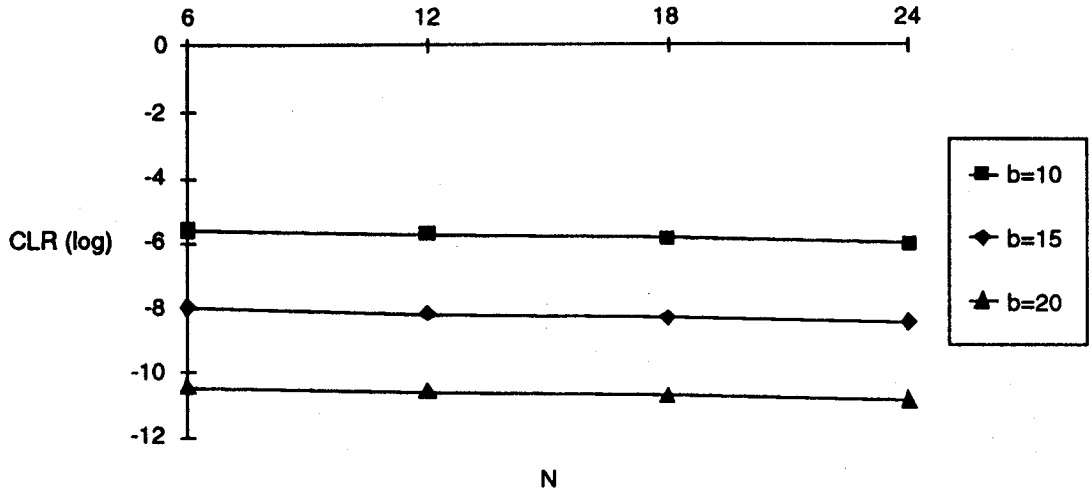

Figure C-21. Cell Loss Ratio vs. Number of Input Lines for Different Buffer Sizes

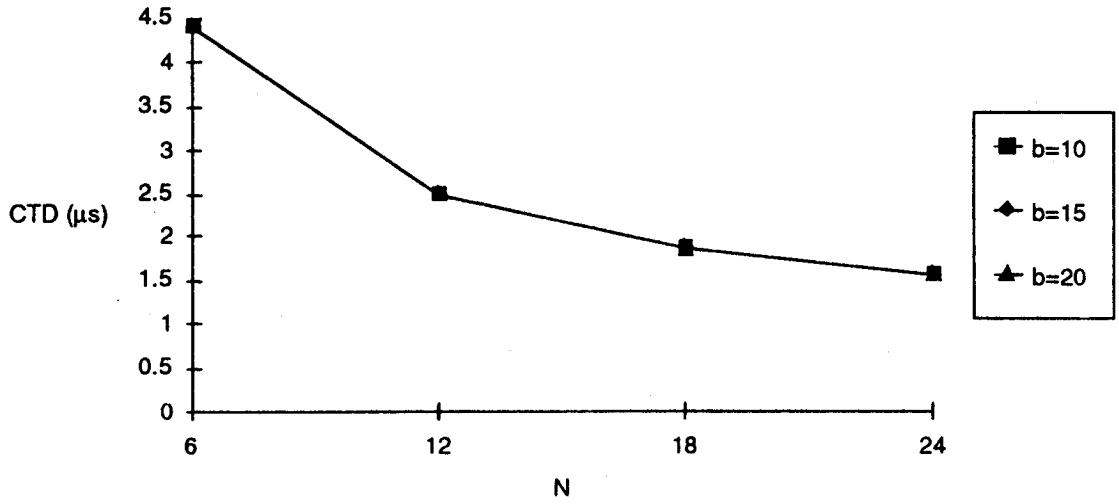

Figure C-22. Cell Transfer Delay vs. Number of Input Line for Different Buffer Sizes

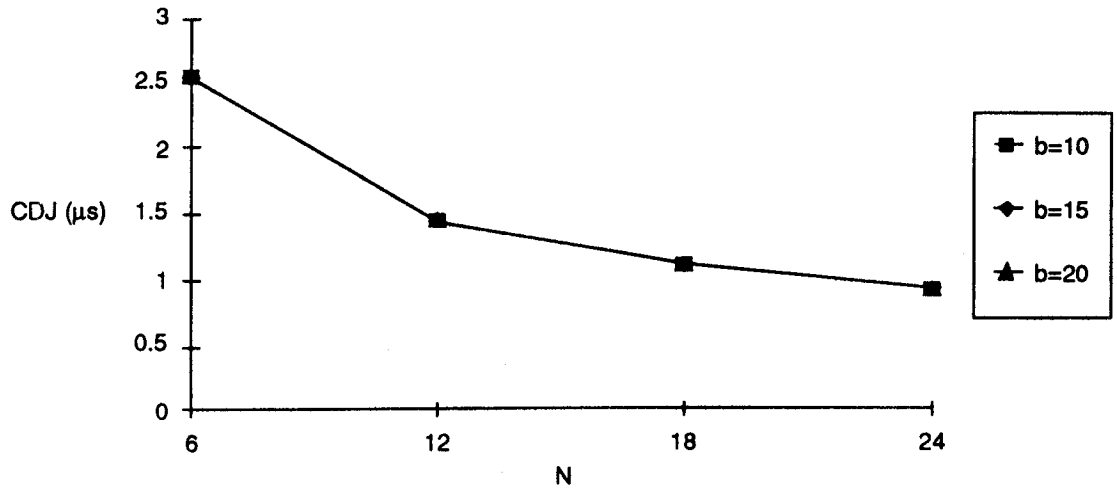

Figure C-23. Cell Delay Jitter vs. Number of Input Lines for Different Buffer Sizes

9.1.5 Ratio of Average Burst Size and Buffer Size vs Number of Input Lines

Figures C-24, C-25, and C-26 depict the impact of the ratio of average burst size and buffer size to the ATM performance. These results show that when the ratio of average burst size and buffer size is less than 1, the performance degradation due to adding new lines is not very significant. However, when the ratio of average burst size and buffer size is larger than 1, the performance degradation due to adding new lines is noticeable. This suggests that the average burst size should be included in the traffic characteristics, which are parameters used to set up a connection through the network.

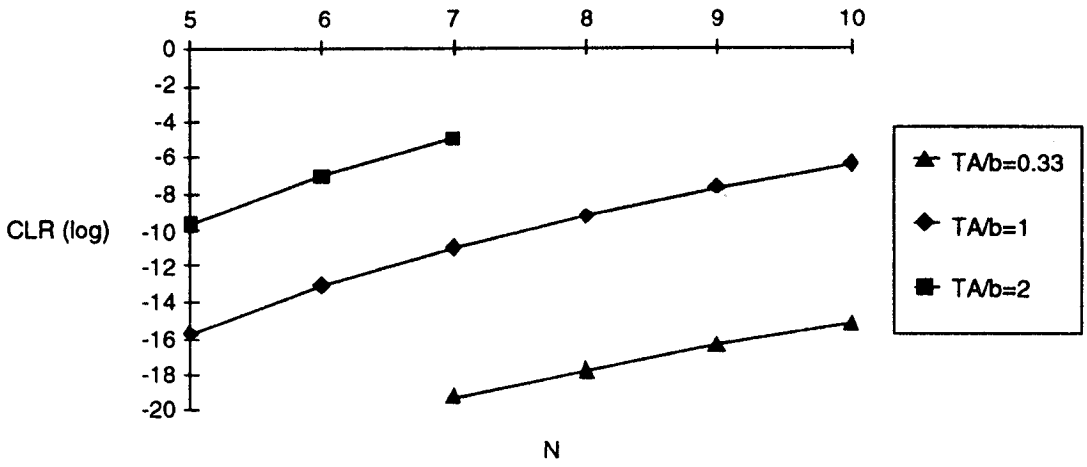

Figure C-24. Cell Loss Ratio vs. Number of Input Lines for Different Ratios of Average Burst Size and Buffer Size

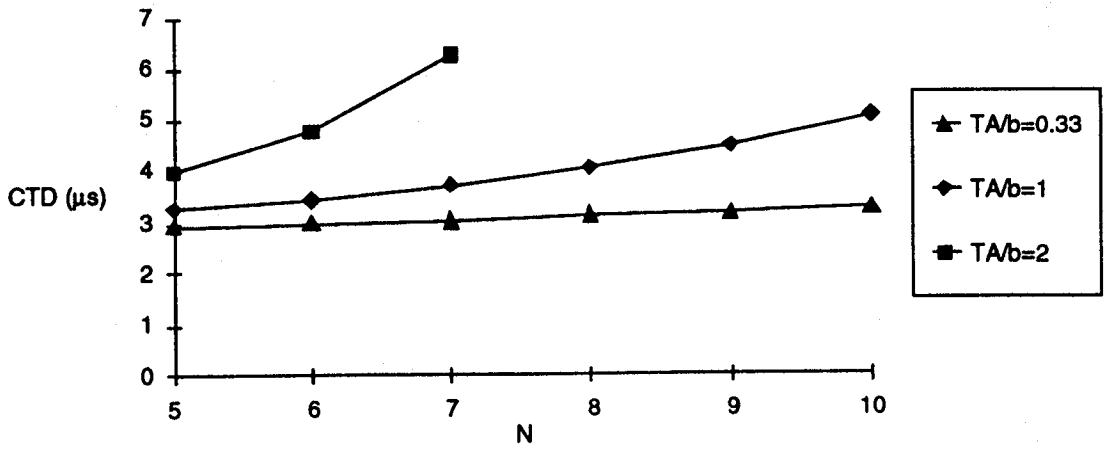

Figure C-25. Cell Transfer Delay vs. Number of Input Lines for Different Ratios of Average Burst Size and Buffer Size

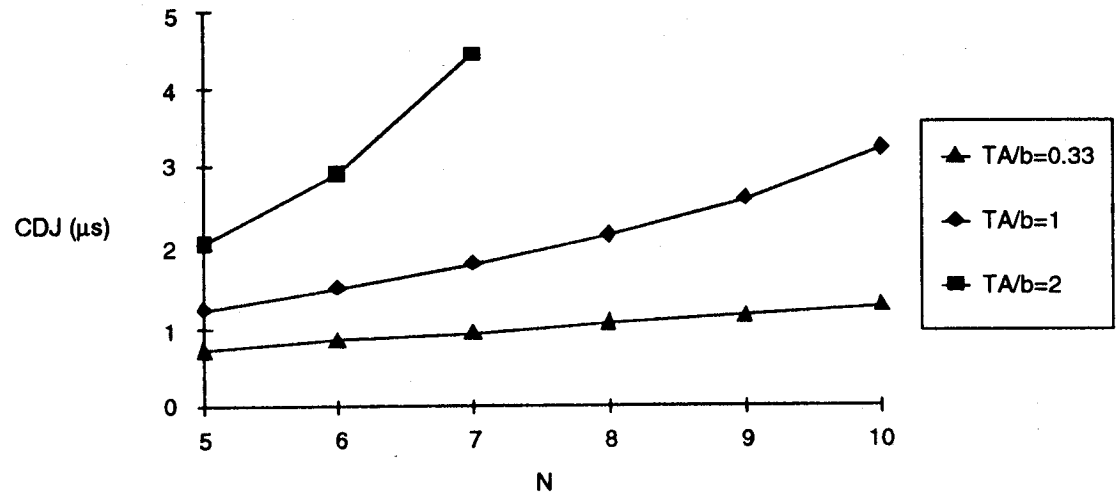

Figure C-26. Cell Delay Jitter vs. Number of Input Lines for Different Ratios of Average Burst Size and Buffer Size

9.1.6 Priority and Buffer Allocation

Figures C-27 to C-33 illustrate the performance comparisons among different models. The number of lines interfaced with the earth station is 10. There are two trunk lines carrying high-priority cells and 8 bursty lines carrying low-priority cells. For CLR, the high-priority cells perform best for Model D followed by Model C. This is easy to understand since all the buffers in Model D can be used by the high-priority cells. Model C also performs well when the number of buffers assigned to high-priority cells is large. For CLR of low-priority cells, Models D and B perform best while Model C performs worst. The reason is the inefficient use of the buffer. In Figure C-28, average CLRs for Models A and C are compared. The CLR is worse when the buffer is divided into two areas. The reason is

because the unutilized buffer in one area cannot be used for another priority of cells. The conclusion so far is that complete partitioning should not be used if CLR is the main requirement.

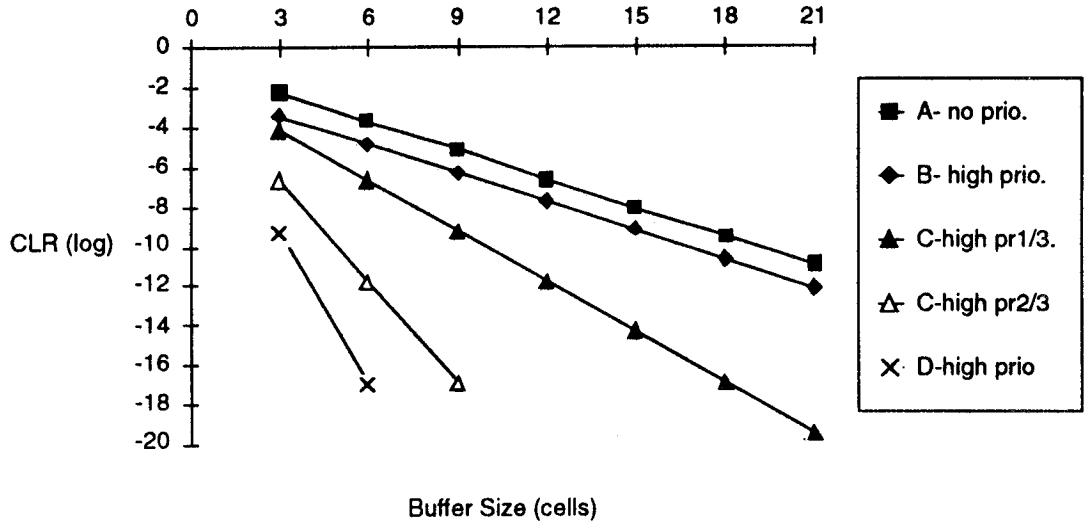

Figure C-27. Cell Loss Ratio of High Priority Cells vs. Buffer Size for Different Models

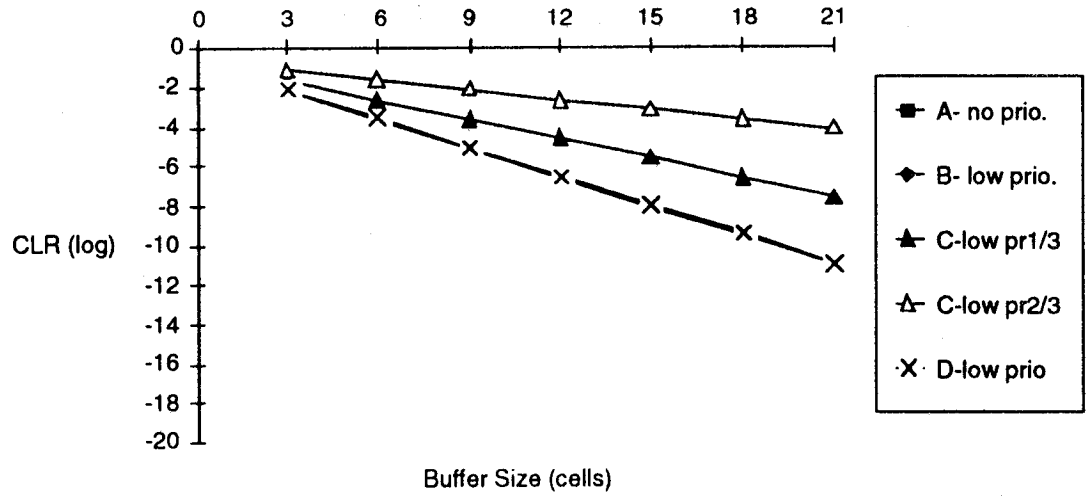

Figure C-28. Cell Loss Ratio of Low Priority Cells vs. Buffer Size for Different Models

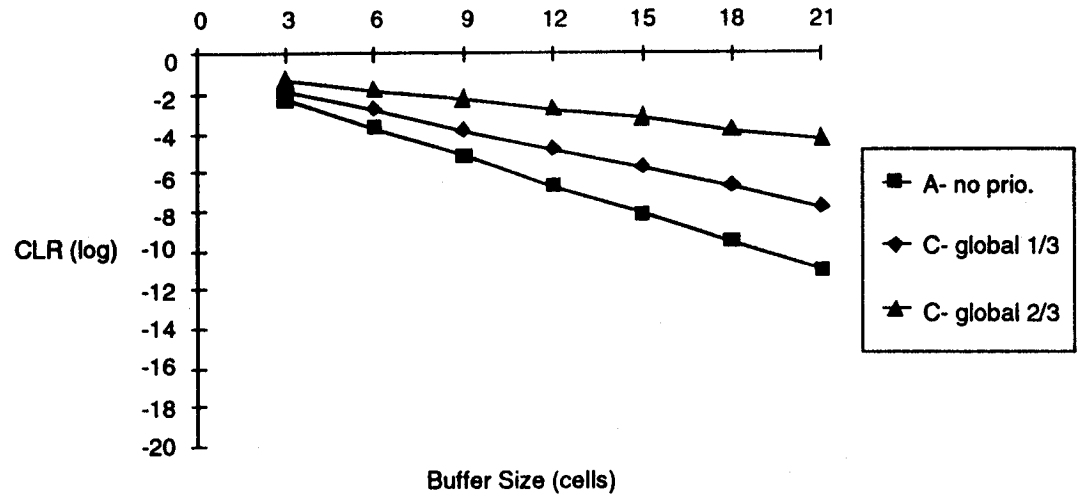

Figure C-29. Average Cell Loss Ratio vs. Buffer Size for Models A and C

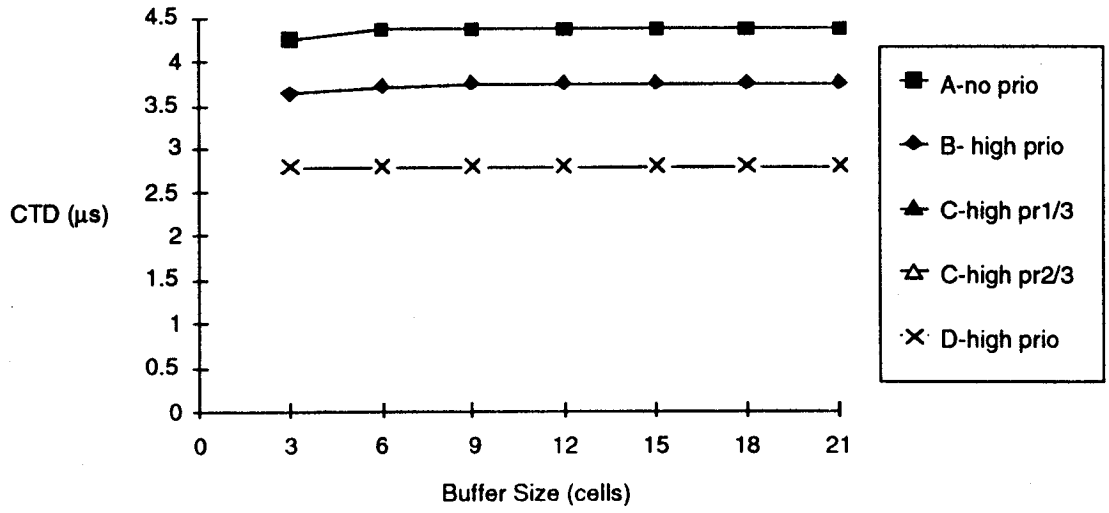

Figure C-30. Cell Transfer Delay of High Priority Cells vs. Buffer Size for Different Models

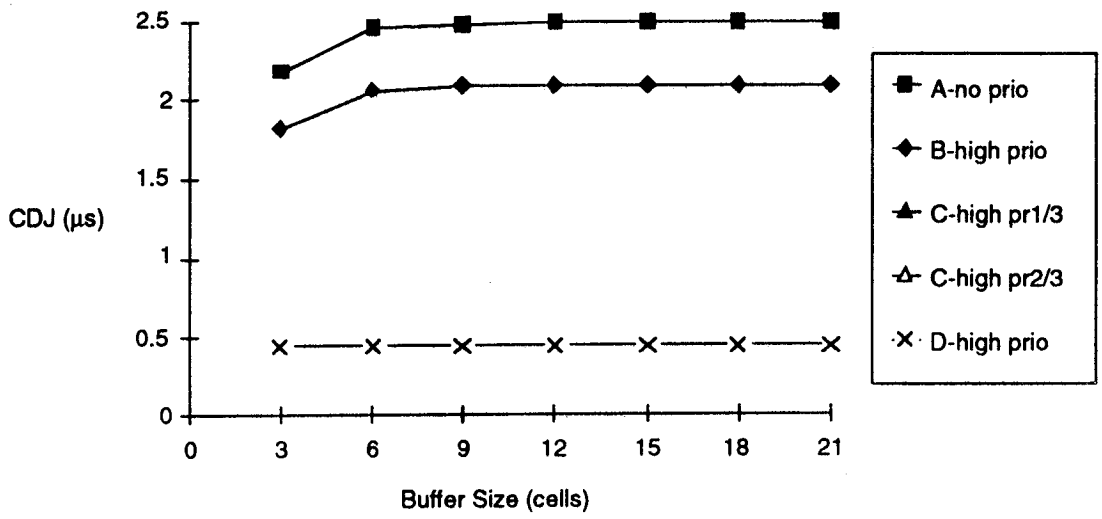

Figure C-31. Cell Delay Jitter of High Priority Cells vs. Buffer Size for Different Models

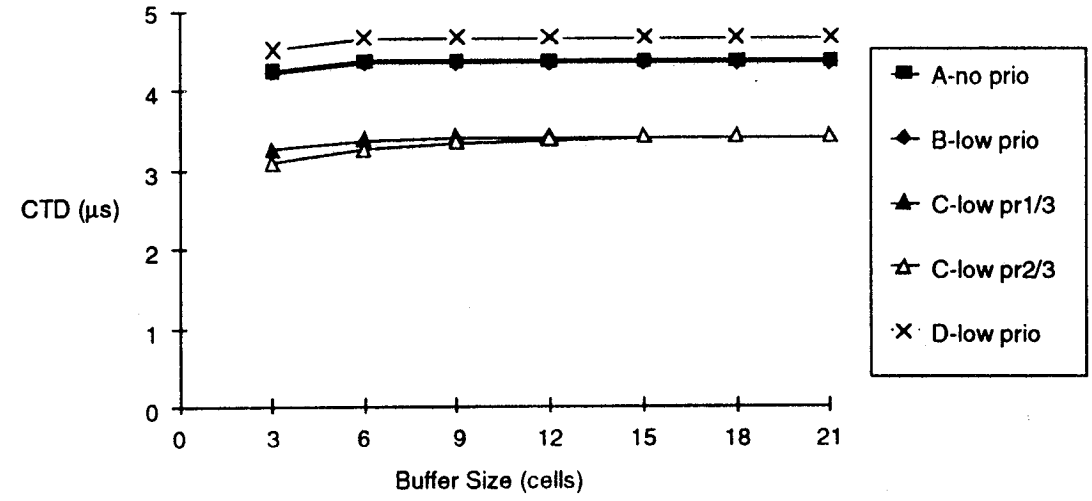

Figure C-32. Cell Delay Jitter vs. Number of Input Lines for Different Ratios of Average Burst Size and Buffer Size

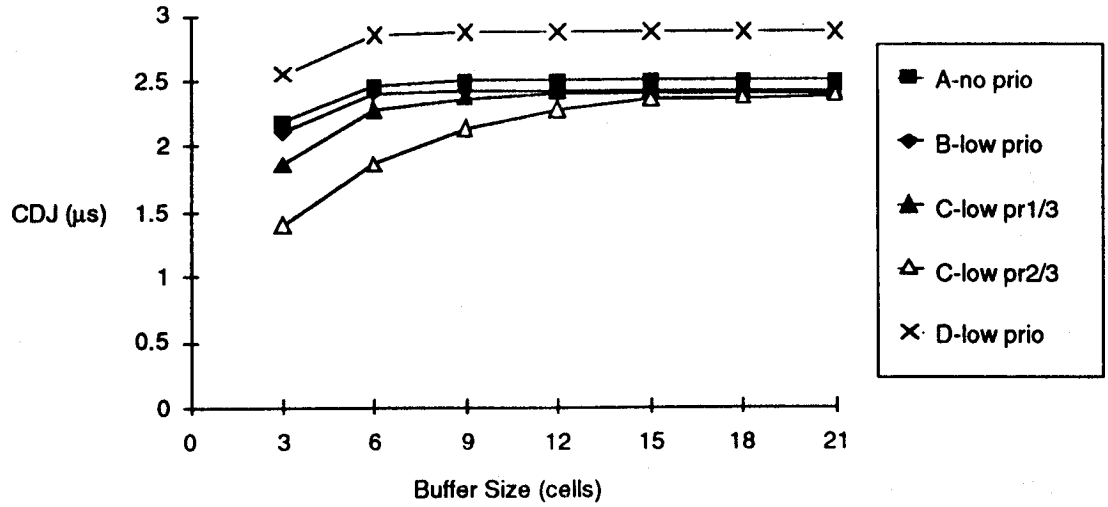

Figure C-33. Cell Transfer Delay of Low Priority Cells vs. Buffer Size of Different Models

Remember that CTD of cells only counts non-lost cells. From Figures C-29 and C-30, the CTD of high-priority cells is largely dependent on the queueing discipline instead of the buffer sharing scheme. The result is that Models D and C perform the best for CTD of high-priority cells. The CTD for low-priority cells is not only dependent on the queueing discipline but also on the buffer allocation scheme. It turns out that Model C performs best and Model D performs worst. The reason is that since the buffer size in Model C is smaller, the cell queueing delay is also smaller. Hence the concern for Model D is that the cell delay for the low-priority cells is high.

Figures C-34 to C-39 present the ATM performance of the earth station when interfaced with inhomogeneous traffic lines. As before, there are always two trunk lines carrying high-priority cells, and there are eight bursty lines carrying low-priority cells. At each computation the bursty line uses different burstiness. As expected, Model D performs best in CLR of high-priority cells, and CLR of high-priority cells is irrelevant to the burstiness of the bursty lines. It is interesting to see that Model C performs worse than Models A and B when the burstiness of the bursty lines is low. However, when the burstiness is high, the CLR of high-priority cells for Models A and B starts degrading and eventually Model C outperforms Models A and B. The reason is when the burstiness of the bursty lines is low, the buffers for Models A and B will be used by the high-priority cells most of the time. Hence, the CLR of high-priority cells still performs well for Model A and B. However, when the burstiness of the bursty lines grows higher, the percentage of high-priority cells among the total cells gets fewer; hence, CLR of high-priority cells degrades rapidly. This shows that the performance of the trunk lines is interfered by the burstiness of other adjacent bursty lines. By using the priority concept, this interference is avoided. For CLR of low-priority cells, Model C performs worst; however, this distinction decreases when the burstiness grows. For CTD and CDJ of the high-priority cells, Models C and D perform best (see Figures C-35 and C-36). For

CTD and CDJ of the low-priority cells, Model C performs best; while Model D performs worst (see Figures C-37 and C-38).

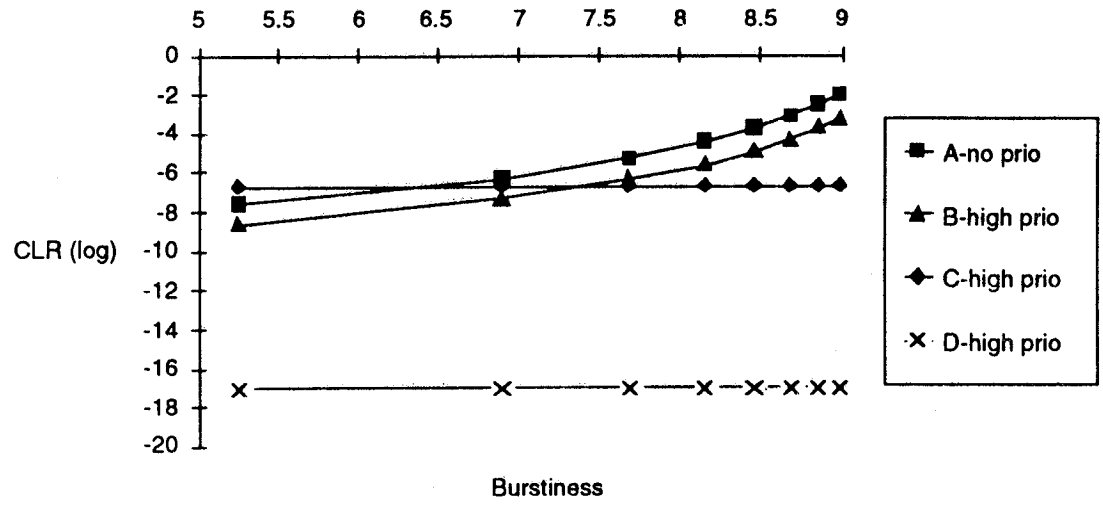

Figure C-34. Cell Loss Ratio of High Priority Cells vs. Burstiness for Different Models

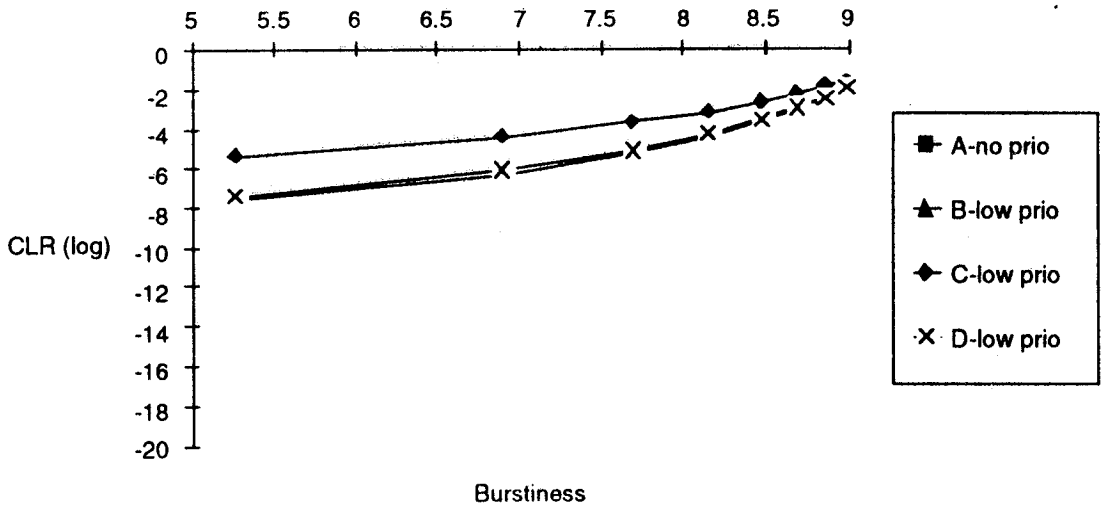

Figure C-35. Cell Loss Ratio of Low Priority Cells vs. Burstiness for Different Models

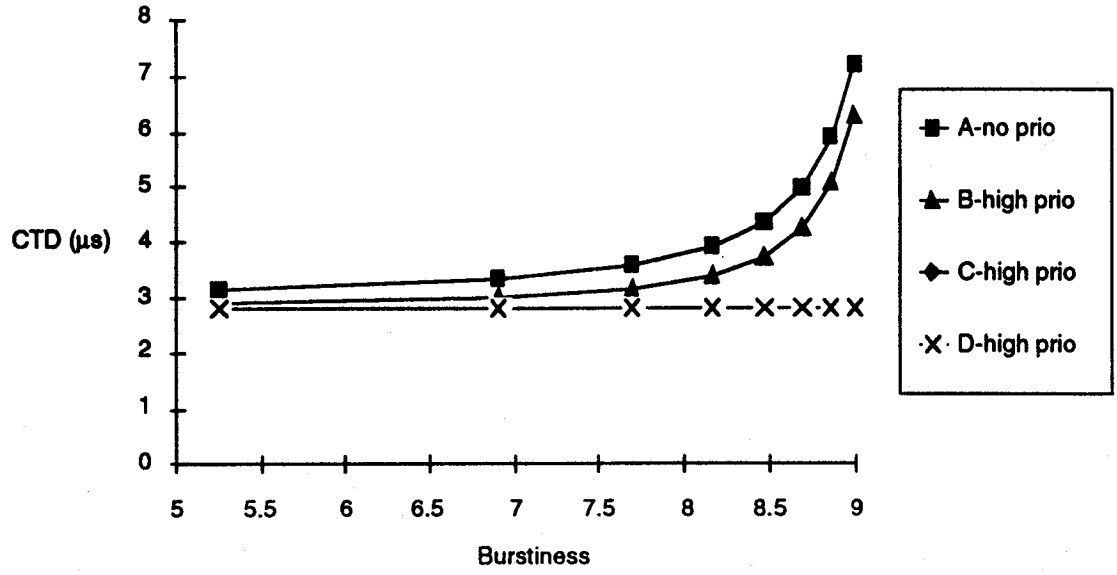

Figure C-36. Cell Transfer Delay of High Priority Cells vs. Burstiness for Different Models

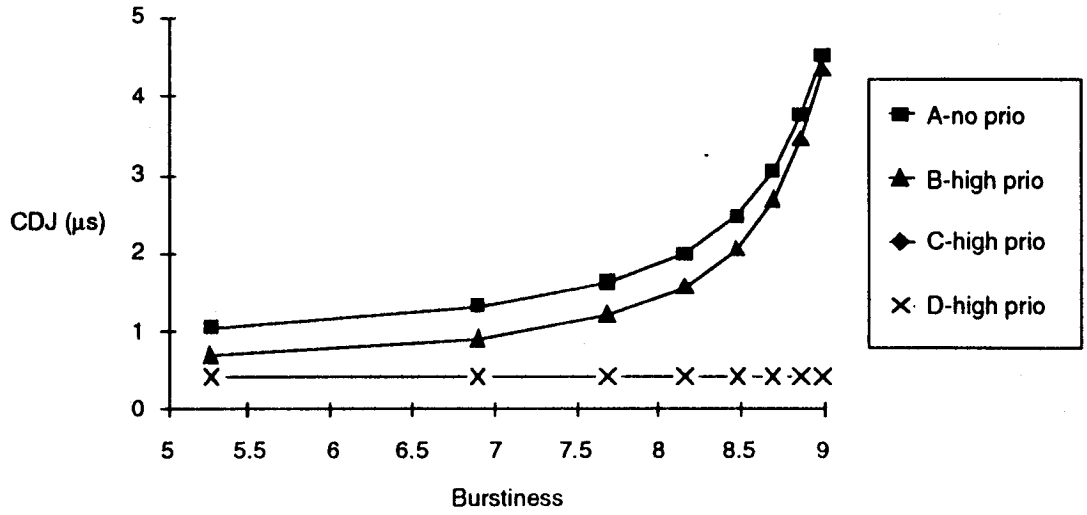

Figure C-37. Cell Delay Jitter of High Priority Cells vs. Burstiness for Different Models

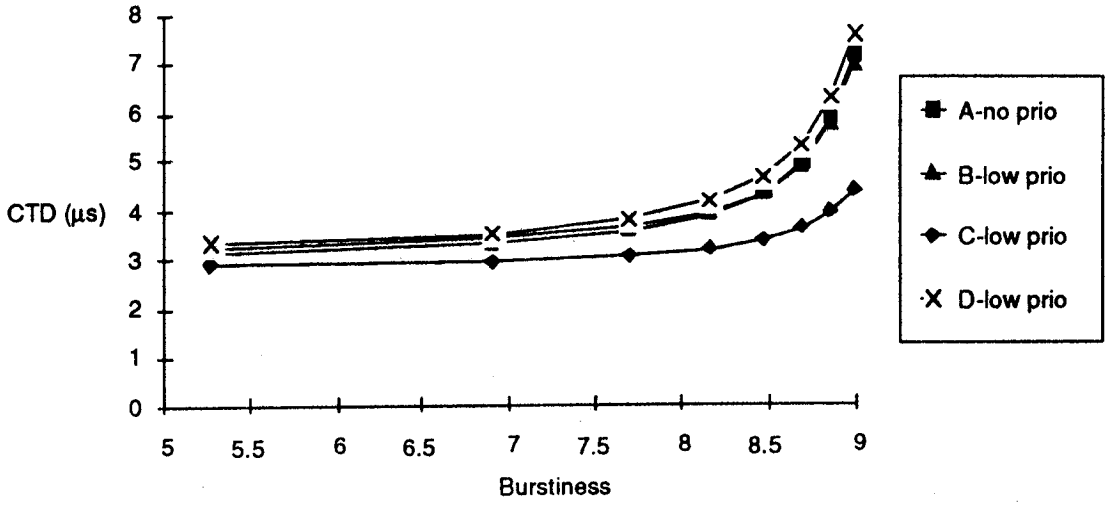

Figure C-38. Cell Transfer Delay of Low Priority Cells vs. Burstiness for Different Models

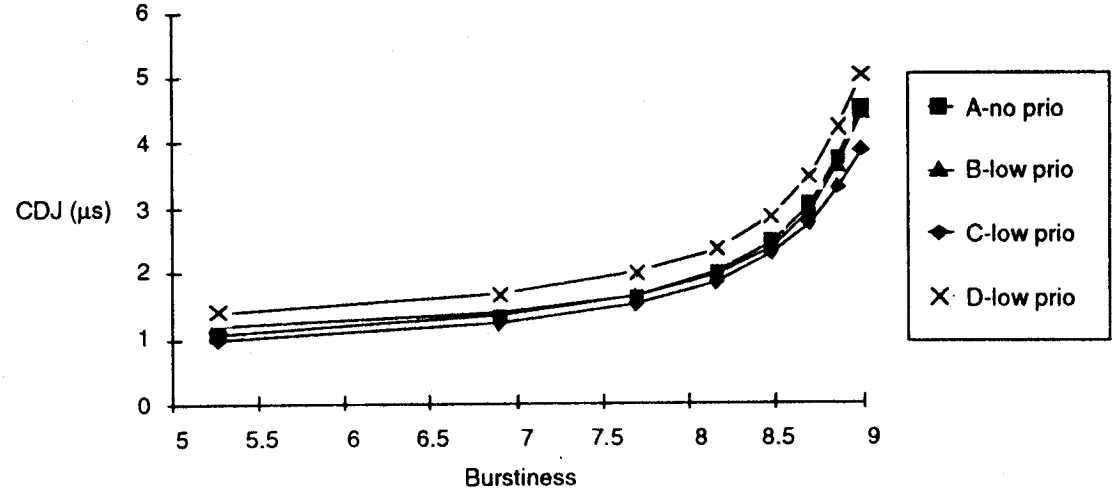

Figure C-39. Cell Delay Jitter of Low Priority Cells vs. Burstiness for Different Models

9.2 Further Research

In this subsection, the applications for the queueing models will be demonstrated. The first one is to extend the queueing model to the space segment so that the ATM cell performance within the satellite B-ISDN system can be measured. The second one is to apply this queueing model to the admission control, which belongs in the preventive control category of the B-ISDN.

9.2.1 The Space Segment

9.2.1.1 Sensitivity of On-Board Buffer Size Versus Output Transmission Rate

In these computations, there are eight earth stations in the satellite network. There are three types of earth stations: trunk, bursty, and mixed. The CLR for each earth station is less than 10^{-9}. The output utilization of the earth stations is between 66.75% and 80%. As can be seen in Figure C-40, the cell loss ratio of the space segment is sensitive with respect to the output transmission rate. When the output transmission rate of the space segment is equal to 2, the number of buffers required on board to achieve CLR < 10^{-9} is more than 28. When the output transmission rate of the space segment is equal to 3, the number of buffers needed on board to reach CLR < 10^{-9} is less than 12. The output utilization for output rate equal to 2 is around 70%, and the output utilization for output rate equal to 3 is about 50%. Hence, to maintain the QOS of ATM cells on board the satellite, increasing the output transmission rate the number of buffer requirements on board will be decreased significantly with the cost of decreasing efficiency.

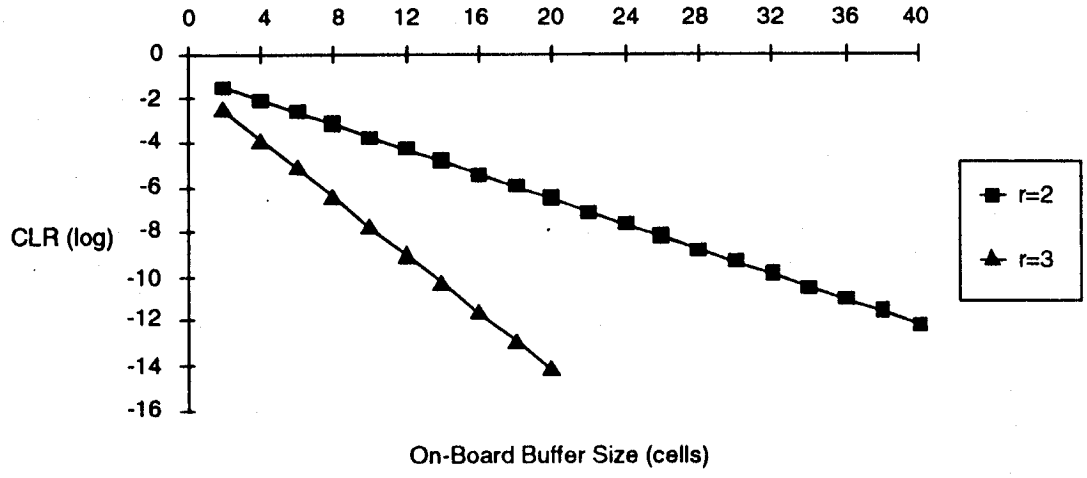

Figure C-40. Space Segment Cell Loss Ratio vs. On-Board Buffer Size for Different Output Transmission Rates

9.2.1.2 CLR of the Satellite Syetem

The overall CLR for the satellite system can be computed if CLR of the earth station and CLR of the space segment are known. (The cell loss due to satellite transmission errors is ignored.) Denote the CLR for the earth station as CLR(E) and CLR for the space segment CLR(S). The overall CLR can be calculated as follows.

$$CLR = 1 - (1 - CLR(E))(1 - CLR(S)), \text{ or}$$

$$CLR = CLR(E) + CLR(S) - CLR(E)CLR(S)$$

This result is shown in Figure C-41. It can be seen that the overall CLR is always dominated by the larger one of CLR(E) and CLR(S). Hence, to improve the overall CLR, the enhancement of CLR for both the earth stations and the space segments must be achieved.

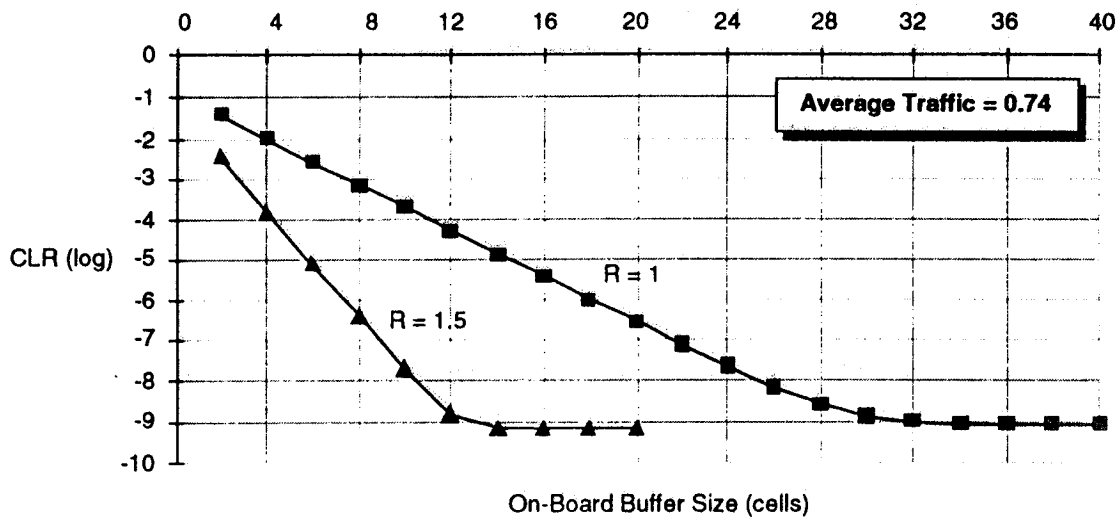

Figure C-41. System Cell Loss Ratio vs. On-Board Buffer Size for Different Output Transmission Rates

9.2.2 Admission Control Based on the Traffic Descriptors

Connection admission control is done at the call setup phase. The setup message for signaling contains traffic characteristic information (peak rate, average rate, burstiness, and peak duration) and a required QOS class. Based on these parameters, connection admission control determines whether the call is accepted or not based on the knowledge of the network loading, required bandwidth, delay/loss performance requirements, and parameters subject to usage parameter control. The function of usage parameter control is to monitor the traffic volume entering the network so that the parameters which have been agreed upon will not be violated. After the bandwidth is allocated, the allocated bandwidth is maintained while the call continues. In the example provided below, the connection admission control is performed based on CLR only. Since the calculation of CLR also depends on the network loading and required bandwidth, observing the change of CLR can also perceive the change of the network loading and the required bandwidth. The initial configuration is that the earth station has 20 connections. Now, the signaling message comes to the earth station and requests a bandwidth so that the 21st connection can be set up. The earth station will use the queueing model and calculate the CLR based on the traffic descriptors. It can be seen from Figure C-42 that when the buffer size is 28 and CLR is less than 10^{-9}, the line with burstiness 5.26 is accepted, and the other line with burstiness 8.86 is rejected.

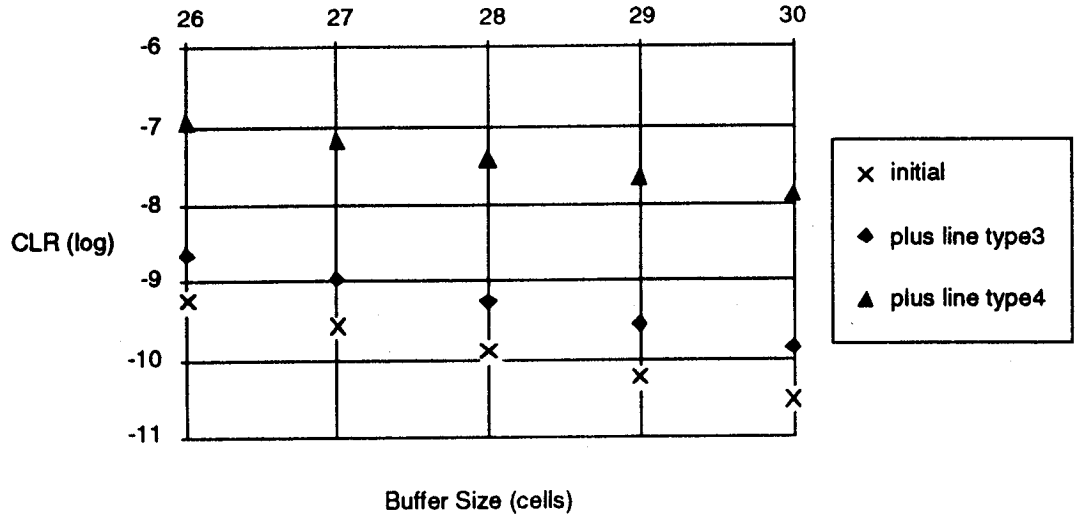

Figure C-42. Example of Connection Admission Control

Alternately, admission control can be performed with the use of a neural network. In this approach, the computation results (the training sets) will be provided to the neural network. After the neural network is trained, it can decide whether to accept the call based on the traffic descriptors only. However, the computation of the ATM performance is still the prerequisite for the training sets.

Appendix D

Traffic Simulation of Fast Packet Switched Networks

1 Introduction

This appendix presents a performance performance analysis of B-ISDN traffic through the advanced satellite communications system, where the earth station is equipped with ATM interfaces and the space segment with a fast packet switch.

In ATM, the packet called a "cell" consists of a 5-byte header and 48-byte information payload. The routing information of an ATM cell consists of a 24-bit virtual path identifier and virtual channel identifier (VPI/VCI) at user network interface (UNI) and a 28-bit VPI/VCI at the network node interface (NNI). To utilize the strength of the VPI with a local significance and to avoid VPI retranslation on-board, the satellite virtual packet (SVP) concept has been introduced in "ATM Cell Packetization Issues and Alternatives". The SVP is created by appending a header, which contains a routing tag for the on-board switch, to one cell or a group of cells destined to the same downlink beam. In this section, the pros and cons of using the SVP concept are reviewed.

Since the satellite resources are bandwidth limited, the design focus is to increase the downlink beam utilization. To achieve this objective, one approach is to increase the throughput of the on-board fast packet switch. Different fast packet switch architectures have been proposed to improve the switch throughput [D-1][D-2] . Two fast packet switching architectures are considered in this paper: the input queueing fast packet switch with a nonblocking switching fabric and the output queueing fast packet switch with a nonblocking switching fabric.

For the input-queueing fast packet switch, the throughput is constrained by the head of line blocking problem. Three schemes are studied in this paper to improve the throughput: the increase of the input buffer size, the increase of searching depth of the input queue to resolve the output contention [D-3], and the increase of the switch speed. For the output-queueing fast packet switch, the incoming packets are not stored in the input buffers, the switch has to operate N times faster than the line speed to avoid the output contention problem, where N is the size of the switch.

The main focus of this section is the performance analysis of the ATM cells (the cell transfer delay, cell delay jitter, cell delay distribution) through the multiplexer at the earth station and the fast packet switch at the satellite using the SVP concept.

2 Trade-off Analysis of Ground SVP Preprocessing and On-Board SVP Processing

Satellite virtual packets (SVPs) are created by appending a header to one cell or a group of cells destined to the same down link beam at the earth station (see Figure D-1) for unified routing, control and management purposes. The header of the SVP is termed as the satellite virtual label (SVL). The SVP format alternatives, the SVP sizing issues, and the SVL format have been discussed in "ATM Cells Packetization Issues and Alternatives" section. The trade-off between ground preprocessing and on-board processing are reviewed.

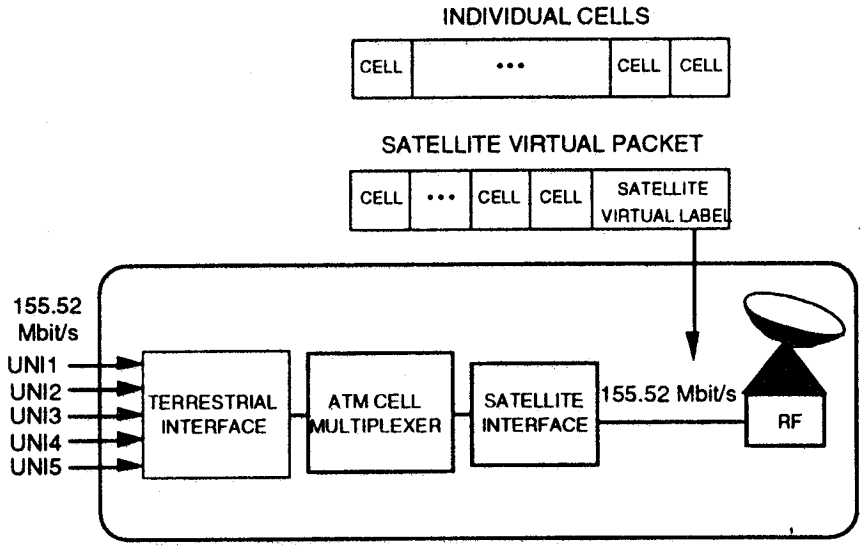

Figure D-1. ATM Earth Station Configuration

a. VPI/VCI Processing

If ground ATM cell preprocessing is performed, the routing tag of the cells for the on-board switch can be prepended at the earth station. Therefore, VPI/VCI processing is not necessary at the space segment since the packet can route through the switch using the prepended routing tag.

b. ATM HEC Processing

To perform VPI/VCI processing, the ATM cell header has to be decoded first. Within the 5-byte ATM cell header, there is a 1-byte header error control (HEC). The HEC decoding procedure has to follow the CCITT recommendation I.321. If ATM cells are preprocessed at the earth station, the space segment only has to decode the SVL instead

the cell header. The FEC of the SVL can be designed in such a way it is more suitable for the satellite environment.

c. Ground Preprocessing Complexity and On-Board Processing Complexity

For pure ATM cell transmission, the earth station has to provide the ATM physical layer function, as per CCITT Recommendation I.321. The cell delineation (cell synchronization) and HEC verification are two functions of the ATM physical layer. To preprocess the ATM cells and format them into SVPs, the earth station has to provide ATM physical layer and ATM layer functions. One of the functions in the ATM layer is VPI/VCI translation.

If the ATM cells are preprocessed at the earth stations, no ATM layer functions are required at the space segment since the SVP can be routed and processed using the SVL. Another advantage of preprocessing the ATM cells is that no cell delineation and cell HEC verification in the ATM physical layer functions are necessary on-board since the ATM cells are the payload of the SVPs. If the SVPs are in synchronization, it is guaranteed that the ATM cells are also in synchronization.

If ATM cells are not preprocessed at the earth station, the space segment has to provide all the functions of the ATM physical layer and ATM layer.

d. Link Utilization Efficiency

Pure ATM cell transmission through the satellite links achieves the highest link utilization. However, as discussed above, the space segment processing complexity for pure ATM cell transmission is high.

The SVL of the SVP is an overhead. The larger the SVL, the worse the satellite link utilization efficiency. For a fixed SVL, if the number of the packets contained by the SVP is larger, the link utilization efficiency is higher.

e. Reserved Slot for Circuit Emulation Services

With SVP concept. it is possible to reserve cells for circuit emulation services as the reserved slots used in circuit switching. The result is cell delay jitter is reduced. This is an effective scheme as long as the cells do come in to the earth station periodically so that the reserved cells in the SVP are not wasted.

f. Bit Interleaving

Protection of ATM cells and SVPs from burst errors have been discussed in "Protection of ATM Cell Headers from Burst Errors" section. In order to effectively combat the burst errors, the interleaving depth can not be too short. Therefore, the bit interleaving delay is an unavoidable component of the cell end-to-end delay. By performing SVP formatting at the earth station, the bit interleaving procedure can be naturally achieved among the cells in the same SVP. The SVP packetization delay is equivalent to the writing phase of the bit interleaving delay of cells; the SVP transmission delay is equivalent to the reading phase of the bit interleaving delay of cells. If bit interleaving

of ATM cells are required at the earth stationIn conclusion, performing SVP group preprocessing is more advantageous over pure ATM cell transmission.

3 Performance Analysis of SVP Ground Preprocessing and On-Board Processing

In this subsection, the cell delay through a multiplexer and a fast packet switch are collected from the simulation models to study the performance of SVP transmission through the satellite B-ISDN.

A brief description of the satellite B-ISDN simulation model (see Figure D-2) is given below. The number of earth stations, uplink beams, and downlink beams is assumed to be the same as the size of the on-board switch. Each earth station is interfaced with five user network interfaces (UNIs) or network node interfaces (NNIs). Each UNI or NNI generates cells following Possion distribution. The uplink and downlink access schemes are using time-division multiplexing (TDM). The link transmission speed is 155.52 Mbit/s. The on-board switch speed is (speed-up factor * 155.52 Mbit/s), where speedup factor is the ratio of the switch speed to the link speed. The cell slot time is equal to 2.726 μs when the transmission rate is 155.52 Mbit/s. In the performance figures, CDJ designates the cell delay jitter, CTD the cell transfer delay, u the link utilization, d the checking depth, n the number of downlink beams, and s is the speedup factor of the switch.

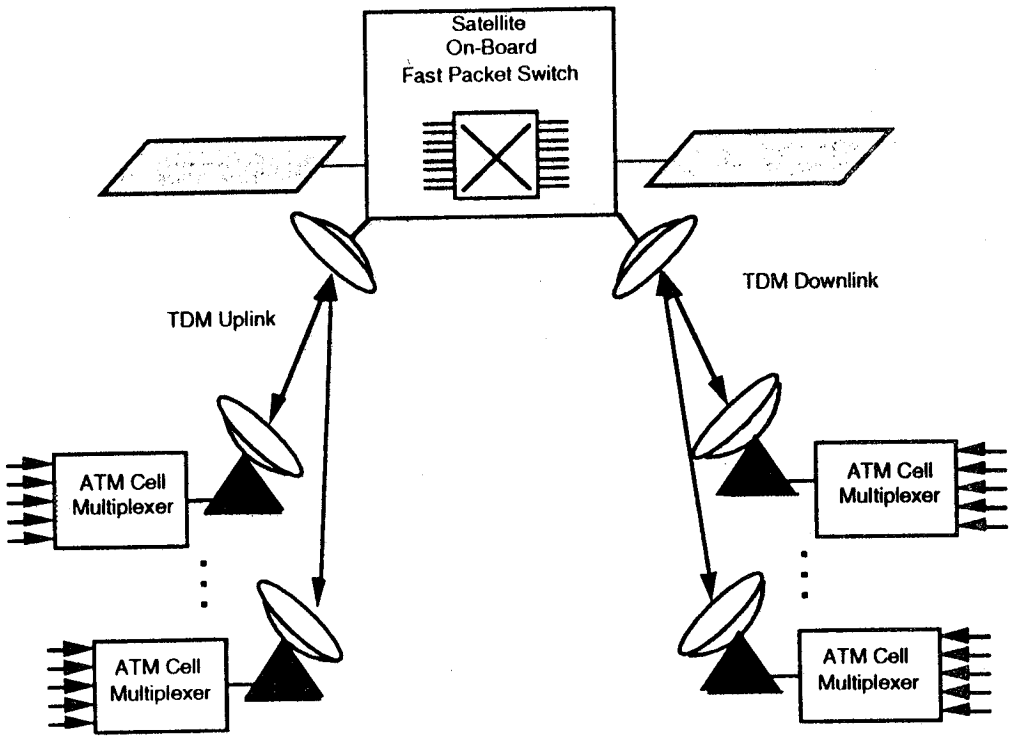

Figure D-2. Satellite B-ISDN Simulation Model

3.1 SVP Transmission Versus Cell Transmission

Earth stations in the satellite B-ISDN are interfaced with different user network interfaces (UNIs) and network node interfaces (NNIs). To increase the transmission link efficiency and share the satellite transmission link among different users, the advanced earth station functions as a statistical multiplexer. At the earth station, it is assumed that the buffer size of the multiplexer is infinite. Under this assumption, its cell loss ratio is zero.

The first set of results illustrate the effect of different sizes of SVPs on the cell delay performance and the buffer size requirements at the earth station for different transmission link utilization and different numbers of downlink beams.

Figure D-3 shows the cell delay jitter, cell transfer delay, cell 90 percentile delay of the statistical multiplexer at the earth station for different sizes of SVPs and different output link utilizations when the number of downlink beams is 8. As shown in this figure, for different sizes of SVPs, the increase of the cell delay is proportional to the size of SVPs. Based on these results, the size of SVPs should be small if cell delay is an important quality of service parameter. Also as shown in this figure, when the output link utilization is higher, the delay performance of the SVP transmission improves while the delay performance of the single-cell transmission degrades. The reason is that when the output link utilization is higher, the probability that an SVP will be filled with cells destined to the same downlink beam within a given time is also higher. Hence, the delay performance of the SVPs improves when the output link utilization is higher. Basically, the cell delay for the SVPs consists of three elements: the packetization delay (the time required to fill the SVP with cells), the waiting time for transmission in the queue, and the transmission time. From Figure D-3, most of the cell delay for the SVPs occurs during the packetization process.

In conclusion, the cell delay performance for single-cell transmission is determined by the waiting time in the queue for transmission, i.e., the queueing delay. The cell delay performance for SVP transmission is largely determined by the packetization delay. This implies that if SVP transmission concept is used for the satellite network, the satellite network link has to be operated at very high utilization (above 80%); therefore, the on-board fast packet switch throughput has to be higher than 80%. The other alternative is to provide some corrective measures such that the delay of SVP transmission through the earth station can be minimized.

Figure D-4 shows the cell delay jitter, cell transfer delay, cell 90 percentile delay of the statistical multiplexer at the earth station for different SVP sizes and different output link utilizations when the number of downlink beams is 16. Compared Figures D-3 and D-4, it can be concluded that the cell delay for SVP transmission is proportional to the number of downlink beams. The reason is, as previously mentioned, the cell delay for SVP transmission is dominated by the packetization delay. If there are more downlink beams, the filling rate for each SVP destined to different downlink beam is reduced under the condition that the output link utilization is kept constant.

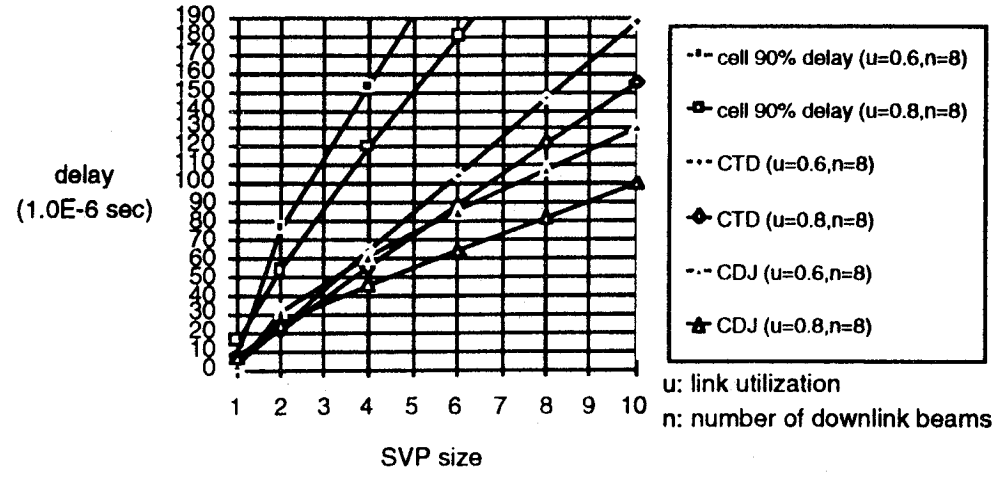

Figure D-3. Cell Delay Performance vs. SVP Sizes for Different Output Link Utilization

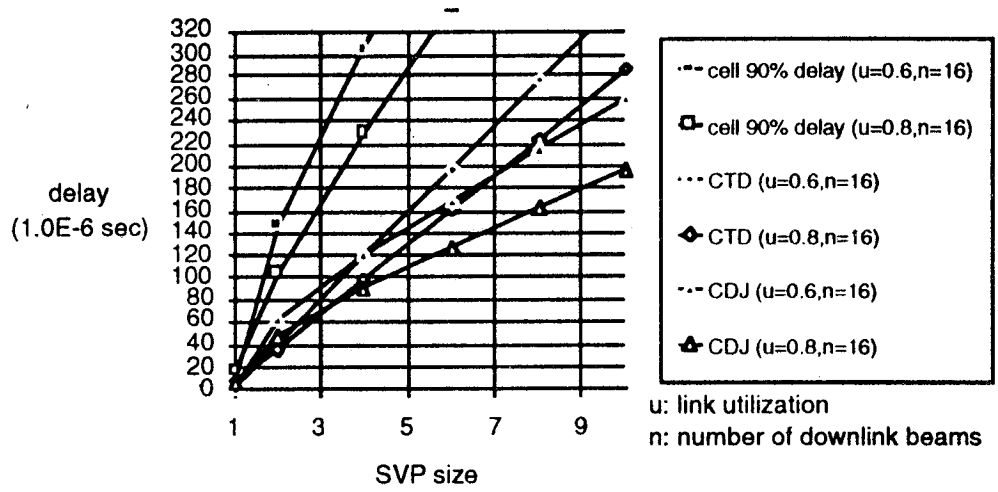

Figure D-4. Cell Delay Performance vs. SVP Sizes for Different Output Link Utilization

Figures D-5 and D-6 show the cell delay probability mass function and cell delay cumulative distribution function for different SVP sizes when the number of downlink beam is 8 and the output link utilizations are 60% and 80%, respectively. From the cell delay distribution, the buffer size requirement for the earth station to achieve a certain cell loss ratio can be derived. For example, for a 4-cell SVP and utilization equal to 0.6, the probability that the cell delay is 400 μs is about 10^{-4}. Hence, to have a cell loss ratio of 10^{-4}, the size of buffer required for the statistical multiplexer is about 37 SVP

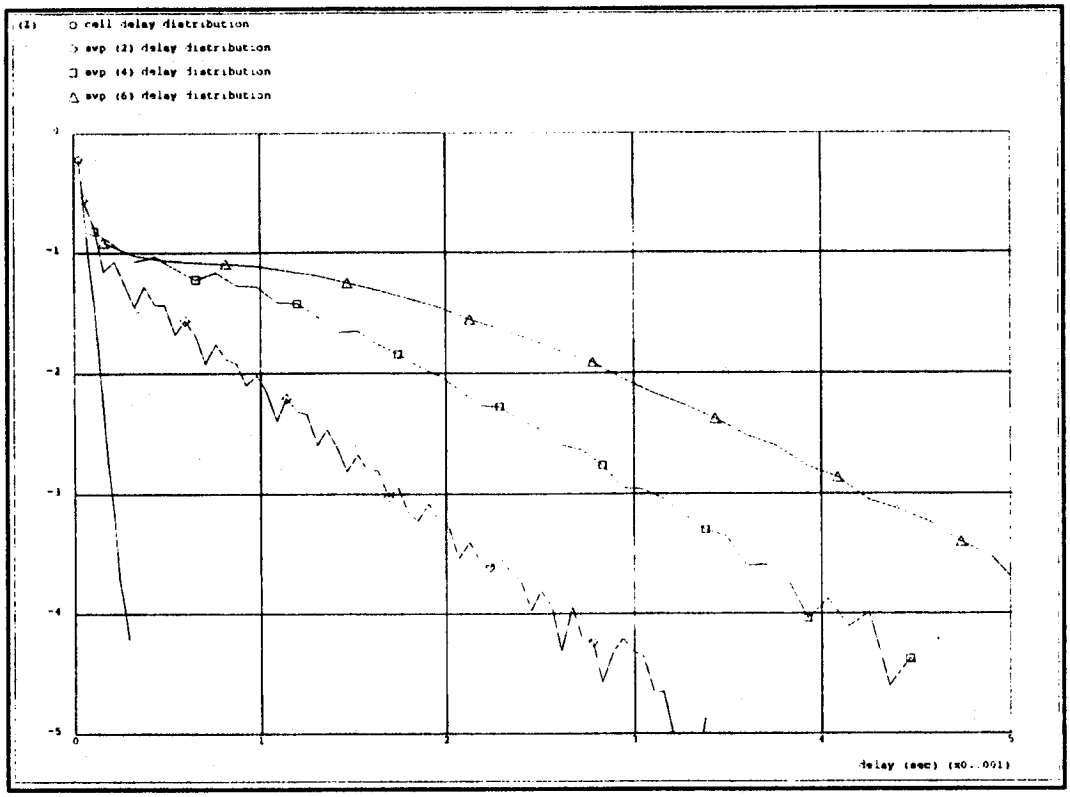

Figure D-5a. Cell Delay Distribution for Different SVP Sizes When the Output Link Utilization is 60 Percent

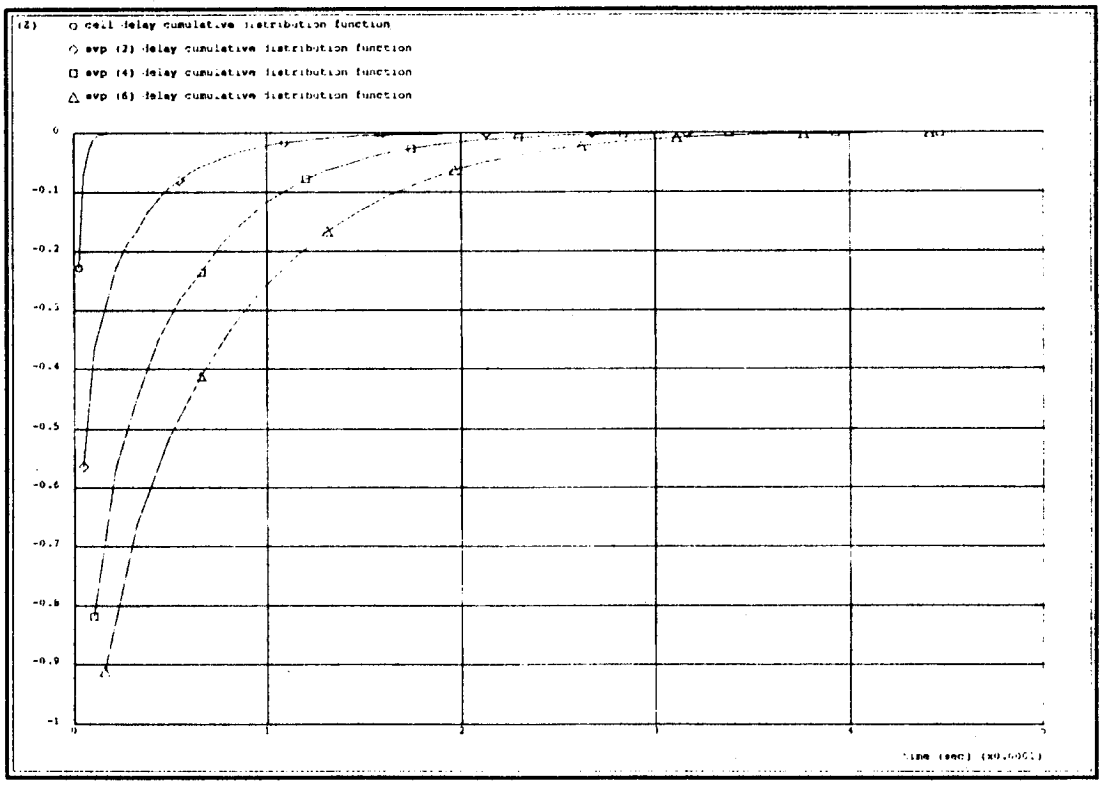

Figure D-5b. Cell Delay Cumulative Distribution for Different SVP Sizes When the Output Link Utilization is 60 Percent

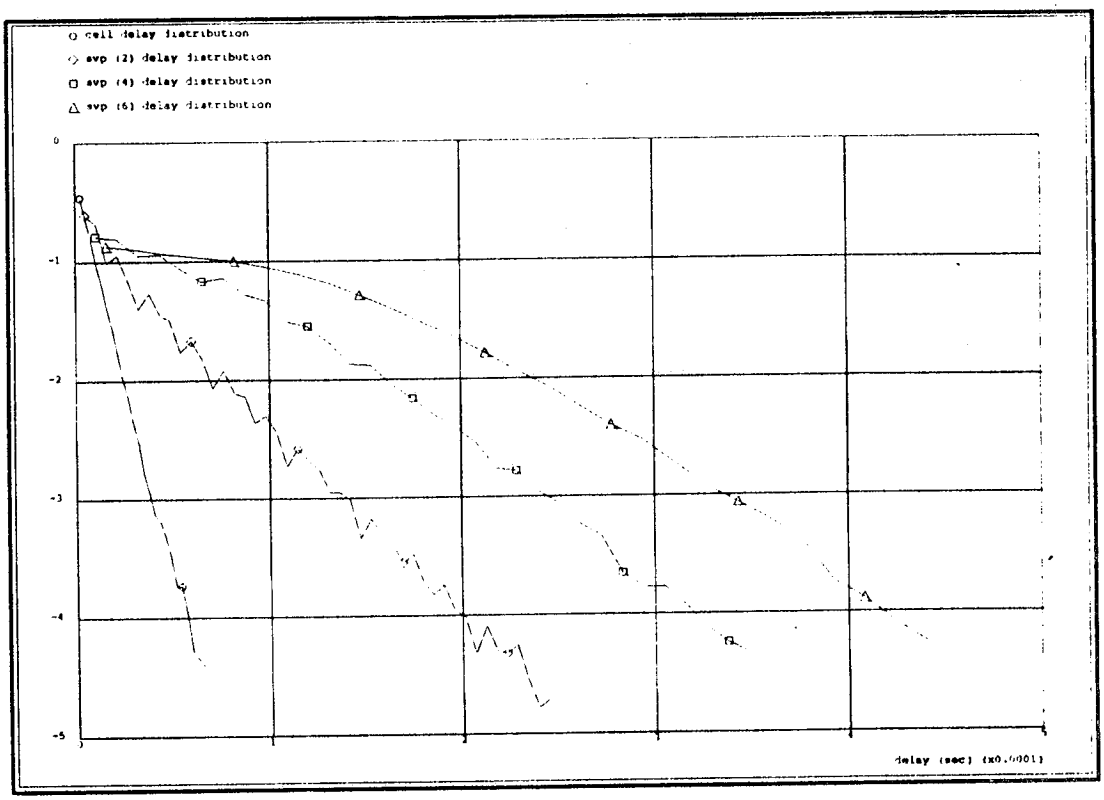

Figure D-6a. Cell Delay Distribution for Different SVP Sizes When the Output Link Utilization is 80 Percent

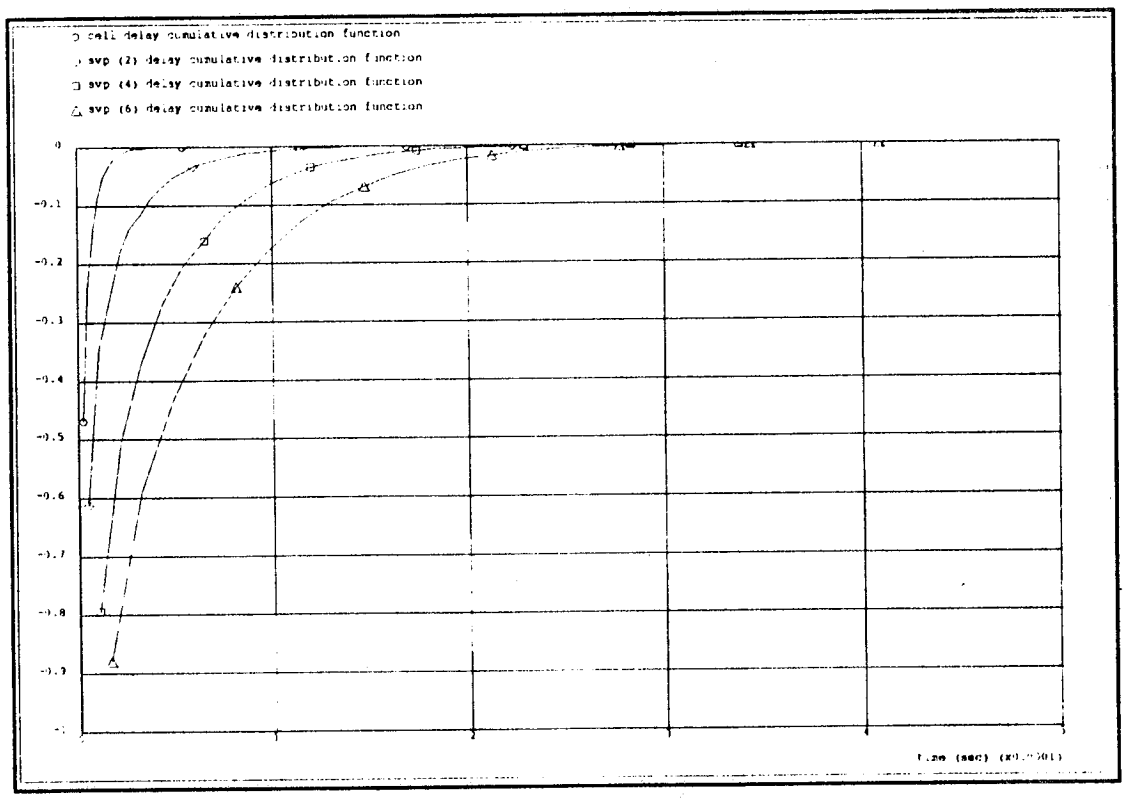

Figure D-6b. Cell Delay Cumulative Distribution for Different SVP Sizes When the Output Link Utilization is 80 Percent

packets. Comparing Figures D-5 and D-6, with higher utilization, the cell delay distribution curves of the SVPs are approaching closer to the cell delay distribution for single-cell transmission.

3.1.1 Corrective Measures to Reduce the SVP Packetization Delay

As mentioned in the previous subsection, the packetization delay of formatting the SVP at the earth station dominates the cell delay through the earth station. There are several ways of resolving this problem.

The first scheme is to use a SVP with a variable size and there is a timer associated with each SVP. If the timer expires, the SVP will be sent out with the cells currently in the SVP. The minimum size of SVP is one cell and a constraint can be placed on the maximum size of SVP. However, this variable packet scheme changes the whole satellite network processing requirement compared to the fixed size. For example, the on-board switch has to be operated in mini-slotted mode, the packet transfer protocol through the switch has to know the beginning and the end of a packet, the timing signals generated for memories and interface processors are more complex, just to name a few. One variation is to segment a variable-length SVP into several fixed-size SVPs on-board so that the on-board processor still operates in the fixed-slot mode. Nevertheless, the memory requirement and processing complexity make this approach very unattractive for on-board application.

One feasible alternative of the above scheme is to allow only several sizes of SVPs to be used within the satellite network, for example sizes 1 and 2. In this situation, the overall processing requirement is reduced since the packet size can be determined using a 1-bit size indication field. In this example, the operation of the switch has to be able to operate in single-slot mode and two-slot mode concurrently. Assume the slot size is the same as the size of single-cell SVP. If the incoming packet is a 2-cell SVP, then the operation of the switch has to be modified to accommodate that the packet size is larger than the switch slot size. One solution is that the input port has to reserve two slots in advance in order to successfully transmit the 2-cell SVP. This modification can be achieved easily with the token ring reservation scheme [D-3]. In this situation, the token generator generates two streams of tokens: one stream is for the current slot and one stream is for the next slot. Each token represents one output port. The 2-cell SVP has to reserve two tokens in two different streams for the same output port simultaneously to successfully transmit the SVP. The output port examines the size indication field and assemble the SVP back to the original size after the second cell has been received. If the incoming packet is a single-cell SVP, the operation of the on-board switch is normal. The single-cell SVP only processes the token at the first token stream, i.e. the current-slot token stream. The other solution is to segment a 2-cell SVP into two single-cell SVPs on-board. After segmentation, the header of the 2-cell SVP has to be copied to each single-cell SVP. The output port examines the size indication field and performs assembly if necessary.

The most feasible approach is to use a fixed size SVP approach. To reduce the SVP packetization delay, there is a timer associated with each SVP. If the timer expires and the SVP has not been filled with cells, the SVP will be padded with empty cells and sent out. The disadvantage of this approach is that the link utilization efficiency is low if the percentage of empty cells among the data cells is high. The other problem is that the padding of empty cells increases the multiplexer output link utilization. The multiplexer transmission rate, i.e., the uplink transmission rate and the timer value for the SVPs have to be designed very carefully. Otherwise, although the packetization delay can be minimized, the queueing delay at the multiplexer is increased due to the packet arrival rate is increased. It is possible that the multiplexer queue will be saturated if the timer-value is very small due to the nature of statistical multiplexing.

Based on the above discussion, another set of simulation is performed to study the impact of the timer value to the fixed-size SVP delay at the earth station. Note that in this set of simulation, the uplink SVP transmission at the earth station is synchronized based on the SVP size. The previous simulation results and the following simulation results, the uplink SVP transmission at the earth station is synchronized with single-cell size. The results show that there is no difference for the cell delay jitter of SVP at the earth station. The cell transfer delay difference between these two cases is the synchronization delay for transmission.

Figure D-7 shows the the effect of different timer value on the multiplexer output link utilization for two different input link utilizations when the SVP size is 2. The input link utilization is defined as the sum of the link utilizations (information rate) of the UNIs. It can be seen that with a smaller timer value the multiplexer utilization increases significantly.

Figure D-8 shows that the effect of different timer value on the cell delay jitter through the multiplexer when the SVP size is 2 and the input link utilization is 0.6. By reducing the time-out value, the SVP packetization delay is reduced; however, the SVP queueing delay is increased since the output link utilization is increased. There is a time-out value for each input link utilization to optimize the cell delay through the earth station under Possion traffic assumption. Since the input link utilization is low, the timer value can be largely reduced to minimize the cell delay jitter. In this situation, the packetization delay jitter has been reduced from 28 μsec to 12 μsec while the queueing delay jitter is only increased from 4 μsec to 6 μsec when the timer is reduced from infinite to 25 μsec. Figure D-9 show that the effect of different timer value on the cell delay jitter through the multiplexer when the SVP size is 2 and the input link utilization is 0.8. Since the output link information utilization is already high, the timer value can not use a small value otherwise the multiplexer will reach saturation and the queueing delay will be increased exponentially. The use of a 50 μsec timer reduces the cell delay jitter form 24 μsec to 19 μsec, which is not very significant.

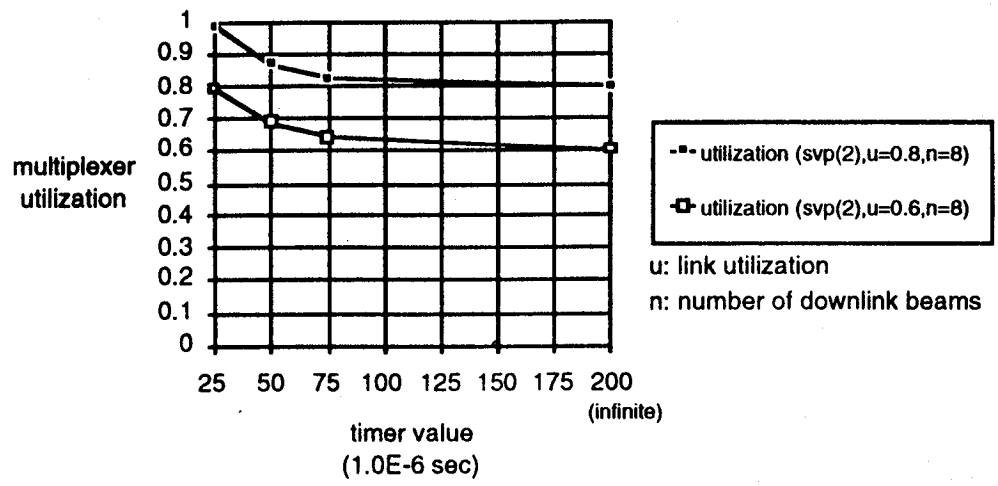

Figure D-7. Multiplexer Utilization vs. Time-out Value for Different Input Link Utilization

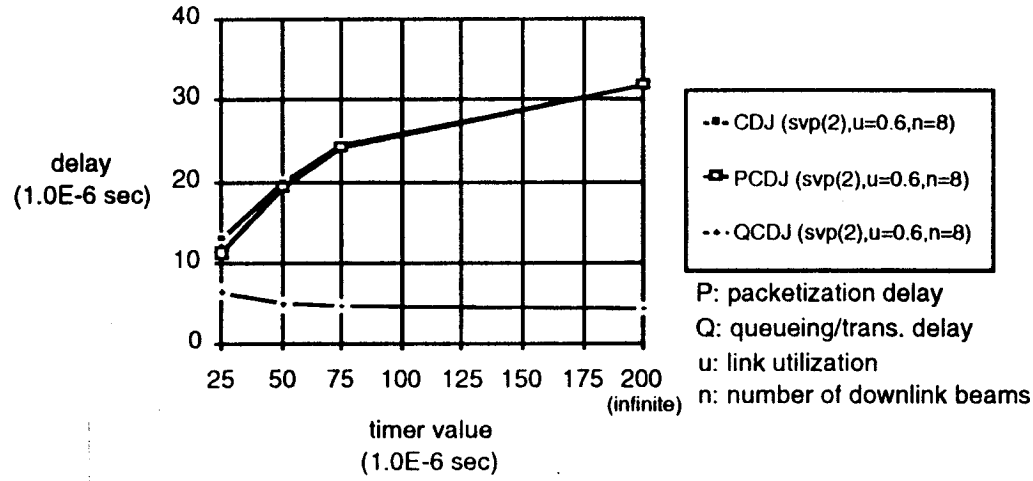

Figure D-8. Cell Delay Jitter vs. Time-out Value When the Input Link Utilization is 60 Percent

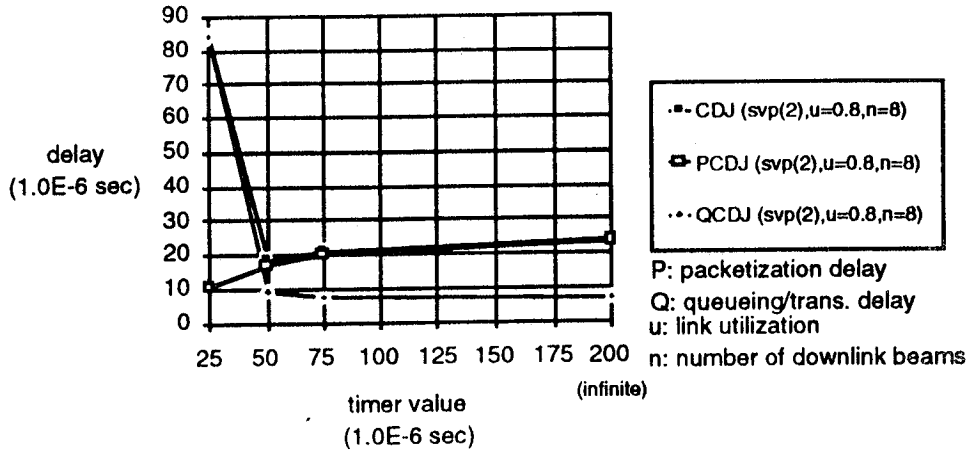

Figure D-9. Cell Delay Jitter vs. Time-out Value When the Input Link Utilization is 80 Percent

Figures D-10, D-11, and D-12 show the performance of 4-cell SVPs for different timer values and different input link utilizations. When the input link utilization is 0.6, the use of a 67.5 μsec timer reduces the CDJ from 60 μsec to 28 μsec. When the input link utilization is 0.8, the use of a 100 μsec timer reduces the CDJ from 45 μsec to 36 μsec. Figures D-13, D-14, and D-15 show the performance of 6-cell SVPs for different timer values and different input link utilizations. When the input link utilization is 0.6, the use of a 125 μsec timer reduces the CDJ from 85 μsec to 45 μsec. When the input link utilization is 0.8, the use of a 150 μsec timer reduces the CDJ from 65 μsec to 51 μsec.

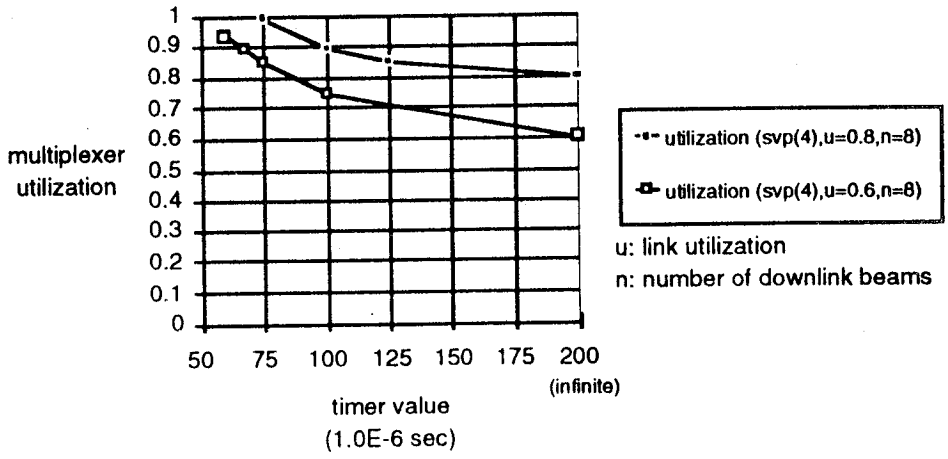

Figure D-10. Multiplexer Utilization vs. Time-out Value for Different Input Link Utilization

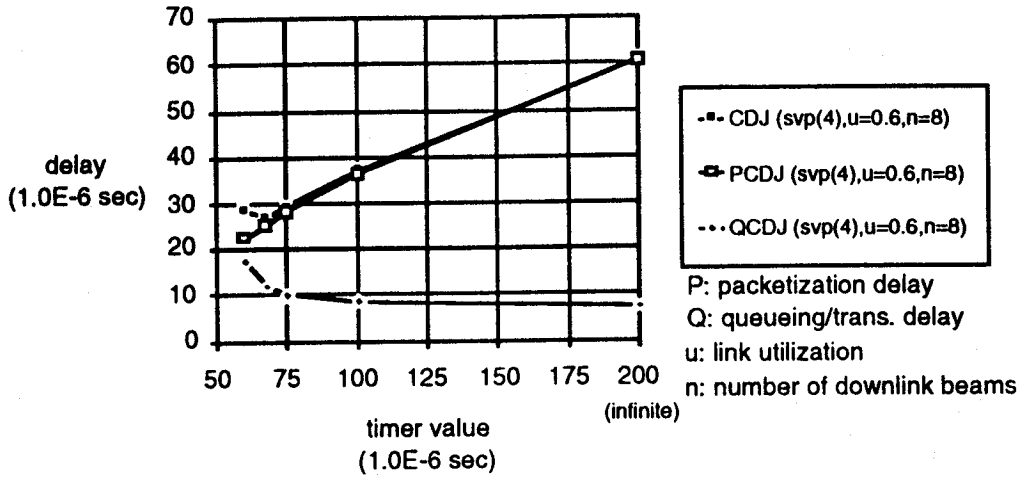

Figure D-11. Cell Delay Jitter vs. Time-out Value When the Input Link Utilization is 60 Percent

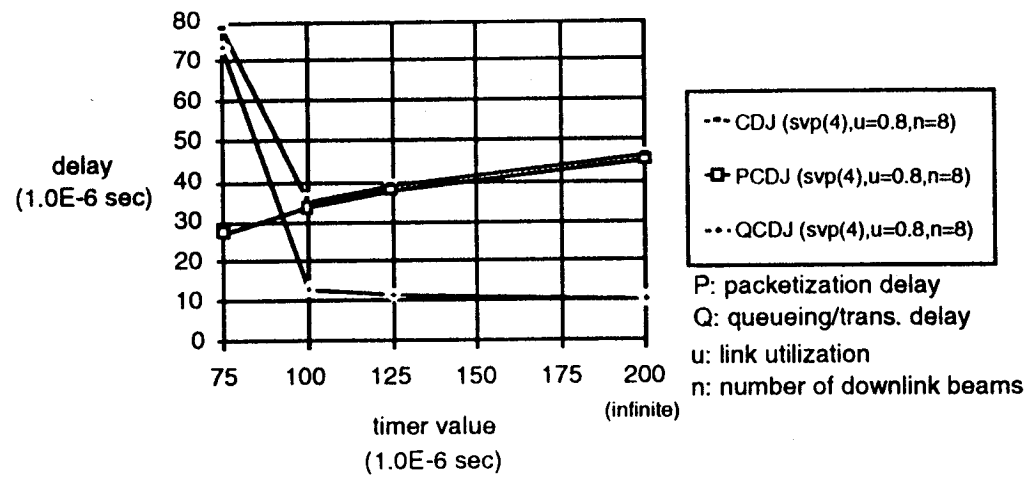

Figure D-12. Cell Delay Jitter vs. Time-out Value When the Input Link Utilization is 80 Percent

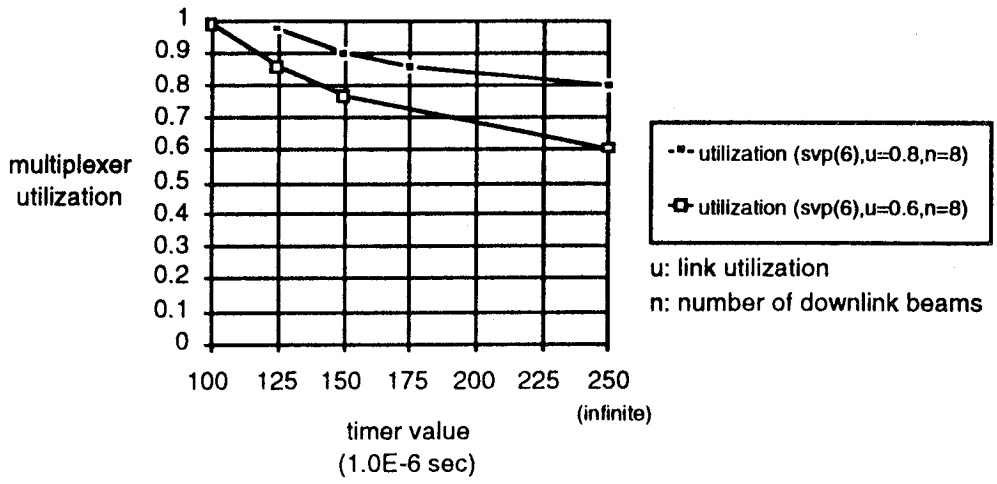

Figure D-13. Multiplexer Utilization vs. Time-out Value for Different Input Link Utilization

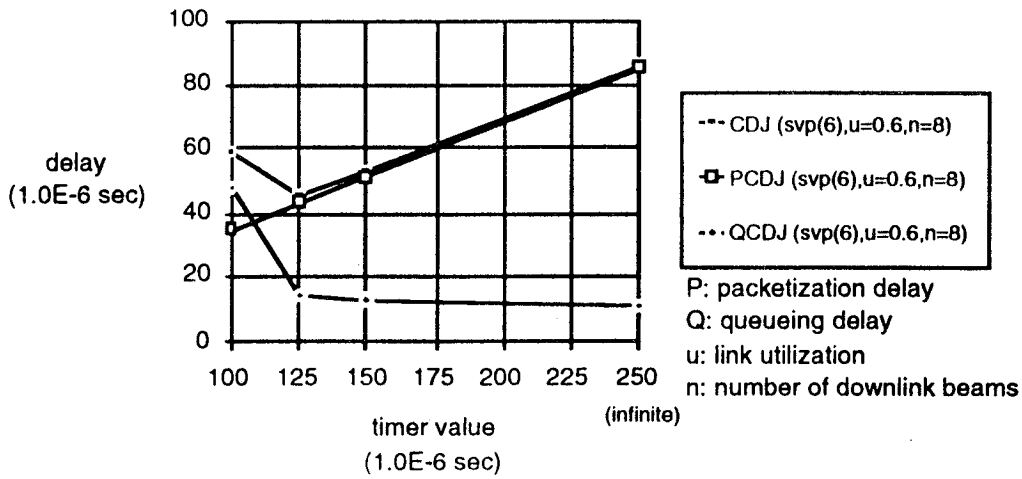

Figure D-14. Cell Delay Jitter vs. Time-out Value When the Input Link Utilization is 60 Percent

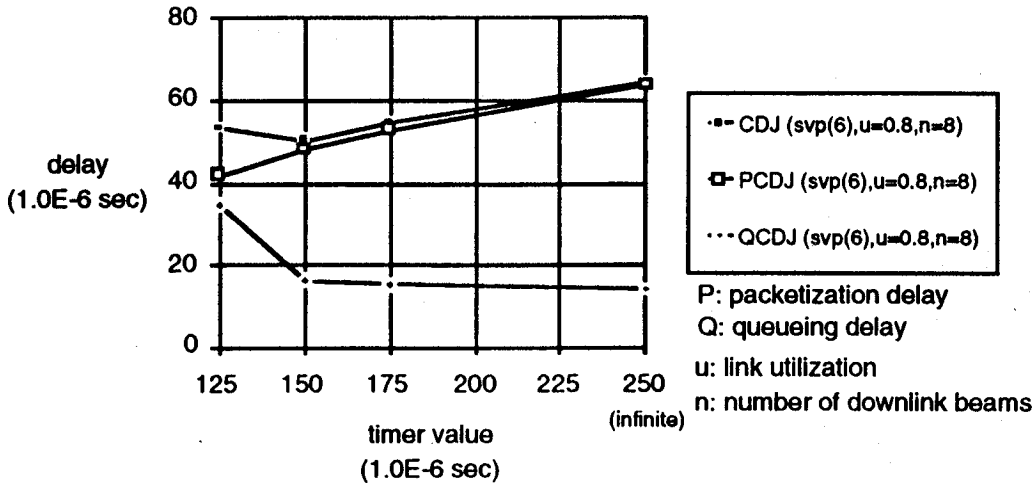

Figure D-15. Cell Delay Jitter vs. Time-out Value When the Input Link Utilization is 80 Percent

From these figures, the conclusion is that if the input link utilization is low, then using a small timer value can largely reduce the cell delay jitter at the earth station. If the input link utilization is already high, the effect of using a timer is not significant. However, it is always a good practice to put a timer for each SVP to prevent that a SVP packetization process may take a long time.

3.2 On-Board Fast Packet Switch Performance

For the on-board fast packet switch, since the mass and power are the constrained design factors for satellites, the buffer size of the switch has to be finite. Also since the satellite communications system is both power and bandwidth limited, the downlink beam resource has to be used very efficiently. One way of increasing the downlink beam utilization is to increase the throughput of the switch. The on-board fast packet switch architectures are assumed to be the banyan type switching networks [D-1][D-2]. The switching fabric is assumed to be unbuffered and point-to-point nonblocking. Since output blocking is an unavoidable situation of a packet switch, it is assumed that packets are either buffered at the input ports or packets are buffered at the output ports.

In this subsection, the impact of formatting SVPs at the earth station to the performance of the on-board fast packet switch is analyzed. The effectiveness and feasibility of different schemes to improve the on-board switch throughput are also examined.

3.2.1 Input Buffering

In this scheme, the packets are buffered at the input ports. The throughput of a switch with FIFO input queue is limited by the head of line blocking problem. It has been shown that the theoretical throughput for the input buffering nonblocking point-to-point switch with infinite buffer size is about 0.58 [D-4]. Head of line blocking is a side effect of the results of output blocking. That is, if one packet at the head of queue can not be transmitted due to output blocking, this packet hinders the delivery of the next packet in the queue due to the FCFS nature of the queue, even though the next packet can be transmitted to the destination without any blocking. Three schemes have been studied to improve the throughput of the switch. The first is to increase the buffer size. Intuitively, the larger the buffer size, the better the packet loss ratio, but the worst delay performance. The second method is to use a non-FIFO queue. If the first packet is blocked due to output blocking, the scheduling algorithm will also examine the packets at the back of the first packet. The number of packets examined each time depends on the preset window size or the checking depth. For a normal FIFO operation, the checking depth is 1. The third method is to operate the switch at a higher speed than the link speed.

3.2.1.1 Increasing Buffer Size

The buffer requirement for the on-board switch is determined by measuring the cell delay distribution using an infinite buffer size. From the delay distribution, the buffer size requirement for a specific cell loss ratio can be calculated as shown in the example in Section 3.1. Figures D-16 and D-17 show the cell delay jitter, cell transfer delay, cell 90 percentile delay for different SVP sizes through the on-board switch. The delay degradation of a larger SVP through the fast packet switch is much less than that through the multiplexer at the earth station because, unlike the multiplexer, no packetization process is required at the switch. The cause of delay degradation of a larger SVP is that a SVP transmission time through the switch is larger than a single-cell packet transmission timer.

When the utilization is low, the single-cell delay performance is better than the SVP delay performance because the SVP transmission time through the switch is n times larger than the cell transmission time, where n is the size of the SVP. In general, under the same conditions, the packet delay through a switch for a packet with n cells is n times larger than that for a single cell.

From Figures D-16 and D-17, when the link utilization is higher (approaching the 0.6 switch throughput for an 8 x 8 switch), the cell delay performance degrades much quicker than the SVP delay performance. Eventually, the SVP delay performance is better than the single-cell delay performance because the process of formatting cells into SVPs at the earth station disturbs the traffic pattern. Note that the traffic pattern coming to the simulated satellite network is assumed to be random. Under this scenario, the probability that n SVPs destined to the same destination arrive at the

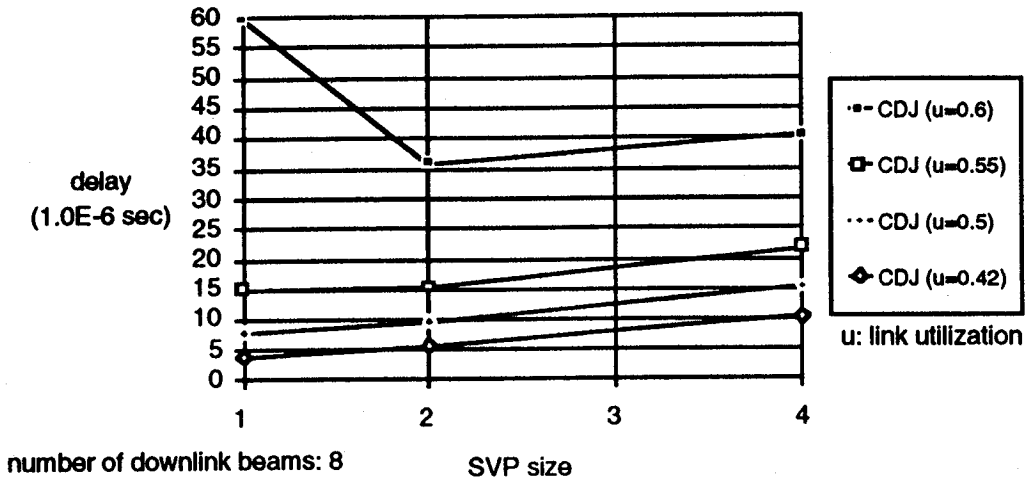

Figure D-16. Cell Delay Jitter vs. SVP Size for Different Input Link Utilization

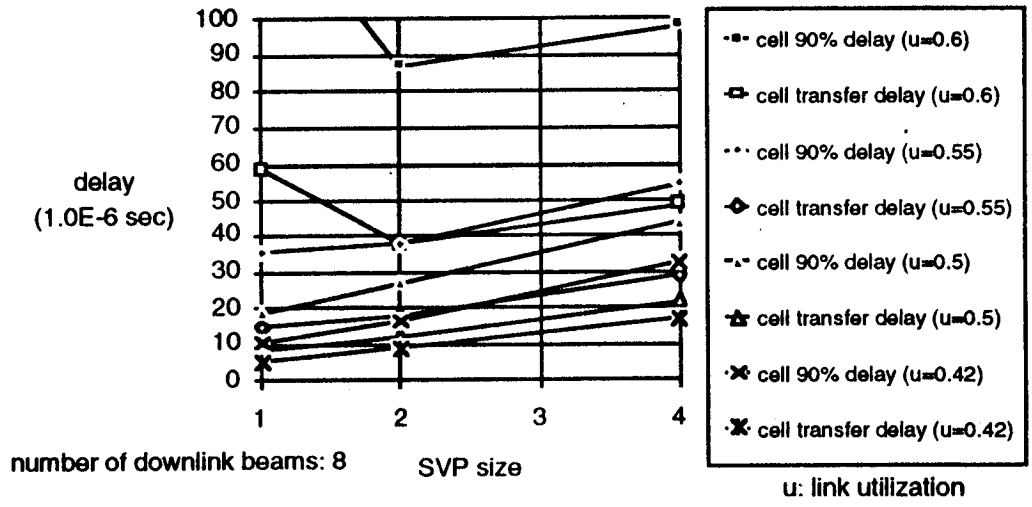

Figure D-17. Cell 90 Percent Delay and Cell Transfer Delay vs. SVP Size for Different Input Link Utilization

input port of the switch continuously is less than that of n cells, where n > 1. Hence, the output contention problem is reduced. This discovery is useful if the traffic coming to the earth station can be assumed to be random as in the case of packet-switched traffic. In this case, the maximum throughput of the on-board switch is increased by formatting the SVP at the earth station since the output contention problem of the switch is reduced.

The cell delay distribution for different SVP sizes and different link utilization is shown in Figure D-18. The results show that when the utilization is low, the single-cell packet has the best performance. When the utilization is high, the 2-cell SVP has the best delay performance. This is because when the link utilization is low, most of the time the packets are very sparse on the transmission link; hence, the packet service time through the switch dominates the delay performance. When the link utilization is high, the probability of continuous arrivals of the packets with the same destination at the input port (called P_c) intensifies the output contention problem. When P_c is high, the average queue length is higher and each packet will experience a higher queueing delay. The larger the SVP size, the smaller the P_c. With a smaller P_c, the average queue length is also reduced. Although the average queue length is reduced, the average queue delay is the sum of the average queue length and the packet service time. Therefore, there is a trade-off between the packet service time and the SVP size to optimize the cell delay under the same input link utilization. In conclusion, the cell delay distribution through the switch is determined by two factors: the packet transmission delay and the probability of continuous arrivals of packets with the same destination at the input port.

From the cell delay distribution provided in Figure D-18, using the extrapolation scheme the buffer size required to achieve CLR 10^{-9} is around 100 for single-cell transmission when the link utilization is only 0.55. To achieve the same CLR, the buffer size requirement will be increased exponentially when the link utilization approaches the switch throughput 0.6. Therefore, increasing the buffer size is not effective for the fast packet switch to improve the throughput and it is not proper for the satellite environment.

3.2.1.2 Non-FIFO Queue

Using a non-FIFO queue with a checking depth 2, the switch throughput can be increased from 0.6 to 0.73. Using a non-FIFO queue with a checking depth 3, the switch throughput can be increased from 0.6 to 0.79. The switch throughput improvement decreases when the checking depth increases.

The cell delay jitter, the cell transfer delay, and the cell 90 percentile delay of the switch for checking depth 2 and 3 are depicted in Figures D-19, D-20, D-21, and D-22. Compared these figures with Figure D-16, the non-FIFO queue is an effective scheme for increasing the downlink beam utilization. In implementation, the maximum checking depth is determined by the processing speed of the routing tag of each packet at the input port queue. Since the scheduling algorithm has to finish the operation in one packet slot time and the number of packets processed by the scheduling algorithm is proportional to the switch size, the maximum checking depth allowed for a smaller switch is larger than that of a larger switch. Since the size of the on-board switch is between O(10) and O(100), the non-FIFO queue with a preset checking depth is proper for satellite application.

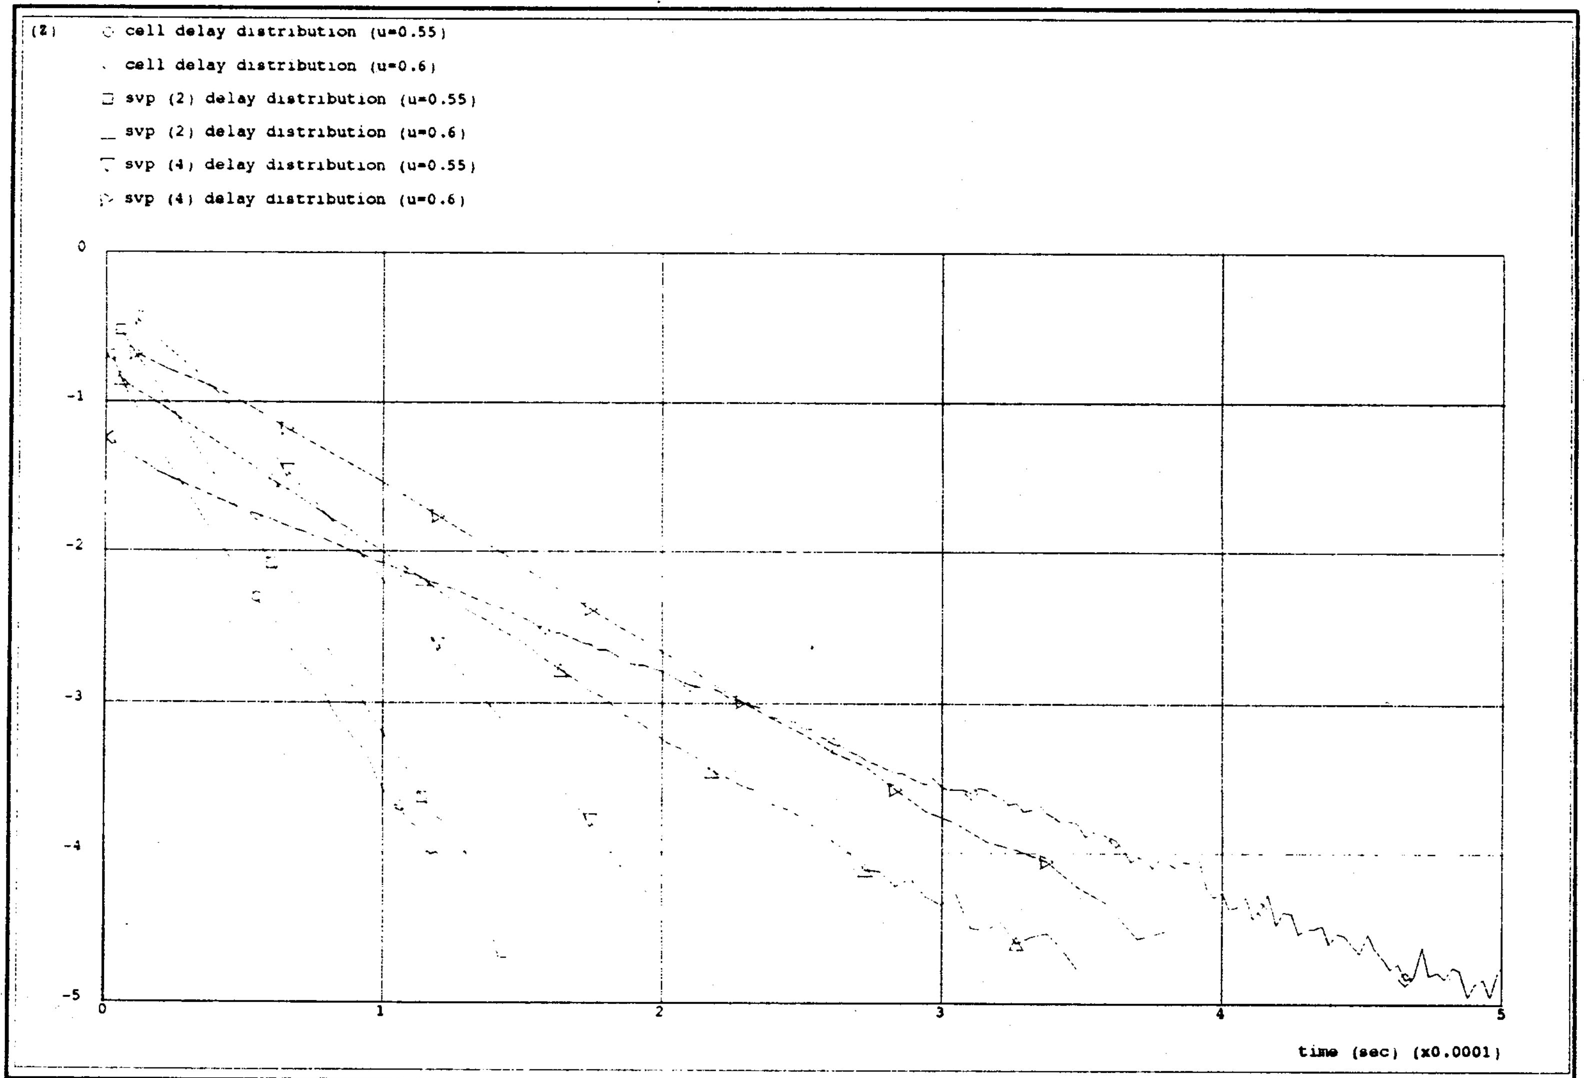

Figure D-18. Cell Delay Distribution for Different SVP Sizes and Input Link Utilization

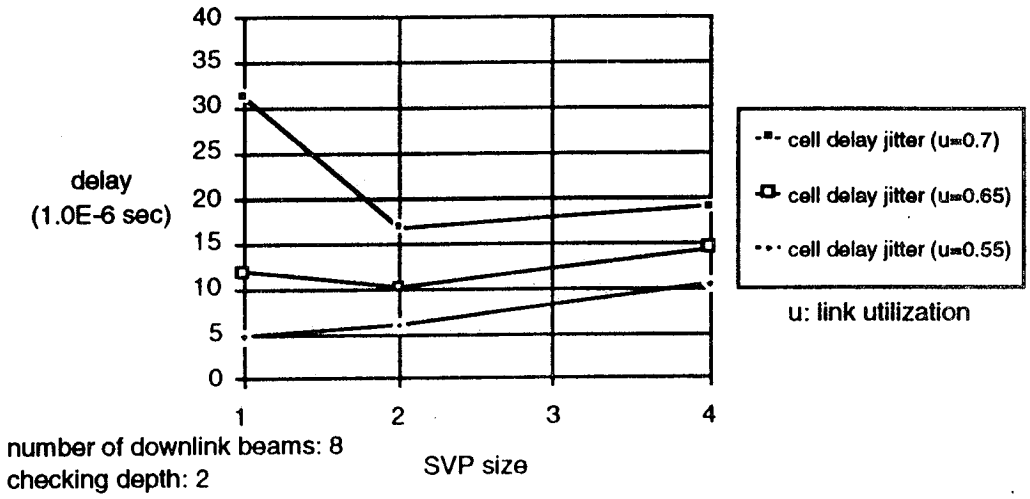

Figure D-19. Cell Delay Jitter vs. SVP Size for Different Input Link Utilization

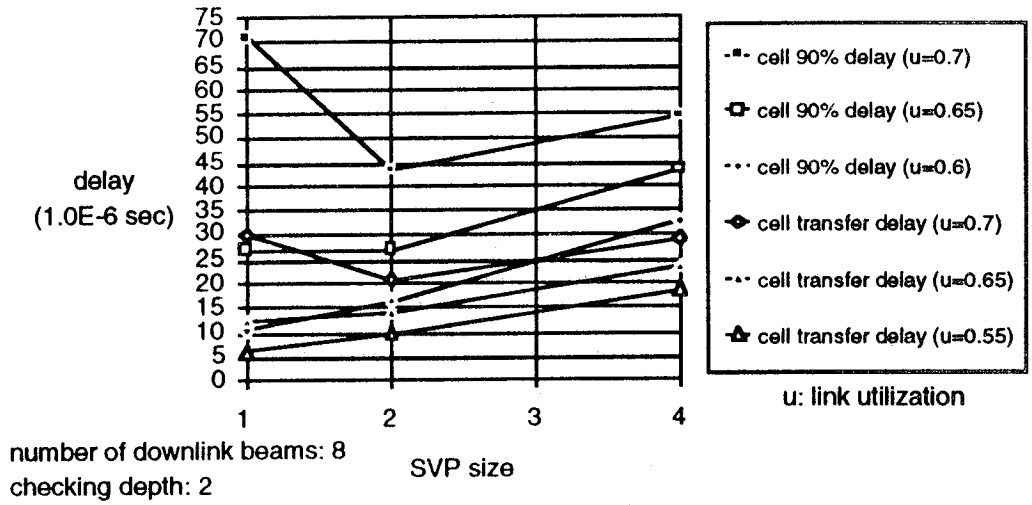

Figure D-20. Cell 90 Percent Delay and Cell Transfer Delay vs. SVP Size for Different Input Link Utilization

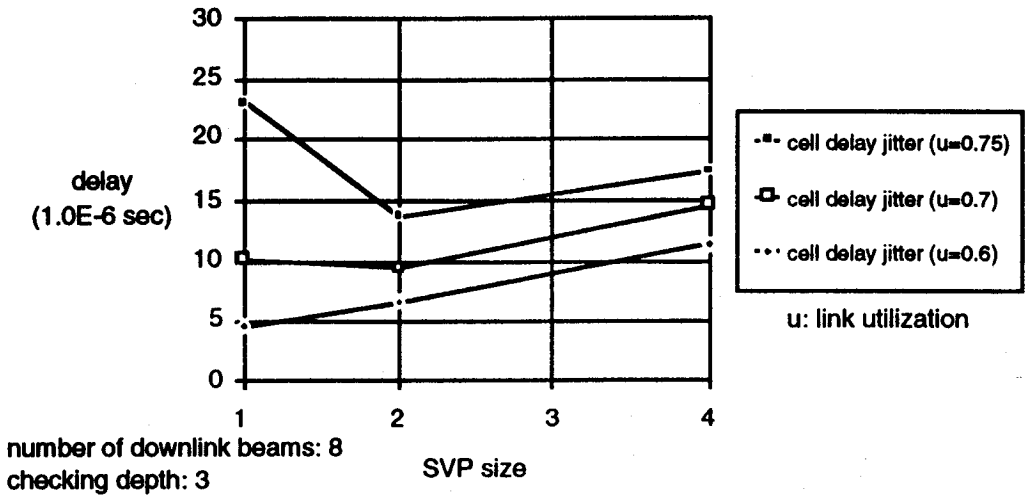

Figure D-21. Cell Delay Jitter vs. SVP Size for Different Input Link Utilization

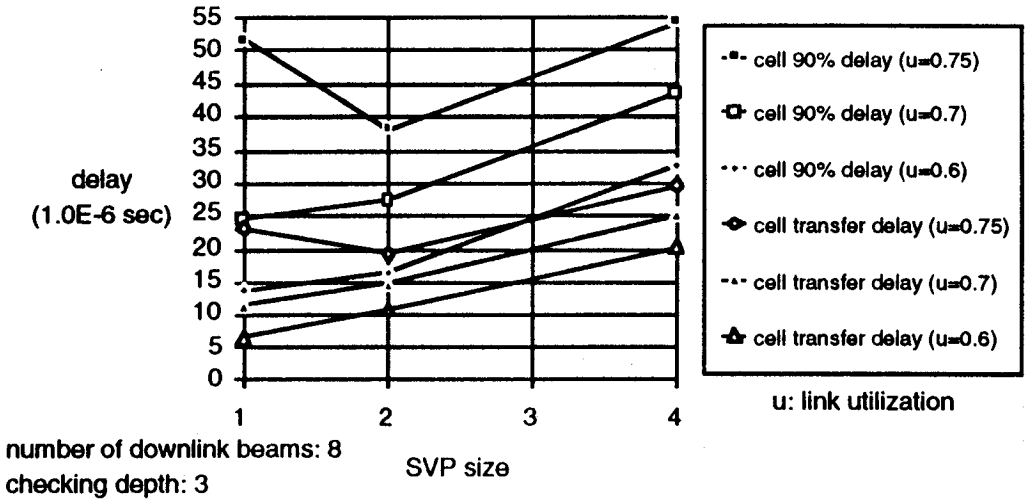

Figure D-22. Cell 90 Percent Delay and Cell Transfer Delay vs. SVP Size for Different Input Link Utilization

The delay distribution comparison for different SVP sizes and different link utilizations when the preset checking depth is 2 is provided in Figure D-23. Basically, as discussed before, the delay distribution is determined by two factors: the packet transmission delay and the probability of continuous arrivals of packets with the same destination at the input port queue (called P_c). For a larger packet, the transmission delay is larger. However, for a larger packet, Pc becomes less; hence, the output contention problem is reduced and the result is the average queue length is smaller. Note that the effect of P_c on the cell delay when the link utilization is low is small since the average input queue length is short.

The delay distribution comparison for different SVP sizes and different link utilizations when the checking depth is 3 is provided in Figure D-24.

3.2.1.3 Increasing Switch Speed

The third scheme is to operate the switch at a faster speed than the link speed so that more incoming packets can be processed in one packet slot time (pkt size/link speed) and the output contention problem is reduced. If the output contention problem is reduced, the input queueing delay is also reduced. In this scheme since the switch speed is greater than the downlink speed, to effectively improve the downlink beam utilization, buffering is required at the output ports to hold the packets. The output queue performs as a statistical multiplexer and the speed of the multiplexer is the same as the downlink speed.

It is possible to combine the non-FIFO queue scheme and the speedup scheme so that the trade-off among throughput, the delay,and the hardware cost can be optimized. Two configurations are simulated to show the effect of speedup: the first one has a speedup factor of 1.5 and checking depth of 1, and the second has a speedup factor of 1.25 and checking depth of 2. Note that the maximum achievable throughput for both configurations is around 0.9.

The cell delay jitter, the cell transfer delay, and the cell 90 percentile delay of the switch for both configurations are depicted in Figures D-25, D-26, and D-27. For single-cell transmission, when the utilization is less than 0.85, the delay performance for both configurations is about the same. When the utilization approaches 0.9, the first configuration performs better. This is because the packet transmission time for a speedup factor 1.5 is less the packet transmission time for a speedup factor 1.25. For SVP transmission, the delay performance are about the same for both configurations and all utilizations. Note, as previously mentioned, that there is a trade-off between the packet service time through the switch and SVP size to optimize the packet delay. Therefore, the combination scheme, using speedup and non-FIFO queue, is adequate for single-cell transmission and lower utilization, and it is appropriate for SVP transmission for all utilizations.

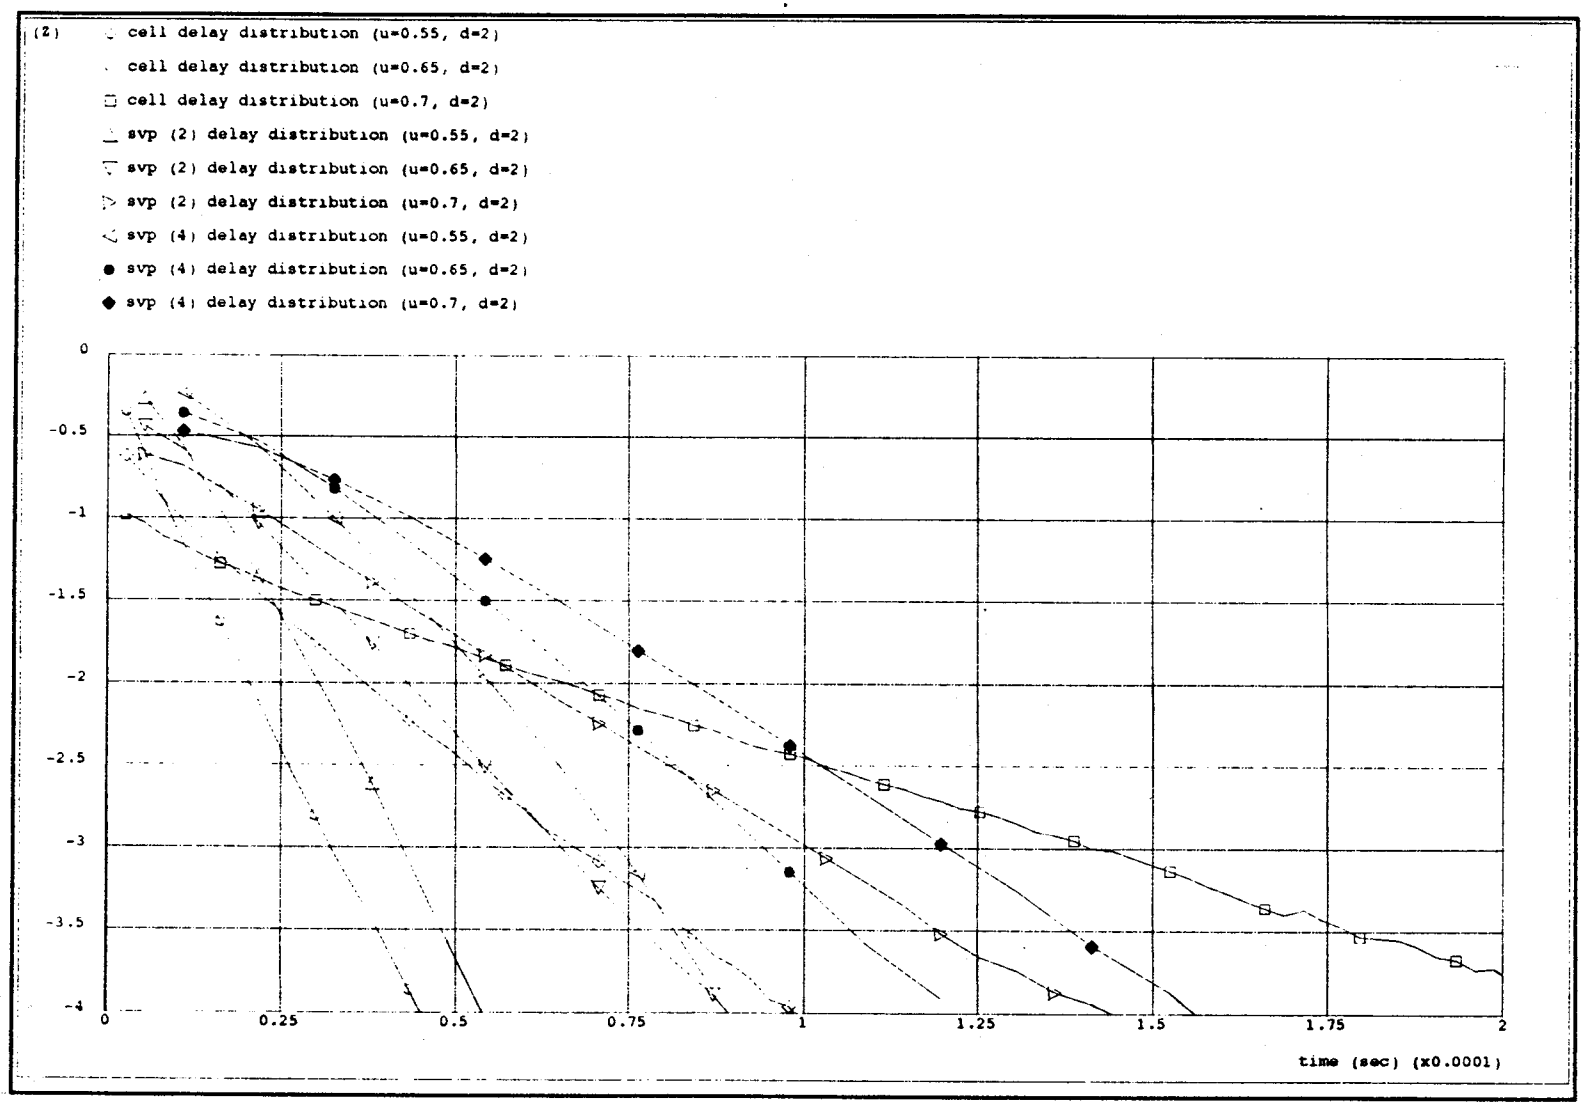

Figure D-23. Cell Delay Distribution for Different SVP Sizes and Link Utilization for Checking Depth of 2

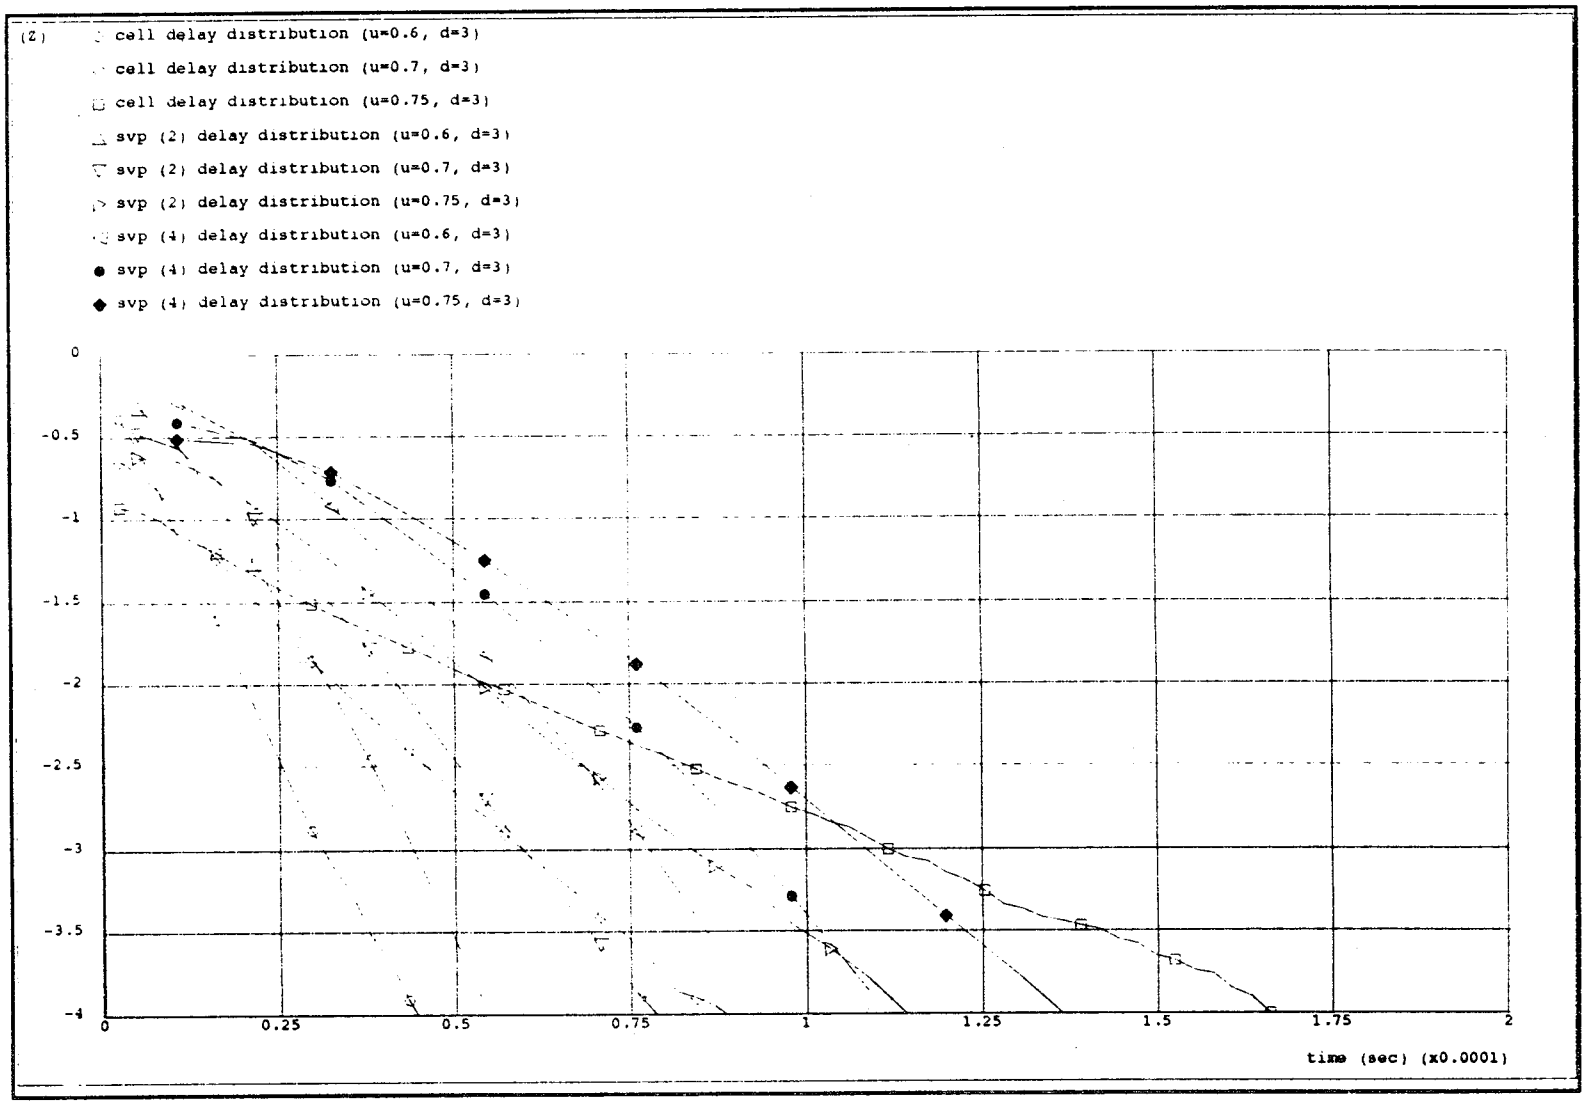

Figure D-24. Cell Delay Distribution for Different SVP Sizes and Link Utilization for Checking Depth of 3

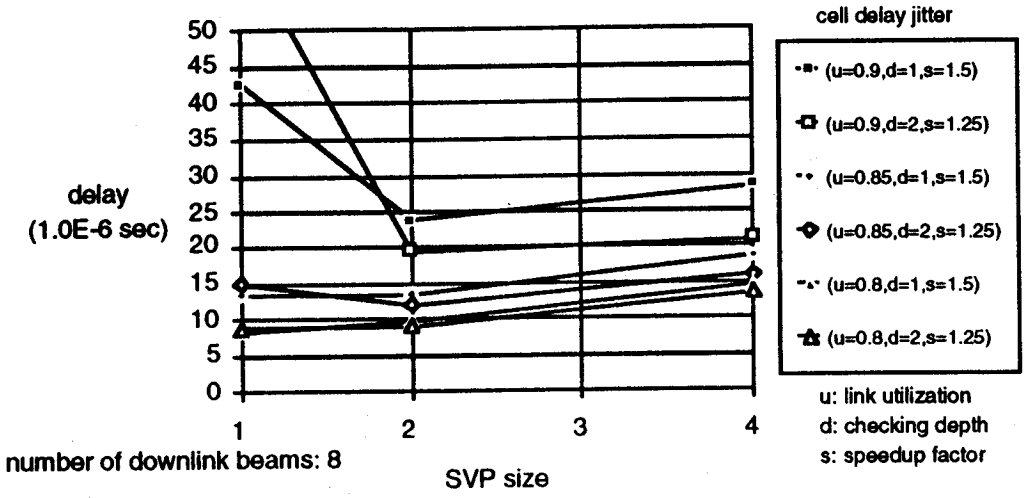

Figure D-25. Cell Delay Jitter vs. SVP Size for Different Link Utilization

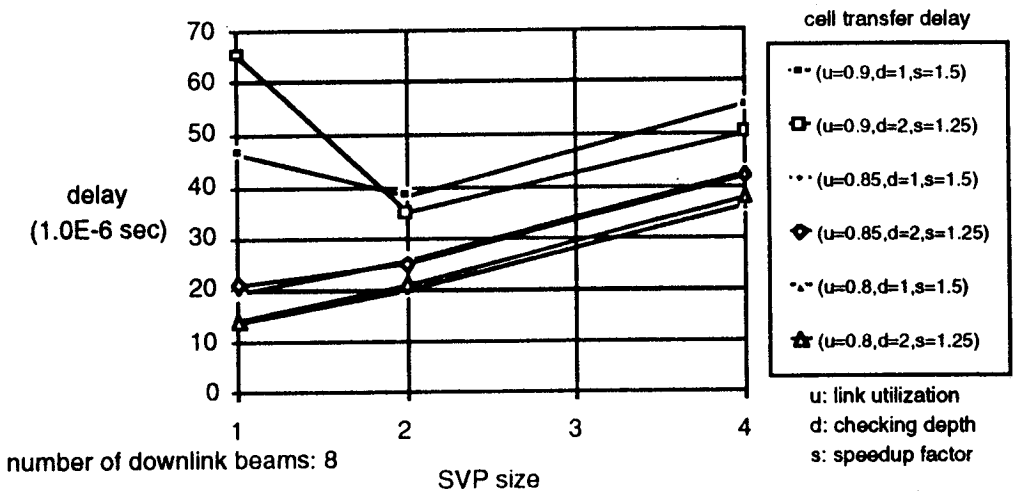

Figure D-26. Cell Transfer Delay vs. SVP Size for Different Link Utilization and Configurations

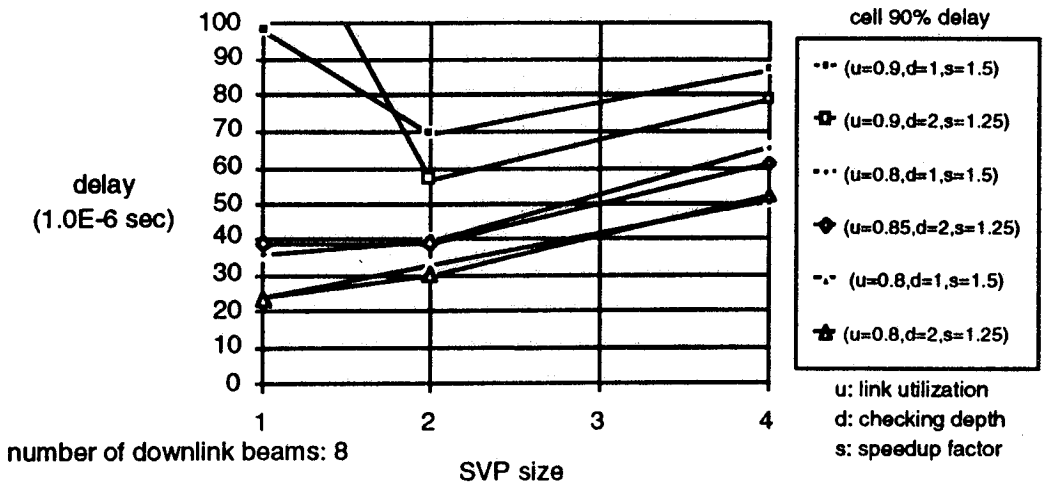

Figure D-27. Cell 90 Percent Delay vs. SVP Size for Different Link Utilization and Configurations

The delay distribution comparison for both configurations when the SVP sizes and link utilization are parameters are provided in Figures D-28, D-29, and D-30.

3.2.2 Output Buffering

In this approach, the packets are buffered at the output ports. To resolve the output blocking problem, either the switching fabric has to operate at a speed faster than the line speed, as in the case of the banyan-type network, or there must be a disjoint path between any input-output pair and the output port has a multiple-entry port buffer such as the knockout switch [D-5]. In this section, only the banyan-type switch is considered. If the switching speed is N times faster than the line speed (where N is the size of the switch), all the packets destined to the same output port during the same slot time can be buffered at the output port. However, this approach is only feasible if the link speed is low and the switch size is small. In conclusion, it is not feasible to use output buffering alone to increase the throughput of the banyan-type switch.

4 Concluding Remarks

In this section, the pros and cons of preprocessing ATM cells into SVPs at the earth station have been studied. Formatting cells into SVPs at the earth station can avoid on-board VPI/VCI processing and HEC processing, simplify the space segment complexity without introducing much hardware at the earth station, and bit interleaving of cell headers can be naturally achieved. Based on the performance analysis of SVP transmission through the earth station, to fully utilize the SVP concept without affecting the delay quality, the uplink and downlink has to operate at a very high

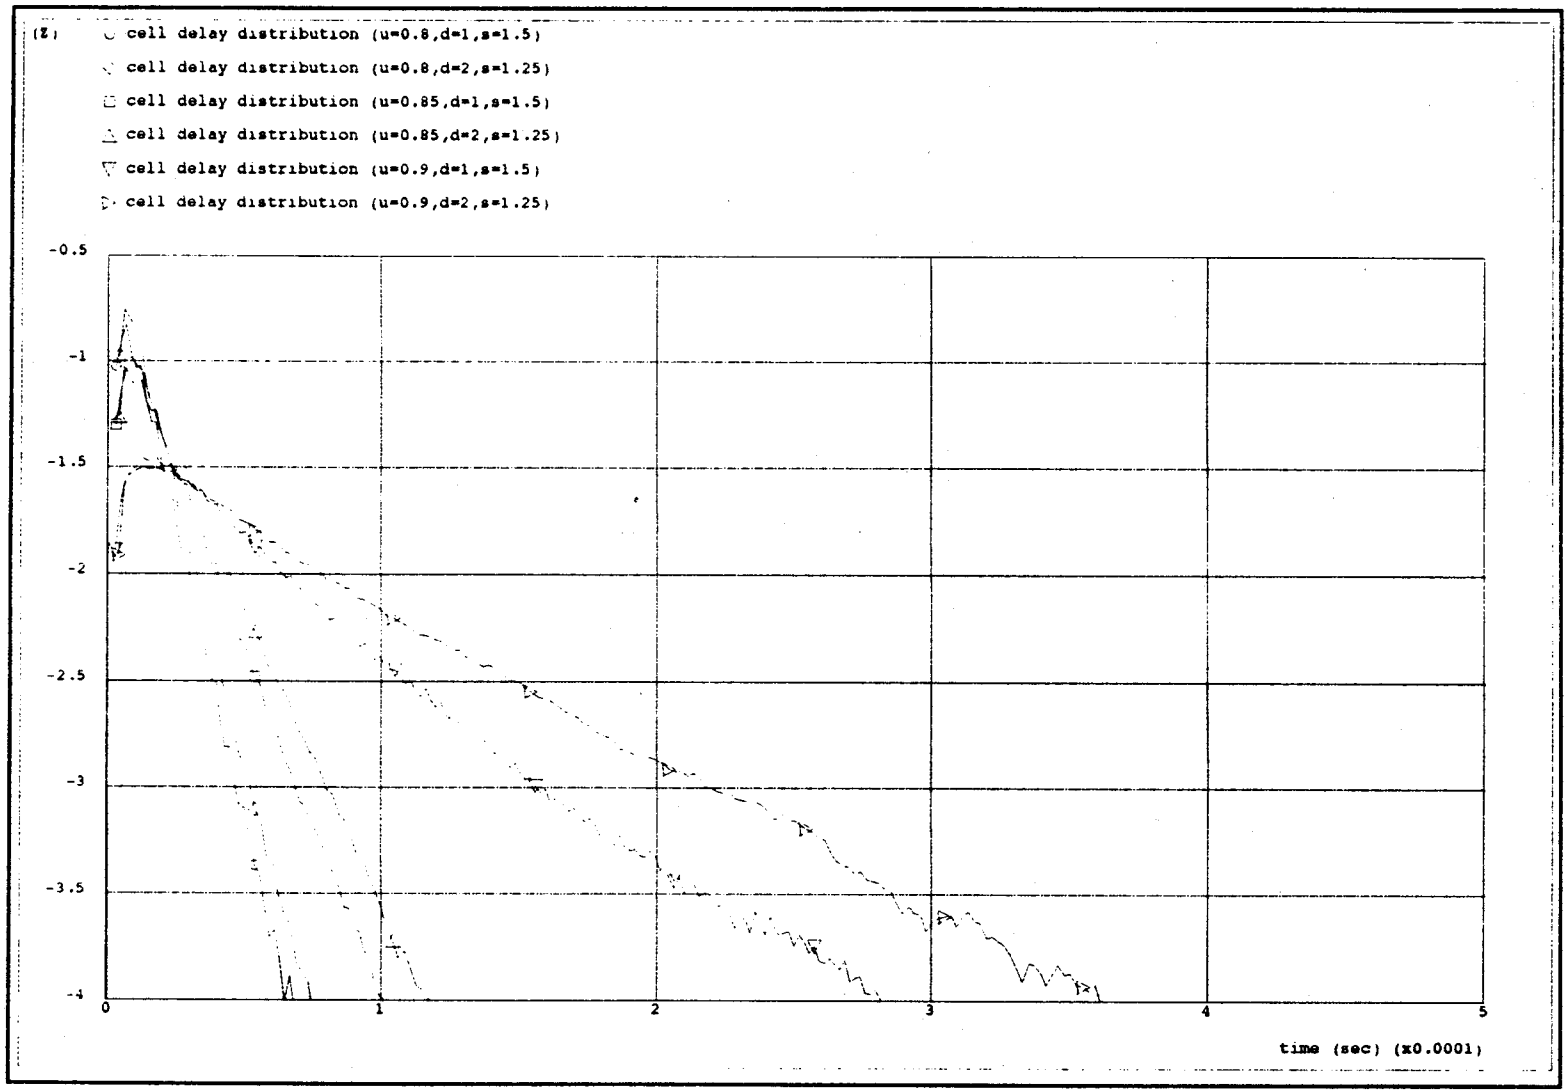

Figure D-28. Cell Delay Distribution for Different Link Utilization and Configurations

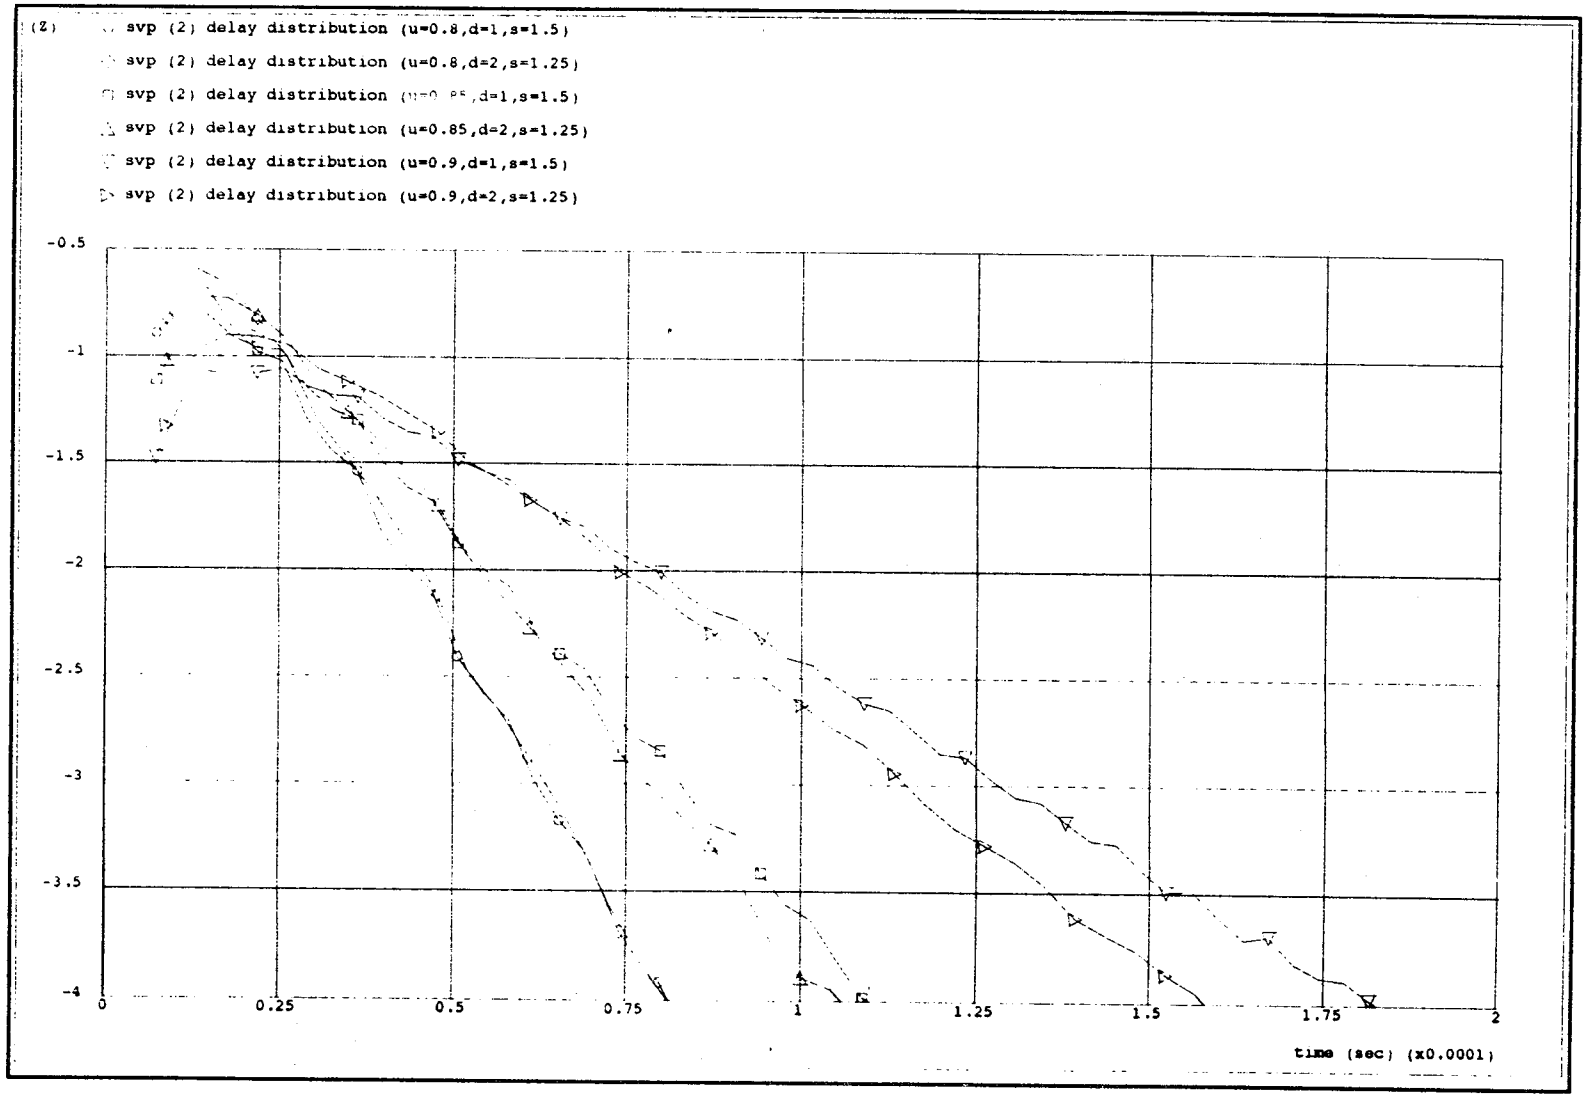

Figure D-29. Cell Delay Distribution for Different Link Utilization and Configurations for SVP Size of 2

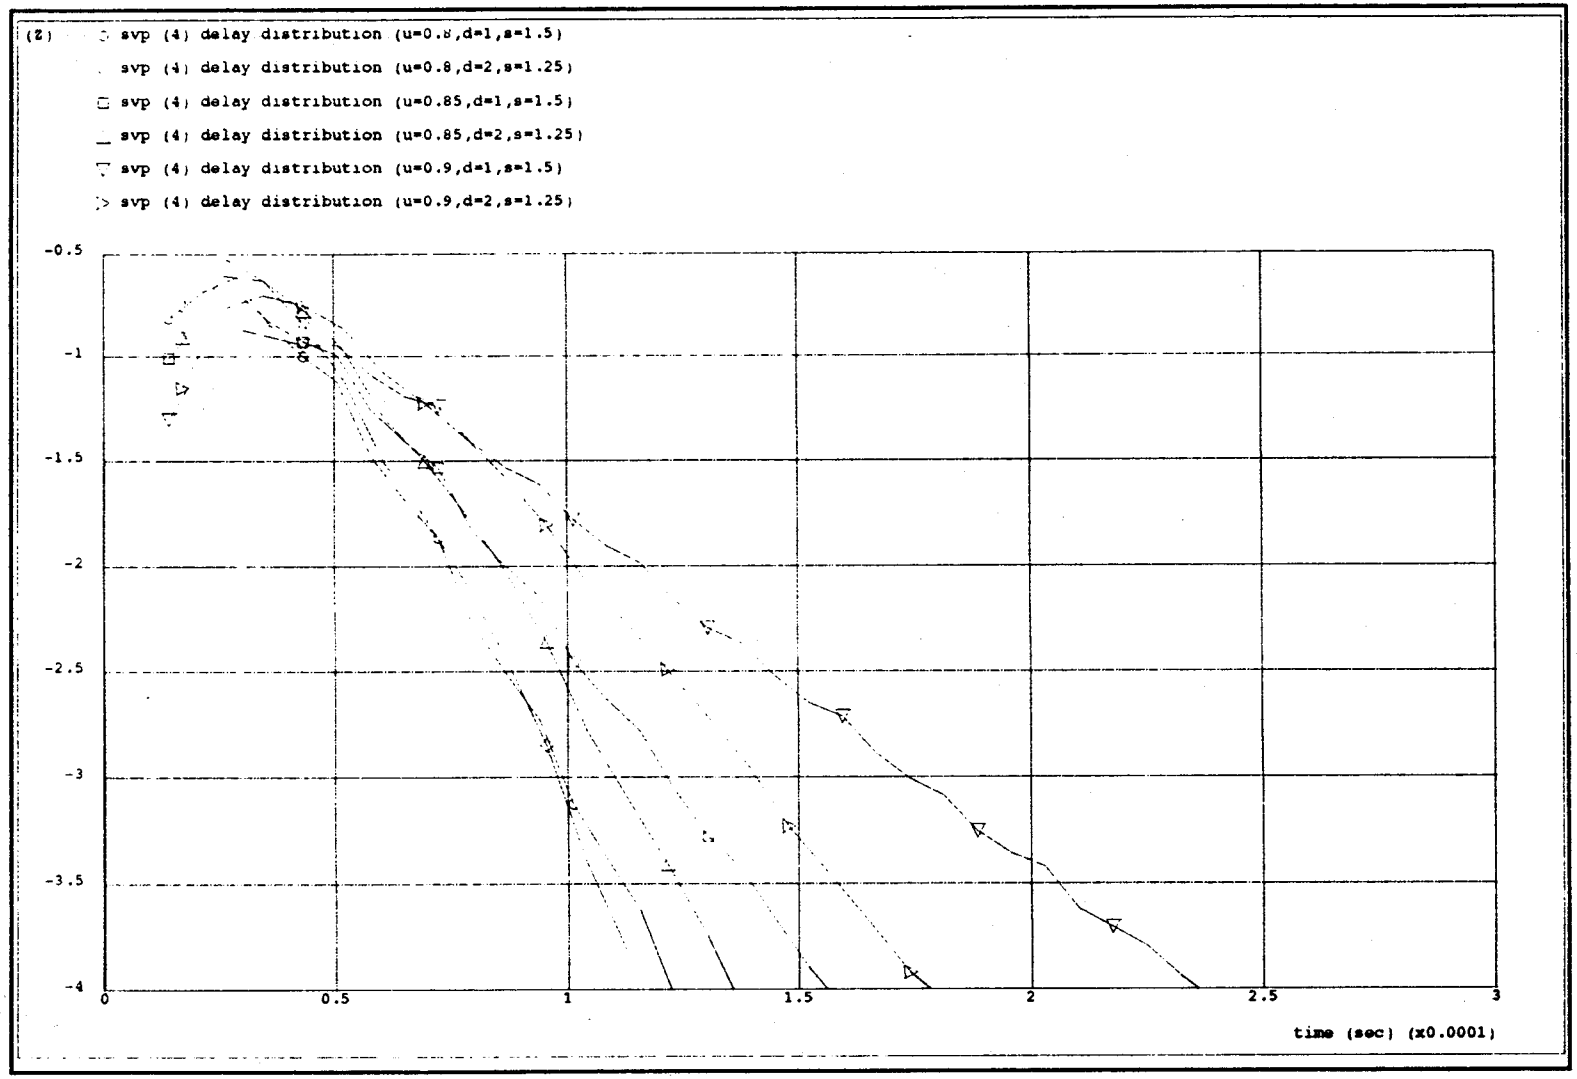

Figure D-30. Cell Delay Distribution for Different Link Utilization and Configurations for SVP Size of 4

utilization(above 80%) and the SVP size should be kept small (≤ 4 cells). Otherwise, a corrective measure to reduce the SVP packetization delay such as using a timer associated with each SVP should be used. It has been shown that the timer approach is more effective if the input link utilization is low. Nevertheless, it is always a good practice to use a timer for each SVP to prevent that a SVP may have a large packetization delay.

It is also found that the SVP formatting at the earth station has the effect of reducing the probability of continuous arrivals of the packets with the same destination at the input port. The result has shown to be effective in reducing the output contention problem of the on-board switch. This suggests that spacing is one of the necessary requirements for satellite B-ISDN congestion control. The spacing process is one of the traffic shaping functions used to ensure that cell streams coming into the network do not exceed the negotiated value between the subscriber and the network. The spacing mechanism does not send the packets with the same destinations back by back, since the peak value for this stream is the same as the satellite transmission link. Other packets with different destinations can be inserted between two packets with the same destination. The result of spacing is that output contention of the on-board switch is reduced.

To provide high uplink and downlink utilization within the satellite network, the throughput of the on-board fast packet packet must be greatly improved. Several schemes have been examined in this section. Increase buffer size is not effective for the satellite environment. The non-FIFO queue with a preset checking depth at the input port to resolve HOL blocking has shown to be effective. However, the throughput improvement is still very limited since it is not practical use a large checking depth. To significantly increase the throughput, the speedup scheme must be used. It has been found that the non-FIFO queue in conjunction with speedup scheme can significantly increase the throughput of the switch with a reasonable hardware cost. Therefore, the combination of input queueing, a non-FIFO queue with a preset checking depth, speedup, and output queueing optimizes the performance of the switch and the hardware cost.

This section has analyzed the single-cell and SVP performance through the satellite B-ISDN for point-to-point traffic. Extension of this work to point-to-multipoint traffic is left for future research.

5 References

[D-1] F. A. Tobagi, "Fast Packet Switch Architectures for Broad-band Integrated Services Digital Networks," Proceedings of the IEEE, vol. 78, no. 1, pp. 133-167, Jan. 1990.

[D-2] H. Ahmadi and W. E. Denzel, "A Survey of Modern High-Performance Switching Techniques," IEEE JSAC, vol. 7, no. 5, pp. 1091-1103, Sep., 1989.

[D-3] B. Bingham and H. Bussey, "Reservation-Based Contention Resolution Mechanism for Batcher-Banyan Packet Switches," Electronic Letters, vol. 24, no.13, pp. 772-773, June 1988.

[D-4] M. Karol, M. Hluchyj, and S. Morgan, "Input vs. Output Queueing on a Space-Division Packet Switch," IEEE Trans. on Communications, vol. 35, pp. 1347-1356, Dec. 1987.

[D-5] U. Yeh, M. Hluchyj, and A. Acampora, "The Knockout Switch: A Simple, Modular Architecture for High-Performance Packet Switching," IEEE JSAC, vol. 5, pp. 1274-1283, Oct. 1987.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	January 1993	Final Contractor Report

4. TITLE AND SUBTITLE

On-Board B-ISDN Fast Packet Switching Architectures
Phase 1: Study

5. FUNDING NUMBERS

WU-None
NAGW-4528

6. AUTHOR(S)

Faris Faris, Thomas Inukai, Fred Lee, Dilip Paul, and Dong-Jye Shyy

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

COMSAT Laboratories
22300 COMSAT Drive
Clarksburg, MD 20871

8. PERFORMING ORGANIZATION REPORT NUMBER

E-7561

9. SPONSORING/MONITORING AGENCY NAMES(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
Lewis Research Center
Cleveland, Ohio 44135-3191

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

NASA CR-189144

11. SUPPLEMENTARY NOTES

Project Manager, William D. Ivancic, Digital Systems Technology Branch, NASA Lewis Research Center, (216) 433-3494.

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified - Unlimited
Subject Category 17

12b. DISTRIBUTION CODE

13. ABSTRACT (*Maximum 200 words*)

Broadband ISDN (B-ISDN) extends the service integration concept to include future high speed, high volume network services as well as to accommodate narrowband ISDN in its infrastructure. The impact of B-ISDN services on satellite communications can be significant, requiring the same or similar quality of communications services provided by the terrestrial B-ISDN network. The purpose of this study is to investigate potential satellite applications, alternate satellite network architectures, system design issues, and onboard switching/processing options for satellite B-ISDN services. Special attention is focused on fast packet switching architectures and switch design details to exploit the feasibility of implementing a fast packet switch in the future advanced communications satellites.

14. SUBJECT TERMS

B-ISDN; Fast packet switch; Time division-multiplead; Asynchronics transfer mode (ATM)

15. NUMBER OF PAGES

504

16. PRICE CODE

A22

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	