N94-22778

DESIGN AND IMPLEMENTATION CONSIDERATIONS OF A MSAT PACKET DATA NETWORK

Fouad G. Karam, Telesat Mobile Inc.
1145 Hunt Club Road, Ottawa, Ontario, Canada, K1V 0Y3
613-736-6728/Fax613-736-4548
Terry Hearn, Westinghouse Electric Corp.
410-993-1346/Fax410-765-9745
Doug Rohr, Westinghouse Canada
Arthur F. Guibord, Telesat Mobile inc.

ABSTRACT

The Mobile Data System, which is intended to provide for packet switched data services is currently under development. The system is based on a star network topology consisting of a centralized Data Hub (DH) serving a large number of mobile terminals. Through the Data Hub, end-to-end connections can be established between terrestrial users on public or private data networks and mobile users. The MDS network will be capable of offering a variety of services some of which are based on the standard X.25 network interface protocol, and others optimized for short messages and broadcast messages. A description of these services and the trade-offs in the DH design are presented in this paper.

1. INTRODUCTION

The mobile data services which are provided by TMI consist of two types :The basic services which are based on the standard X.25 network interface protocol, and the Reliable Transaction and the Unacknowledged data delivery services. These will be described briefly in Section 2 while Section 3 deals with presenting the logical and functional architecture of the Data Hub. This includes the protocol processing and associated interfaces, and the MDS network management which handles configuration, fault, accounting, performance and statistics.

During the initial stages of the development, several design trade studies were undertaken to establish the architecture of the central Data Hub. Results from these studies are presented in Section 4. The advantages of implementing the Data Hub as a number of distributed processors rather than a single processor are outlined. This paper also discusses how the chosen architecture provides for flexibility in system growth while meeting the overall availability and performance requirements.

In order to meet performance and system availability requirements, decisions had to be made with respect to computing platforms. As it will be demonstrated, fault tolerant computers are deployed for network and system management, while the real-time processing functions are handled by high throughput multi-processing systems operating in a real-time fashion.

Several failover and failure recovery scenarios will be analyzed and described in Section 5 of this paper. As it will be seen, the system designers ensured that equipment failure shall not cause an overall network failure.

2. OVERVIEW OF MOBILE DATA SERVICES

The Mobile Data System (MDS) provides for packet data transmission to mobile users. High efficiency and cost effectiveness are achieved by a large number of mobile users dynamically sharing the space segment. Data services provided by MDS fall into two categories: Basic, and Specialized. The basic service category within the MDS provides for the establishment of end-to-end virtual circuits between DTEs attached to Mobile Terminals (MTs) and DTEs attached to the Data Hub (DH). The basic service category is composed of two distinct services: X.25 and asynchronous. The X.25 service is compliant with the 1988 version of the CCITT recommendation. The asynchronous service is based on CCITT recommendation X.3, X.28, and X.29. Figure 1 below shows the MDS basic services architecture.

2.1 Protocols and Their Characteristics

The motivation to support the Basic and the Specialized Services in a very efficient manner in order to minimize the utilization of the available L-band spectrum has led to the development of a number of protocols designed specifically for this purpose. The satellite protocols were designed to take advantage of the packet data transmission capability of the system. Whenever possible, MTs are assigned channel capacity by the Data Hub. This reduces the need for MTs to compete for capacity on a slotted Aloha random access channel and results in a more efficient use of spectrum. For instance, for any data transfer from the DH requiring an acknowledgement, the DH will allocate the necessary channel capacity, on a TDMA channel, to the MT in question. The system attempts to maximize the use of TDMA rather than slotted Aloha channels since these are much more efficient.

2.2 The MDS Satellite Protocols

A number of system requirements, that affected the choice and design of the satellite protocol architecture, are identified below; these are:

1. Support for a large number of MTs
2. Optimization of the satellite resources
3. Support for various types of user traffic varying from short messages to large file transfers
4. Flow and congestion control
5. The support of a priority scheme
6. Satellite protocol modularity and flexibility to allow for future growth.

2.3 Satellite Protocols Stack

The Basic and Specialized services are supported via internal MDS protocols operating between the DH and MTs. Figure 2 outlines the architecture for these protocols. The MDS Packet Layer Protocol (MPLP) provides procedures for the setup, maintenance and tear-down of virtual circuits between the DH and MTs. It is responsible for supporting the basic MDS services. The MDS Data Link Protocol (MDLP) provides for the reliable sequenced delivery of packets to the MPLP. Functionally, it is similar to LAP-B Multi-Link Procedure (MLP). The MDS Specialized Services Protocol (MSSP) provides for the multiplexing of application messages over the Reliable Transaction Service (RTS) and the Unacknowledged Data Service (UDS) supported by the MDS Transaction Protocol (MTP) and the MDS Unacknowledged Link Protocol (MULP) respectively. The MTP is used for transaction type data exchange, while the MULP provides for the transmission and reception of unacknowledged data packets to and from MTs.

The Channel Access and Control (CAC) defines a set of procedures for accessing the physical layer. The CAC is mainly responsible for allocating TDMA capacity as well as frame assembly and disassembly of data segments at the DH and the MT. The Bulletin Board (BB) provides for the dissemination of system information from the DH to all MTs. These are: channel definition, protocol parameters, and congestion avoidance indication.

Only the satellite protocols pertinent to user data transmission are discussed.

2.3.1 MDS Channel Access and Control

The MDS CAC specifies a set of procedures to access various MDS channels. The channel types supported in MDS are: (1)the outbound DH-D, TDM channel with fixed frame size, each frame contains variable sized CAC segments, (2) Inbound MT-DT TDMA channel that carries variable sized bursts, each carrying variable sized segments on a reservation basis, (3) Inbound Slotted Aloha random access MT-DRr channel that carries fixed sized bursts, each containing fixed sized TDMA request segments, and (4) Inbound Slotted Aloha random access MT-Drd channel that carries fixed sized bursts, each containing one variable sized small data segment.

The CAC is highly efficient in use of the satellite capacity. It implements an eight-level priority scheme, a congestion control algorithm, as well as a load balancing algorithm over the outbound channels.

2.3.2 MDS Packet Layer Protocol (MPLP)

The MPLP functions as the network layer of the seven layer OSI stack. It is modelled after the ISO 8208 standard. The key features of MPLP are summarized below:

- MPLP supports the X.25 and Asynchronous data services. It also supports various X.25 features like the D, Q, and M bits, the negotiation of flow control parameters, the Fast Select and MDS User Identifier (MUI). All other X.25 optional facilities which are not acted upon by MPLP are conveyed transparently through the network for further treatment at the user side.

- MPLP supports the priority selection on an individual circuit basis. The MPLP user is able to signal one of eight priority level at call setup using the throughput class facility. MPLP conveys this information to the lower layers which will guaranty the corresponding priority of access for the duration of the call.

- MPLP is optimized over a satellite channel by modifying the use of layer 3 Receiver Ready (RR). A specific procedure to deal with layer 3 RR generation is implemented in MPLP.

- MPLP provides for the authentication of every switched virtual circuit that is established between an MT and the DH; the MUI is used for this purpose.

- The MPLP does not support the Restart procedure. However, if a restart request is received from the user, MPLP clears all virtual circuits and resets all permanent connections.

2.3.3 MDS Data Link Protocol (MDLP)

MDLP is a highly efficient link layer protocol which optimizes the use of the satellite resources. The features of MDLP are summarized below:

- The MDLP, in contrast with LAP-B does not support layer 2 RR or RNR as a flow control mechanism. The protocol uses six Protocol Data Units (PDUs).

- MDLP does not require a link layer acknowledgement as in LAP-B. Acknowledgements are withheld until the window is closed or a link layer timer expiry occurs. An MDLP task at the DH or MT can request a selective repeat from its counterpart by forwarding, in a STAT_PDU frame, a bit map of the frames received and thus minimizing the activity over the spacelink.

- In order to satisfy the requirements of the CAC layer, MDLP supports the fragmentation and reassembly of MPLP packets. MDLP uses the "More" bit for this purpose.

- The maximum window size supported by MDLP is 15, allowing up to 15 outstanding frames to be unacknowledged.

- The MDLP at the DH makes use of the MDS TDMA capability. When the DH MDLP is expecting a link layer acknowledgement from the MT MDLP, the latter is explicitly solicited using the Poll bit. In such a case, TDMA request is made on behalf of the MT MDLP and the TDMA allocation is piggybacked on the MDLP frame.

2.3.4 MDS Transaction Protocol (MTP)

The MTP is optimized to support transaction type applications, and a single segment unacknowledged messaging capability. That is, when the data to be exchanged takes the form of a command-response, then MTP is the choice. The MTP is ideal for applications whose message data size does not exceed the maximum CAC segment size (max. 64 bytes).

The MTP is a highly efficient satellite protocol. The MTP user has the ability to specify the response length, the response delay, the number of message repeats to increase the probability of success over a noisy channel, and the delay between repeats.

2.3.5 MDS Unacknowledged Link Protocol (MULP)

MULP is primarily designed for use in applications such as broadcast or multicast of news, weather reports, and financial market information. The major features of the MULP are outlined below:

- MULP supports multiple simultaneous applications.
- MULP supports the fragmentation and reassembly of user messages which can be significantly large (up to 64 CAC data segments).
- Data integrity is guaranteed by MULP. Corrupted user messages are deleted and discarded.
- A MULP user is capable of specifying the number of times, his message is to be repeated. This parameter is highly important, to overcome link errors and guaranty message delivery to the destination.

3. DH LOGICAL ARCHITECTURE

The Data Hub logical architecture is presented in Figure 3. It consists of four functional sub-systems:

- MDS Network Management Sub-system (MNMS)
- Satellite Network Access Controller Sub-system (SNACS)
- Data Channel Unit Sub-system (DCUS)
- Terrestrial Interface Sub-system (TIS)

The MNMS consists of an MNMS controller which provides the management functions of the system namely: MT, configuration, fault, performance, security, and accounting. Databases are stored and maintained locally by the MNMS. Also, a Man-Machine Interface (MMI) in the MNMS provides operator command and display facilities for control and monitoring of MDS operation.

The SNACS consists of a number of Satellite Protocol Processors (SPP) and a number of Network Access Processors (NAP-D). The SPP supports all the satellite protocols except MDLP, while the NAP performs the CAC and MDLP functions. It also supports the BB function used for the dissemination of system information to the MTs.

The DCUS consists of satellite interface channel equipment for MDS data channels. The following functions are provided: encoding, interleaving, scrambling and modulation of DH-D frames, and demodulation, descrambling, de-interleaving, and decoding of received IF signal into the inbound frames.

The TIS provides interface lines to public and private data networks. It also generates call records for basic services and forward them to the MNMS.

4.0 DH HARDWARE ARCHITECTURE

A number of trade studies were undertaken to define and develop a DH architecture. The studies examined trade-offs so the target architecture meets the timing and sizing, reliability, availability, maintainability, and expansibility system requirements.

In order to evaluate different types of processors to perform the SNACS functions, a model to estimate CPU utilization was created. The model included system threads which represent basic and specialized protocol transactions between the MT and the PDN. These threads were then linked into MDS protocol transactions and a total CPU utilization was estimated.

The following architectures were evaluated as possible candidates for the MDS. (1)fault tolerant switching processors, (2)packet switches with open architecture, (3)VAX clusters, and (4) VMEBUS SPP and NAP-D. In this section the architecture of choice is presented. All others were rejected since they did not meet the acceptance criteria.

Figure 4 below shows the DH hardware architecture. The MNMS is based on a fault tolerant computing platform, while the TIS from Northern Telecom DPN-100 packet switch provides redundancy at various levels. Each of the SPP and NAP-D consists of two redundant VME chassis housing a number of Single Board Computers (SBC), which are based on the Motorola MC68040 microprocessor. The NAP-D and SPP communicate via a dual rail Ethernet LANs. The MNMS uses the same LANs to interface with the SNACS. The LAN protocol is TCP/IP. One LAN is designated for inbound traffic (MT to DH) and the other for outbound traffic (DH to MT). The TIS interfaces with the SNACS via a number of high speed X.25 V.35 circuits at 256 kbps each. The NAP-D communicates with the DCUS by means of redundant RS485/RS530 multidrop links running at 800 kbps.

5.0 SYSTEM FAILURE RECOVERY

The design goal for redundancy is to allow for a single component failure and to recover from the failure in a minimum time period by using the redundant component. The MNMS implements a redundancy manager which collects health information from DH components. This information is processed by a rules based knowledge system in order to choose the appropriate action for a given set of detected faults or errors. When the rules have indicated a failed component is required to be switched out of the system and a healthy standby component is available, the following actions are taken:

- the failed component is commanded to go OFF-LINE
- the standby component is updated with configuration data if required
- the standby component is commanded to go ON-LINE
- diagnostics are performed on the failed component
- equipment failures are replaced with Line Replaceable Units (LRUs)
- the failed component is placed in a STANDBY mode of operation.

The NAP-D cages will send health information to the MNMS on a periodic basis. Each NAP-D cage pair is configured in a warm standby mode of operation. Failure to report status to the MNMS, or receiving an unhealthy cage status, will cause a switchover to the standby NAP-D cage. This switchover will not cause active virtual circuits to clear.

The MDS channel units are full duplex units connected to the NAP-D by redundant multi-drop links. Each channel unit may be in the on-line, stand-by or off-line state as determined by the MNMS and the health of each unit. Standby DCUs have their operational code and are waiting for configuration and state transition data.

All SPP cages are in the on-line state. They "usually" carry user traffic on a full time basis. A processor card failure will cause all calls to be re-routed to another operational SPP card by virtue of the loadsharing feature implemented in the SPP software. The SPPs are sized to operate at 50% of their ultimate full load to allow for redundancy.

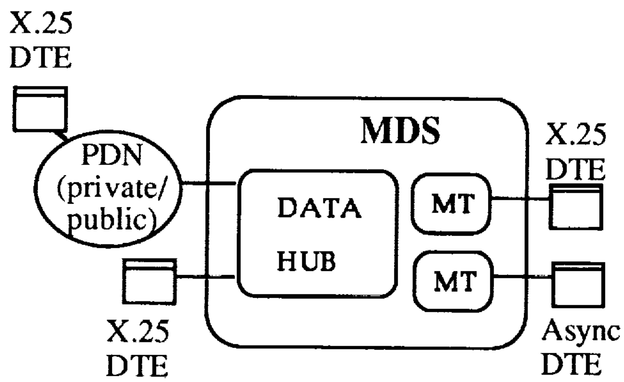

Figure 1: MDS Basic Services

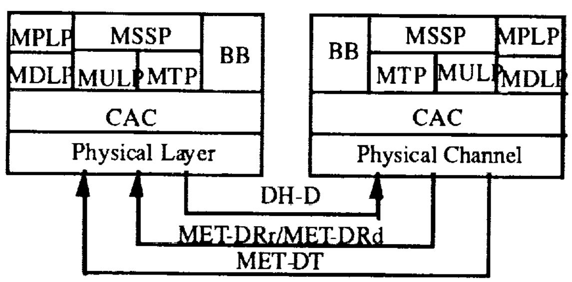

Figure 2: MDS Protocols Stack

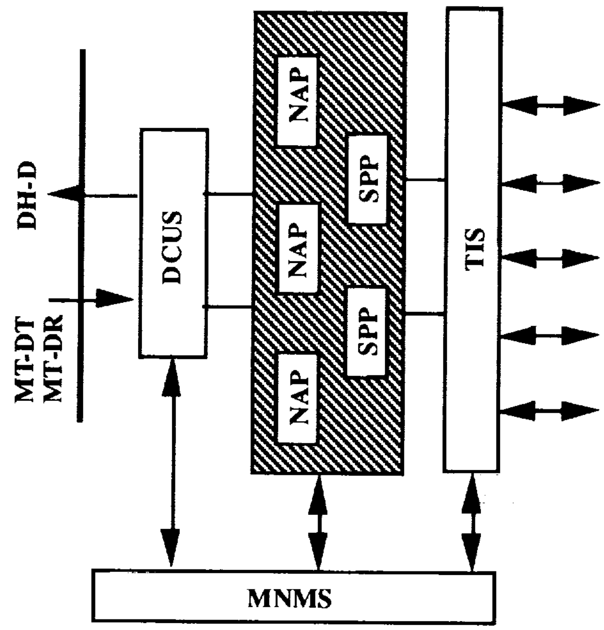

Figure 3: Data Hub Logical Architecture

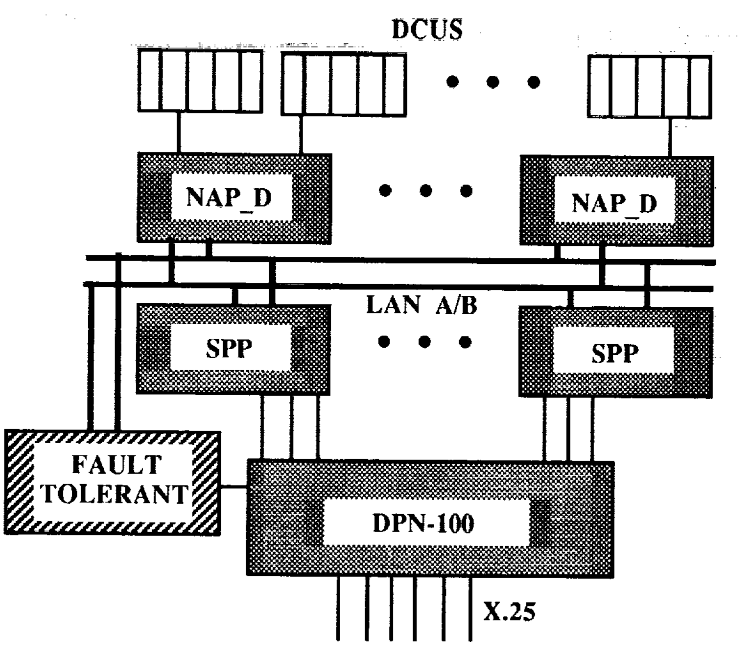

Figure 4: DH Block Diagram