

56/CW/IN/32

1999 NASA/ASEE SUMMER FACULTY FELLOWSHIP PROGRAM

JOHN F. KENNEDY SPACE CENTER  
UNIVERSITY OF CENTRAL FLORIDA

*Performance Evaluation of the NASA/KSC Transmission System*

Kenneth J. Christensen  
Assistant Professor  
Department of Computer Science and Engineering  
University of South Florida

KSC Colleagues: John Schnitzius (NASA) and Chris Kerios (Dynacs Engineering)

**ABSTRACT**

NASA-KSC currently uses three bridged 100-Mbps FDDI segments as its backbone for data traffic. The FDDI Transmission System (FTXS) connects the KSC industrial area, KSC launch complex 39 area, and the Cape Canaveral Air Force Station. The report presents a performance modeling study of the FTXS and the proposed ATM Transmission System (ATXS). The focus of the study is on performance of MPEG video transmission on these networks. Commercial modeling tools - the CACI Predictor and Comnet tools - were used. In addition, custom software tools were developed to characterize conversation pairs in Sniffer trace (capture) files to use as input to these tools. A baseline study of both non-launch and launch day data traffic on the FTXS is presented. MPEG-1 and MPEG-2 video traffic was characterized and the shaping of it evaluated. It is shown that the characteristics of a video stream has a direct effect on its performance in a network. It is also shown that shaping of video streams is necessary to prevent overflow losses and resulting poor video quality. The developed models can be used to predict when the existing FTXS will "run out of room" and for optimizing the parameters of ATM links used for transmission of MPEG video. Future work with these models can provide useful input and validation to set-top box projects within the Advanced Networks Development group in NASA-KSC Development Engineering.

# Performance Evaluation of the NASA-KSC Transmission System

Kenneth J. Christensen

## 1. Introduction

NASA-KSC currently uses three bridged 100-Mbps FDDI segments as its main transmission system, or backbone, for data traffic. The FDDI Transmission System (FTXS) connects the KSC industrial area, KSC launch complex 39 and the Cape Canaveral Air Force Station. NASA-KSC is in the middle of developing a consolidated networking infrastructure based on Asynchronous Transfer Mode (ATM) as an upgrade to the FTXS. This new transmission system is the ATM Transmission System, or ATXS. The KSC Development Engineering (DE) directorate is accountable for this project which is described in [6]. Within the DE directorate, the Advanced Networks Development (AND) group is examining new networking technologies and developing the design specification. The goal of the consolidated networking infrastructure is to reduce the number of disparate networks at KSC and thus achieve cost savings by 1) increased efficiency in the use of the existing in-ground cabling plant, and 2) reduced operations and maintenance requirements. Current transmission of shuttle launch video is via the KSC WideBand and BCDS analog (cable TV) networks. A goal of the shuttle upgrade program is to generate MPEG-2 encoded video and transmit this video via the ATXS.

A key component in the planning, design, and deployment of any large network is performance evaluation, or capacity planning. The goal of capacity planning is to deliver the most cost-effective network designs while ensuring that performance levels will meet specified service requirements (see [8]). While very valuable, building experimental test networks in a laboratory setting is expensive and limited in size and scope. Due to its size and complexity, the entire KSC transmission system cannot be duplicated in a laboratory. Thus, a performance model of the existing and proposed systems needs to be developed and experiments conducted on the model. The results from the experiments can give guidance on the amount of "head room" available on the existing FTXS and the anticipated performance of possible ATM upgrade designs.

During the summer of 1999 (funded as a second-year NASA/ASEE faculty fellowship) a capacity planning study of FTXS and ATXS was conducted. The CACI suite of modeling tools - Predictor and Comnet (see [1]) - was used for performance evaluation. The purchase of the CACI tools was the direct result of a requirements definition and procurement completed in the summer of 1998 (see [2]). The goals of this capacity planning study were:

1. Baseline the topology and data traffic of the existing FTXS
2. Build models of the FTXS and proposed ATXS including video transmission subsystems
3. Collect or acquire representative MPEG video traffic traces and characterize them
4. Conduct experiments on the FTXS and ATXS models to determine:
  - Expected lifetime (or "headroom") of the systems given anticipated traffic growth rates
  - Predicted performance of video traffic
5. Evaluate shaping of video traffic as a means of improving the efficiency of use of network resources

The key deliverables from this study were:

1. CACI Predictor models for the FTXS and ATXS
2. Tools to characterize and import traffic into the Predictor models
3. CACI Comnet models for video transmission on FDDI and ATM networks
4. Characterization of some packetized MPEG-1 and MPEG-2 video clips

These deliverables forward the goal of consolidating video services on existing and future networks at NASA-KSC.

The remainder of this report is organized as follows. Section 2 is a brief overview and motivation for capacity planning. Section 3 is an overview of the FTXS and ATXS network designs. Section 4 describes characterization of MPEG video streams. Section 5 describes the CACI Predictor models and results for FTXS and ATXS based on a "normal day" baseline traffic. Section 6 presents a traffic characterization from the July 22nd scrubbed shuttle launch day. Section 7 describes the CACI Comnet models for video transmission on ATM networks and presents the results of a study of shaping of video traffic. Finally, section 8 is a summary and describes future work.

## 2. Overview and Motivation for Capacity Planning

Capacity planning of computer systems is the process of predicting the future performance of a system. Specifically, capacity planning intends to predict when performance will drop below a desired level of service. For any moderately large or complex system it is not possible to physically build and test design options, hence modeling is required. Capacity planning uses both simulation and analytical modeling methods. Analytical methods use "math formulas" to model and solve for the performance measures of a system. Analytical methods are rooted in queueing theory and require simplifying assumptions to build tractable models. Reference [5] is the classic reference on queueing theory. Simulation methods use discrete event simulation programs executed on a computer to find statistical approximations for the performance measures of a system. Reference [7] is the classic reference on discrete event simulation. Simulation models require much longer execution time than analytical models. However, simulation models do not require many of the simplifying assumptions of analytical models. Another trade-off is correctness, given that the underlying assumptions are accepted, an analytical model can be formally proved correct. However, a simulation model consists of thousands to millions of lines of code and is hence almost certain to contain "bugs", some of which may affect the accuracy of the output. The accuracy of output (i.e., the results) for any model is no better than the accuracy of the input. Given first order approximations for input, the output of both an analytical and simulation model will be poor. This last fact is very important, if many assumptions must be made on the input to a model then faster analytical modeling may be better than simulation modeling since one can study many design options with an analytical model in the same amount of time that only one option can be studied with simulation methods.

Computer networks are modeled as queueing systems. The buffers in network components such as routers and switches are packet (or cell for ATM) queues. Delay in a network occurs when packets are forced to wait for transmission in these buffers. Packet loss occurs when these buffers overflow due to packet arrival rates temporarily exceeding the service rate. Packets arrive with probabilistically distributed interarrival times, often in bursts. A packet is serviced when it is transmitted on an outgoing link. Packet service time is a function of the packet size and the data rate of the outgoing link. For example, a 1500 byte packet has a 1.2 millisecond service time on a 10-Mbps Ethernet. Both packet interarrival times and packet sizes can be modeled with probability distributions. Packet streams are modeled and analyzed as time series and their characteristics determined by time series analysis.

Burstiness can be informally defined as how much the arriving packets are clumped together. A more formal definition of the burstiness of a packet arrival stream is the Coefficient of Variation (CoV) of the packet interarrival times (IAT),

$$\text{CoV of packet IAT} = \frac{\text{Standard deviation of packet IAT}}{\text{Mean of packet IAT}}, \quad (1)$$

where the mean of packet IAT is simply the reciprocal of the packet arrival rate. The utilization of a network is the ratio of measured arrival rate divided by the maximum rate ("bandwidth") of the link. Determining the variance of packet interarrival times requires analyzing packet traces of network traffic. Packet traces contain time-stamps for received packets. It is from these time-stamps that the variance of packet interarrival times is determined. For deterministic traffic the CoV = 0, for random (Poisson) traffic the CoV = 1, and an increasing burstiness is indicated with a higher value of the CoV. As the arrival process increases in burstiness, the saturation knee occurs at lower utilization levels. The knee effect is common to all queueing systems. The challenge of capacity planning is to determine at what utilization level the knee will occur.

## 3. Overview of the FTXS and ATXS Designs

The current data communications infrastructure of NASA-KSC is the FTXS (see [6]). The underlying technology of the FTXS is FDDI, a 100-Mbps shared-medium token ring technology. FDDI is an ANSI standard (see [3]) developed for reliable and efficient data communications based on variable length packets. In an FDDI LAN, all of the attached hosts share the common bandwidth of a 100-Mbps fiber optic ring. In FTXS, three FDDI rings are interconnected via bridges. Organizations, often identified by a KSC building, attach to the FTXS via a router. Figure 1 shows a graphic of the FTXS taken from a Predictor model (the model is described in Section 4). The ATXS (see [6]) is a complementary network to the FTXS with a plan to consolidate some, or all, of the FTXS traffic. The ATXS is based on ATM technology with standards largely set by the ATM Forum (see [2]). The

ATXS is especially intended to carry video traffic that is currently carried on the separate analog BCDS and WideBand systems. The purpose of this study is to evaluate the ATXS for this video traffic consolidation. Figure 2 shows the ATXS network (as a Predictor graphic).

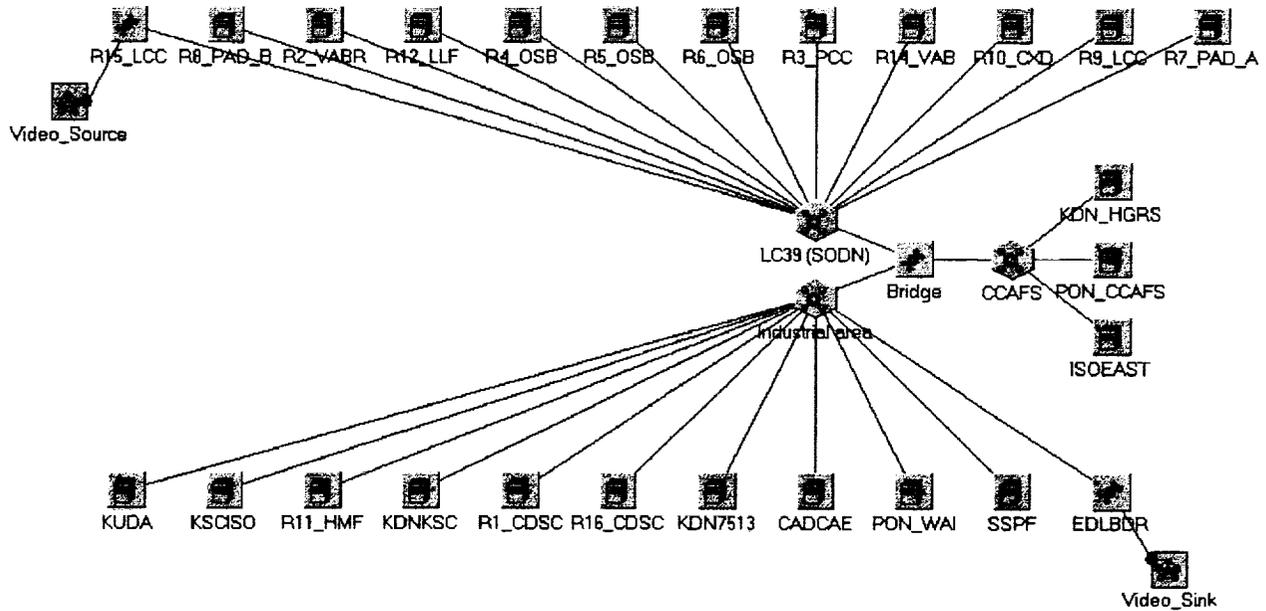


Figure 1 - Predictor graphic of the FTXS

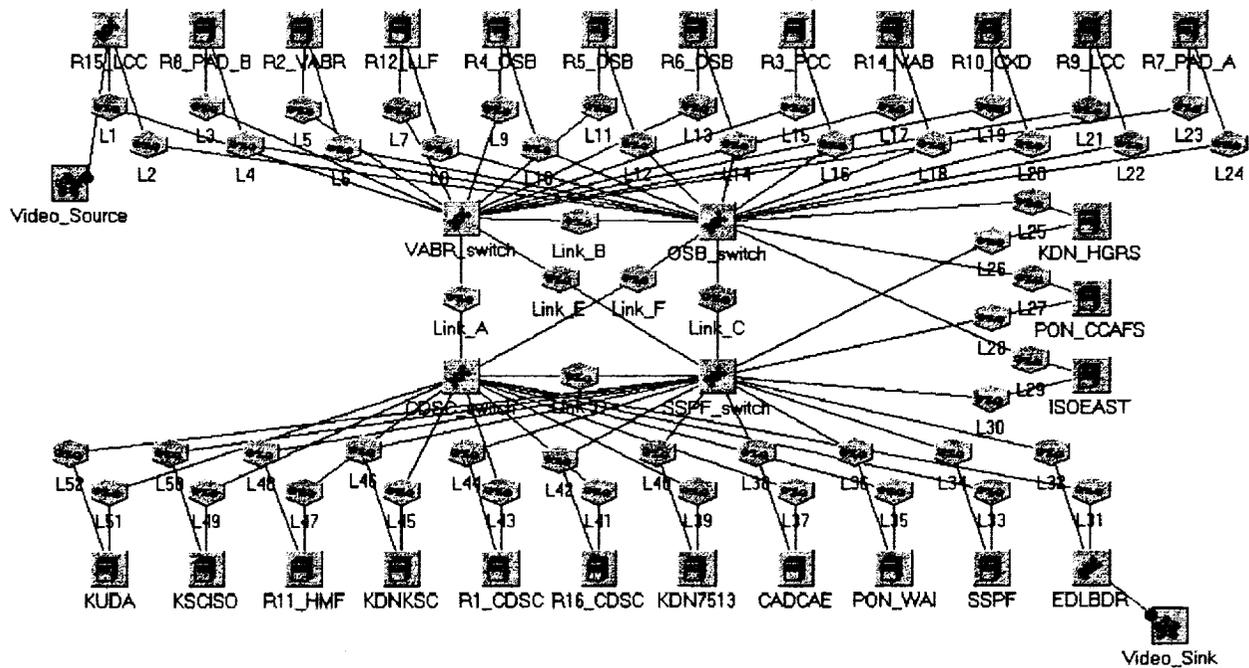


Figure 2 - Predictor graphic of the ATXS

#### 4. Characterization of MPEG Video

MPEG-1 and MPEG-2 are video encoding standards used for digital video, see [9]. A video encoder generates a stream of compressed frames that can be stored as digital data (e.g., as a file on a hard disk) or sent across a network for video transmission. A matching decoder uncompresses the encoded frames and generates a "video out" signal as NTSC or composite video. The encoding and decoding incurs a delay - the amount of delay depends on the implementation of the encoding and decoding hardware and/or software. Additional delay may be incurred in a playout buffer used to compensate for delay variance (called "delay jitter") caused by the variable queuing delays in the network. For video transmission on a network, MPEG frames are packetized into one, or more, packets per frame. MPEG defines three types of frames, I (interpolated), P (predictive), and B (bi-directional). An I frame is an independently compressed video frame, P and B frames are inter-frame compressed in reference to the previous I frame (and to each other). I frames are typically generated every half-second (with 14 P and B frames between each I frame for 30 frames per second video). The loss of an I frame prevents the subsequent B and P frames from decoding correctly and typically results in a half-second "freeze" in the video. The loss of a B or P frame has less effect on the video quality. Since I frames are independently compressed, they are typically larger in size than B and P frames making their loss more likely in a congested network. In addition, for most video encoders the size of the frames is highly dependent on the complexity and amount of motion in the scene on which the camera is focused. For example, a talking-head in front of a solid color background will encode into very small frames (and hence low mean and variance in bandwidth required to transmit the frames on a network), an automobile race with frequent scene changes will result in the opposite. For most encoders, the compression rate is a settable parameter with a higher compression ratio resulting in a lower transmission bandwidth requirement (or less storage), but lower video quality. For MPEG-2 encoding, an average rate of about 6-Mbps results in VCR quality video. Constant bit-rate, variable video quality encoding is also possible where the output rate of the encoder is constant, but video quality varies as a function of the amount scene complexity and motion. Figure 3 shows a video source, encoder, network, and decoder (and assumes one video frame per packet). Note that video frames entering the network are deterministically spaced (typically one every 1/30-th of a second for 30 frames per second video) and that this same timing must exist at the decoder if the decoder is not to be over-run or under-run. An end-to-end network delay of a few milliseconds per packet is considered to be acceptable to most video applications.

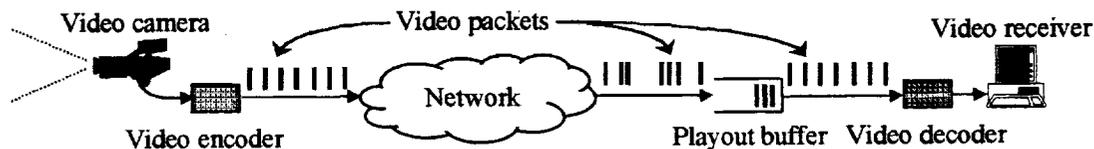


Figure 3 - Digital video on a network - encoding, playout buffering for jitter removal, and decoding

Understanding the characteristics of digital video is necessary in order to predict the network resources required to transmit the video. Significant research has been done in the area of video characterization and traffic modeling. The results in [11] show that the characteristics of a video stream are highly dependent on the content of the video. At the time of this work, there were no MPEG-2 encoders available at KSC. Thus, to study the time series characteristics of MPEG-2 video streams, traces of MPEG-2 encoded 1996 Olympic Games events were obtained (see [13]). Figures 4, 5, and 6 show the bandwidth requirements for 30 seconds of MPEG-2 encoded Olympic Games waterpolo. Increasing the time increment aggregates the time series. At a 30 second increment, the mean bandwidth requirement is 5.94-Mbps. At the smaller time increments the large spikes (one every half-second) correspond to the I-frames, with a peak bandwidth of over 15-Mbps. At the larger time increments, the spikes are not visible, this can result in under-allocation of network bandwidth for video transmission.

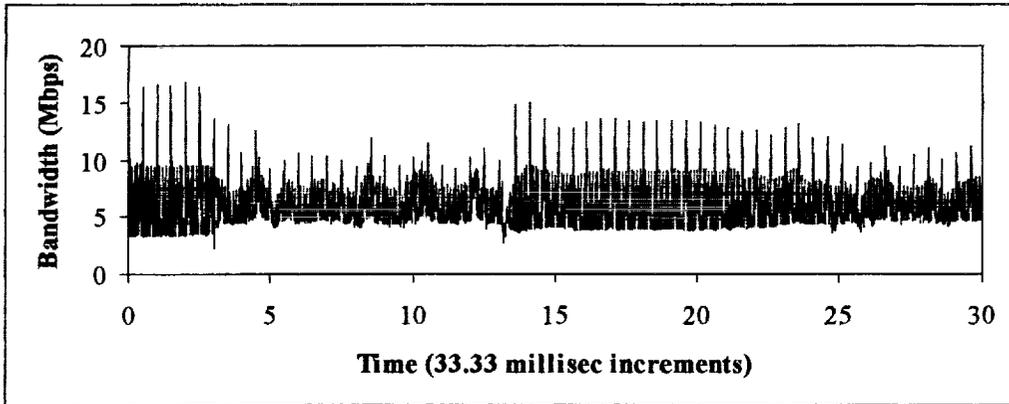


Figure 4 - Bandwidth for 30 seconds of MPEG-2 Olympic games waterpolo (33.33 millisecond increment)

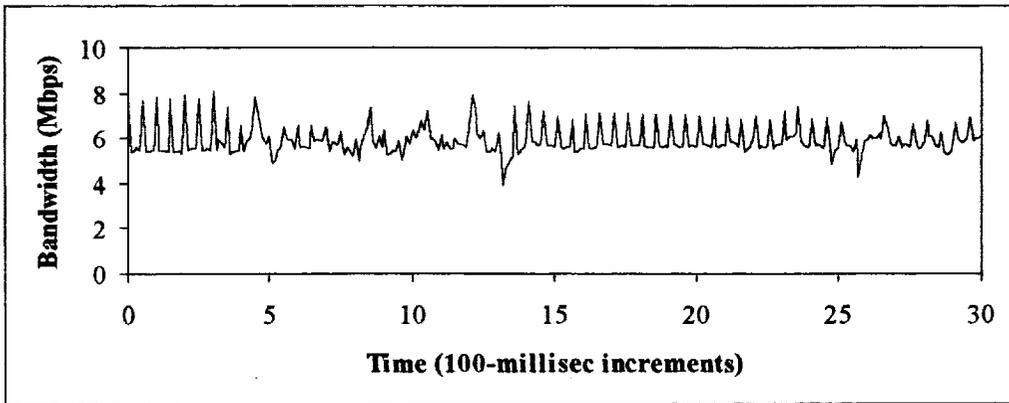


Figure 5 - Bandwidth for 30 seconds of MPEG-2 Olympic games waterpolo (100 millisecond increment)

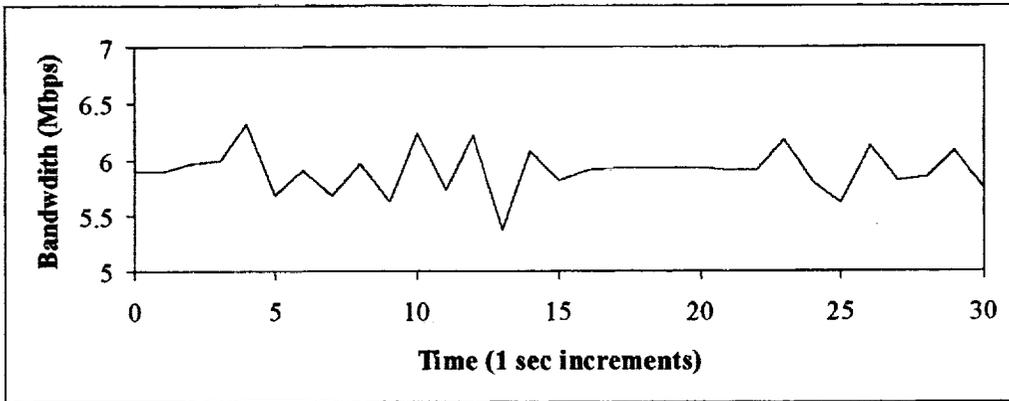


Figure 6 - Bandwidth for 30 seconds of MPEG-2 Olympic games waterpolo (1 second increment)

Table 1 summarizes two key characteristics of the set of Olympic Games MPEG-2 frame size traces. The mean bandwidth required for a video stream is directly proportional to the mean frame size. The burstiness of the video stream is directly proportional to the Coefficient of Variation (CoV) of frame size, which is the ratio of the standard deviation of frame size to the mean of frame size (similar to equation (1) for CoV of packet IAT). It can

be seen in Table 1 that the mean frame size varies very little between video clips (ranging from 24200 to 24552 bytes, a 1.5% difference), however the CoV varies by 41% (ranging from 0.4858 to 0.6861).

Sporting event	Full video clip (40 minutes)		First 10 seconds of video clip	
	Mean frame	CoV of frame	Mean frame	CoV of frame
Archery	24498 bytes	0.6816	24203 bytes	0.6584
Badminton	24200	0.6861	24252	0.6398
Boxing	24504	0.5592	24500	0.6005
Diving	24438	0.5748	24457	0.6651
Sailing	24552	0.4858	24612	0.5145
Swimming	24498	0.6551	24565	0.5855
Table tennis	24363	0.6409	24522	0.5329
Waterpolo	24507	0.5058	24558	0.4617

Table 1 - Key characteristics of Olympic Games MPEG-2 frame traces

## 5. Predictor Modeling of the FTXS and ATXS

A CACI Predictor model can give a first order measure of the capability of the existing FTXS and the future ATXS for transporting video streams. A layer-2 model of the FTXS was manually built as shown in Figure 1. Each of the routers attaching to the FTXS is modeled as a traffic source. A video source is modeled in the LCC building and a video receiver in the EDL building. The FTXS is modified to model the ATXS by replacing the FDDI core with a switched core. The Predictor models compare a shared medium (i.e., that of the shared 100-Mbps FDDI) with a dedicated medium performance of a switched infrastructure. The more advanced functions of ATM including virtual circuits, traffic classes, and traffic shaping cannot be modeled with Predictor.

### 5.1 Traffic collection and import into the Predictor model

Predictor classifies network traffic as packet flows where each flow has a mean and standard deviation of packet size and mean and Squared Coefficient of Variation (SCV) of packets per second. The SCV is defined for a flow of packets in terms of the variance and mean of packet interarrival time (IAT) as:

$$\text{SCV of packet IAT} = [\text{CoV of packet IAT}]^2 = \frac{\text{Variance of packet IAT}}{[\text{Mean of packet IAT}]^2} \quad (2)$$

The SCV of IAT is a measure of the burstiness (or "clumping" tendency) of a packet flow, as described in Section 2. A deterministic packet flow (i.e., packets always have a constant IAT) would have SCV = 0. A Poisson packet flow (i.e., packets have exponentially distributed IAT) would have an SCV = 1. Using Network Associates Sniffer (see [10]) trace files as input, a program named `crunch.c` was written to identify the conversation pairs (between layer-2 MAC addresses) and characterize each conversation for mean and standard deviation of packet size and mean and SCV of packet IAT. The output of the program is in a format suitable for import directly into a Predictor tabulated packet flow. Reference [4] is a NASA technical disclosure filed for this software program. The program has two possible output modes - verbose and CSV (for input to CACI Predictor as a tabulated packet flow). Predictor will import an Expert Sniffer file as traffic input. However it was found via experimentation (and confirmed by CACI technical support) that the Expert Sniffer is "behind" the actual capture. In one experiment, the FDDI Expert Sniffer identified about 10-Mbps of traffic, however the capture file showed approximately 26-Mbps.

The process of collecting and characterizing traffic over a period of many weeks is called "baselining". A baseline is the starting point for a network model. Growth rates, additional traffic sources, changes in network topology, and so on are then made to the baseline to predict future performance of the network. The `crunch.c` program applied to Sniffer capture files can characterize small snapshots of traffic on a LAN. For a longer term view, RMON probes can be used to record the layer-2 mean packet size and packets per second for all conversations in a network. The use of RMON for long term monitoring of shared FDDI bandwidth usage needs to be further investigated. For a fully switched network (e.g., switched Ethernet or ATM), RMON cannot be used.

For a router-centric view of traffic, Cisco NetFlow provides session level information (a session is a TCP connection) for all packet flows through a monitored router. The use of NetFlow traffic collection for input to Predictor models deserves further study.

Over a period of several days in mid June 1999 it was found that the industrial area FDDI ring of the FTXS ranged from about 15% to 25% utilization. The Network Associates FDDI Sniffer monitor (with a 100-millisecond sampling time) was used for these measurements. A single capture of "typical" non-launch day traffic on the industrial area FDDI was made at an approximate 25% utilization level and characterized. The results from the characterization are shown in Figure 7 as `crunch.c` verbose output where all conversation pairs with less than 1% relative utilization have been removed for purposes of clarity. It can be seen that 21 conversation pairs generate a relative 81% of the traffic for this snapshot. One MAC address could not be identified (labeled as "unknown") and is suspected to be a RMON probe or some other monitoring device on the industrial area FDDI ring.

```
-----
Media data rate           =      100 Mbps
Total time of trace      =      9.3434 sec
Total number of pkts / bytes = 93526 pkts / 30429145 bytes
Total (absolute) util    =      26.0541 %
Total number of pairs    =      355 pairs
-----
```

to/from Addr	count	util (a)	util (r)	m_pkt_len	sd_pkt_len	pkt/sec	scv
R9_LCC KSCISO	1600	0.55	2.09	398.25	584.34	171.24	0.97
ISOEAST KDN7513	2597	0.67	2.59	303.04	397.66	277.95	0.61
R16_CDSC R14_VAB	532	0.41	1.58	903.36	662.26	56.94	22.07
R4_OSB KSCISO	548	0.30	1.14	631.59	475.12	58.65	1.38
R14_VAB KSCISO	1201	0.93	3.57	905.22	656.27	128.54	0.88
R15_LCC KDN7513	509	0.40	1.54	917.71	630.73	54.48	19.59
KDNKSC KDN7513	1535	0.33	1.27	251.43	250.29	164.29	1.86
SODN_R6 KSCISO	1675	0.60	2.30	418.47	440.45	179.27	1.14
KDN7513 ISOEAST	2528	0.35	1.35	163.00	182.58	270.57	0.52
KDN7513 R14_VAB	503	0.41	1.59	962.92	620.32	53.83	23.24
KDN7513 KDNKSC	1529	0.48	1.86	369.73	331.12	163.65	1.95
KDN7513 KSCISO	1465	1.05	4.04	838.38	619.88	156.80	1.74
SSPF KSCISO	1071	0.52	1.98	562.07	499.73	114.63	1.36
SSPF PONWAI	3756	2.94	11.27	913.12	646.42	402.00	2.40
unknown PONWAI	411	0.45	1.72	1270.64	433.16	43.99	4.24
KSCISO R9_LCC	2423	1.39	5.35	672.16	566.75	259.33	2.37
KSCISO PONWAI	1966	1.12	4.30	666.07	623.48	210.42	2.78
PONWAI SSPF	3514	1.25	4.80	415.91	569.45	376.09	2.05
PONWAI KSCISO	2714	0.93	3.57	399.89	443.89	290.47	1.83
PONWAI CADCAE	4981	5.78	22.19	1355.68	416.60	533.10	5.56
CADCAE PONWAI	3002	0.30	1.17	118.21	154.34	321.30	2.94

Figure 7 - Characterization of the industrial area FDDI ring FTXS traffic (6/15/99 - baseline 26% utilization)

### 5.2 Experiments and results for FTXS and ATXS Predictor models

The FTXS and ATXS Predictor models were used to predict the level of performance that could be expected for four video MPEG-2 video streams transmitted from the LCC to the EDL. The MPEG-2 waterpolo video clip frame size trace data was packetized into Ethernet (1500 byte maximum size) packets where frames of size larger than 1500 bytes are packetized into multiple back-to-back maximum size packets. The time series characteristics were found to be mean packet size = 1455 bytes, standard deviation of packet size = 208 bytes, packet rate = 506 pkts/sec, and SCV of packet IAT = 13.99 using the tools available at [12]. The data traffic characterized in Figure 7 was imported as the base data traffic representing a high 26% utilization of the FTXS industrial area FDDI ring. Figure 8 shows the mean and 90th percentile video packet delay for one through four video streams and with both no background data traffic and the baseline 26% utilization traffic. Figure 8 shows video packet delay for a growth in baseline data traffic. Figure 9 shows that video packet delay is fairly constant through a 150% growth, of about 65% utilization of data traffic of the 100-Mbps FDDI ring (the four video streams consume about 24% utilization for a total FDDI ring utilization of about 90%). This result shows that FDDI is "stable" for very high (up to 90%)

utilization levels after which packet delay increases sharply. The ATXS model predicts sub-millisecond mean and 90th percentile packet delay for growth in excess of 1000% (an 11 fold increase) in data traffic from the 26% baseline. This result shows that dedicated bandwidth can result in very large aggregate bandwidths for carrying many-to-many traffic. Buffer loss cannot be predicted with Predictor and is studied using Comnet in Section 7.

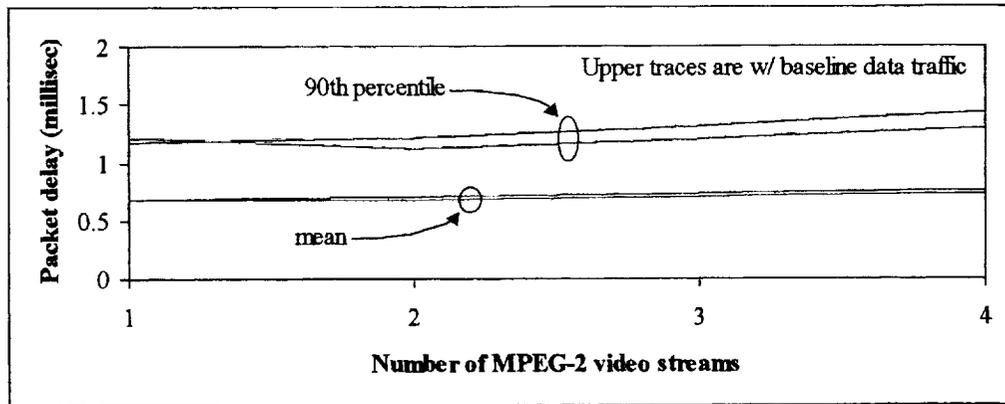


Figure 8 - Video packet delay for FTXS with one through four video streams

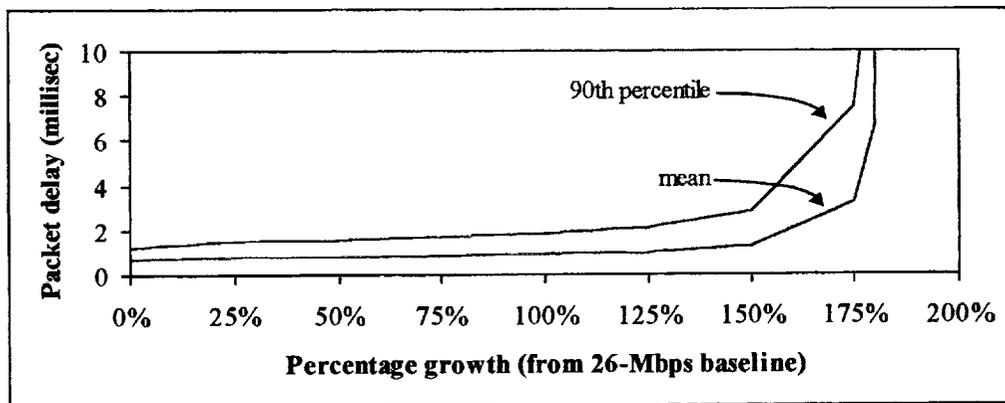


Figure 9 - Video packet delay for FTXS with four video streams and growth in data traffic

## 6. Launch Day Traffic Characterization

The KSC Web server is attached to the FTXS. On launch days the number of queries to this server is greater than on non-launch days. It is important that the internal NASA-KSC network perform "well" on launch days. Traffic data was collected for the 3 hours preceding the scrubbed July 22nd shuttle launch and characterized using the `crunch.c` program. Traffic data was also collected during the day of the scrubbed July 20th launch. An unusually high peak of almost 50% utilization of the FTXS FDDI ring was observed on 11:00am on July 20. Figure 10 shows characterization (`crunch.c` verbose output) of this capture where all conversation pairs with less than 1% relative utilization have been removed for purposes of clarity. It can be seen that 21 conversation pairs generate a relative 82% of the traffic for this snapshot. Figure 11 shows utilization during the night of the scrubbed July 22nd launch. It can be seen that traffic coming from and to the KSCISO router comprises the majority of the traffic on the FDDI ring. Traffic passing through the KSCISO router is also monitored with MRTG. The Sniffer data consists of approximately 10 second snapshots (10 seconds results in a full buffer on the Sniffer tool) taken every 10 minutes. Table 2 shows a breakdown of the protocols passing through the KSCISO router based on analysis (again, with `crunch.c`) of a 1:05am Sniffer trace. The large component of UDP traffic is believed to be from a new (first time used for this launch) streaming video application. It can be observed that

the burstiness of this streaming video is higher than that of the TCP traffic, most of which is presumed to be "normal" HTTP Web traffic.

```

-----
Media data rate           =      100 Mbps
Total time of trace      =      7.2778 sec
Total number of pkts / bytes = 122334 pkts / 45045982 bytes
Total (absolute) util    =      49.5158 %
Total number of pairs    =      322 pairs
-----

```

to/from Addr	count	util (a)	util (r)	m pkt len	sd pkt len	pkt/sec	scv
ISOEAST KDN7513	1241	1.30	2.63	955.95	659.15	170.52	5.25
R16_CDSC SODN R6	1634	1.61	3.26	897.77	628.06	224.52	4.56
R15_LCC KSCISO	1355	0.98	1.97	656.45	616.54	186.18	1.57
SODN_R6 KSCISO	4045	2.18	4.40	490.48	518.27	555.80	2.21
KDN7513 KDNKSC	657	0.60	1.20	824.75	608.81	90.27	2.27
KDN7513 KSCISO	1402	0.90	1.81	581.44	535.36	192.64	1.98
SSPF KSCISO	1587	0.85	1.72	488.87	478.29	218.06	2.21
SSPF PONWAI	3578	3.74	7.56	951.32	695.01	491.63	6.34
KSCISO R5_OSB	625	0.81	1.63	1174.82	522.85	85.88	4.81
KSCISO R2_VABR	974	0.98	1.98	916.42	694.10	133.83	1.72
KSCISO R15_LCC	1468	1.37	2.76	846.38	709.44	201.71	1.88
KSCISO KDNKSC	966	0.96	1.95	908.41	571.51	132.73	2.78
KSCISO SODN_R6	3547	1.46	2.94	373.52	515.45	487.37	2.22
KSCISO PONWAI	4287	3.10	6.26	657.66	610.23	589.05	3.53
PONWAI KDNKSC	664	0.61	1.22	829.10	679.84	91.24	18.81
PONWAI KDN7513	906	0.75	1.51	748.64	554.55	124.49	2.30
PONWAI SSPF	3593	3.48	7.04	882.14	711.52	493.69	4.73
PONWAI KSCISO	4291	1.58	3.18	334.21	405.07	589.60	2.62
CADCAE EDL_BDR	2734	4.43	8.94	1472.42	219.03	375.66	7.20
CADCAE KDNKSC	2807	3.91	7.89	1266.50	509.70	385.69	5.94
CADCAE <new>	2956	4.85	9.80	1493.04	192.20	406.16	6.27

Figure 10 - Characterization of the industrial area FDDI ring FTXS traffic (7/20/99 - launch day 49% utilization)

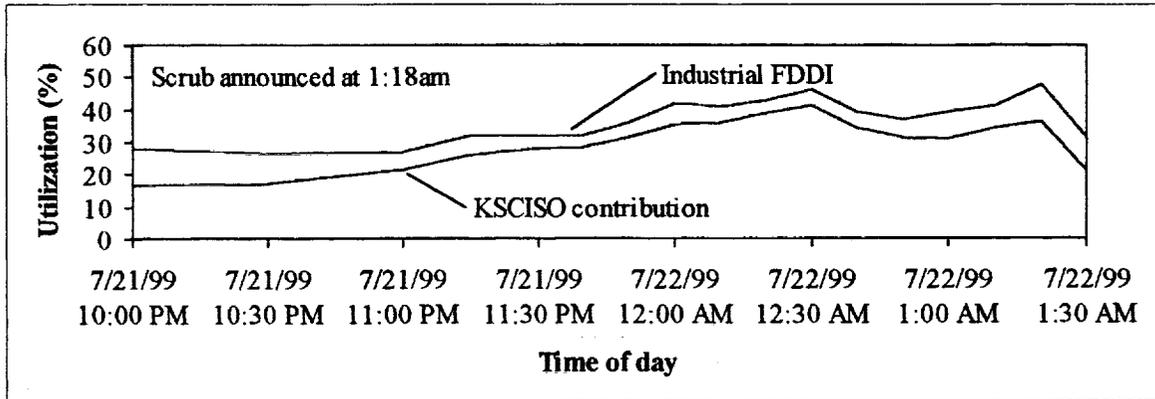


Figure 11 - Industrial FDDI ring utilization for night of July 22nd scrubbed launch (data from crunch.c)

Component	FDDI utilization	Weighted average of SCV
All traffic on FDDI	37.9%	10.19
KSCISO only	26.0	11.39
TCP to FTXS	1.0	3.47
TCP from FTXS	6.4	6.23
UDP to FTXS	2.0	5.71
UDP from FTXS	13.6	21.62

Table 2 - Breakdown of protocols at 1:05am from the July 22nd scrubbed launch (data from crunch.c)

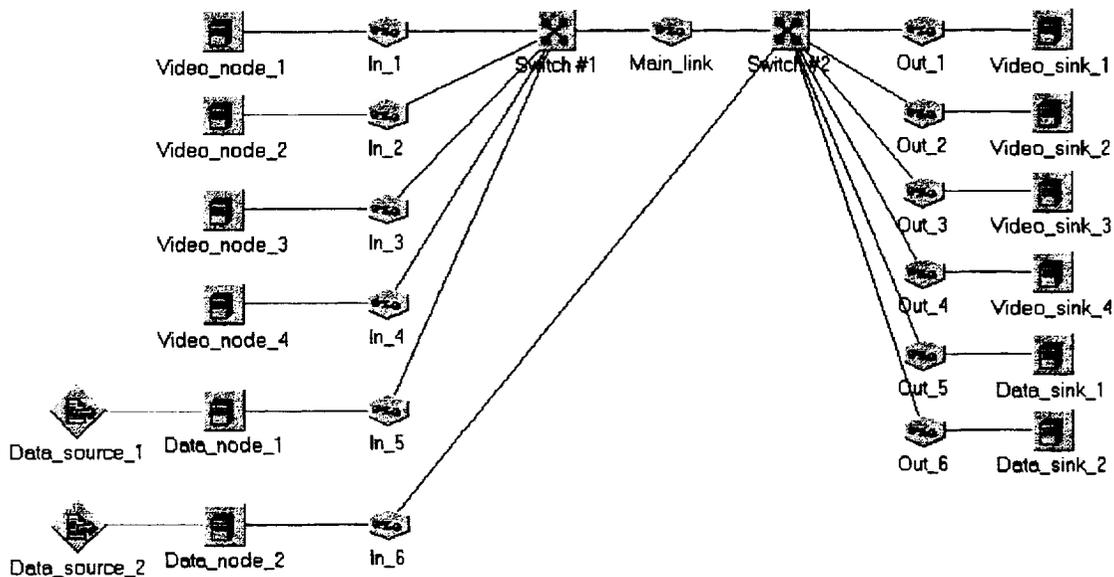
## 7. Comnet Modeling of Video Transmission on ATM

A CACI Comnet model can give a detailed, or “low level”, view of network behavior for a small snapshot of time. This low level view includes an understanding of packet or ATM cell loss due to buffer overflows in switches and routers. A Comnet model was built to study the delay and loss performance of video traffic in a simple ATM network. In particular, the use of shaping of video traffic is to be studied as a means of improving both video and data traffic performance. Figure 12 shows the Comnet model developed to study video traffic performance. Since no MPEG-2 encoded shuttle launch video was available, Olympic Games MPEG-2 video traces were used. Two experiments were defined:

**Experiment #1** - Four streams of video with high CoV (archery, badminton, boxing, and diving)

**Experiment #2** - Four streams of video with low CoV (sailing, swimming, table tennis, and waterpolo)

The video sources are Video\_node\_1, etc. and the video sinks are Video\_sink\_1, etc. The link speeds are all OC-3 and the switch buffer size is set to 512 cells. Shaping (and SAR to/from ATM cells) of the video traffic at the source is done using the Comnet external messages feature. Shaping rates of 6-Mbps through 40-Mbps are used.



**Figure 12** - CACI Comnet model for studying video traffic performance in an ATM network

Figure 13 shows the frame delay of MPEG-2 video for **Experiment #1** as a function of shaping rate. Results are only shown for the case where all frames are successfully delivered. A shaping rate of 6-Mbps is “too slow” and shaping at a rate higher than about 35-Mbps results in cell loss, and thus frame loss, at the switch. Figure 14 shows switch buffer use and its sharp increase at the higher rates of shaping. At a shaping rate of 40-Mbps, very few MPEG-2 frames are successfully delivered. This is not surprising since four times 40-Mbps exceeds the OC-3 line rate. Figure 15 shows frame delay for **Experiment #2**. It can be seen that the lower CoV video streams result in a significantly lower maximum frame delay and slightly lower mean and standard deviation of frame delay.

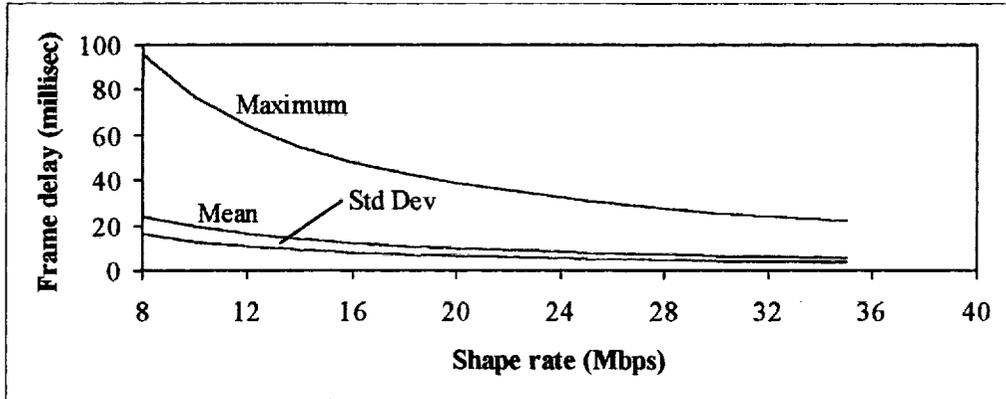


Figure 13 - Frame delay results from shaping experiment #1 for high-CoV video streams

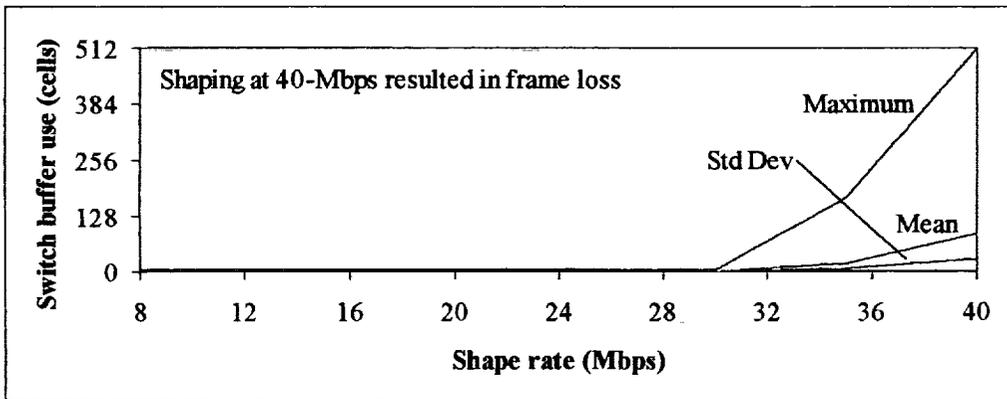


Figure 14 - Buffer use for switch #1 high-CoV video streams (for experiment #1)

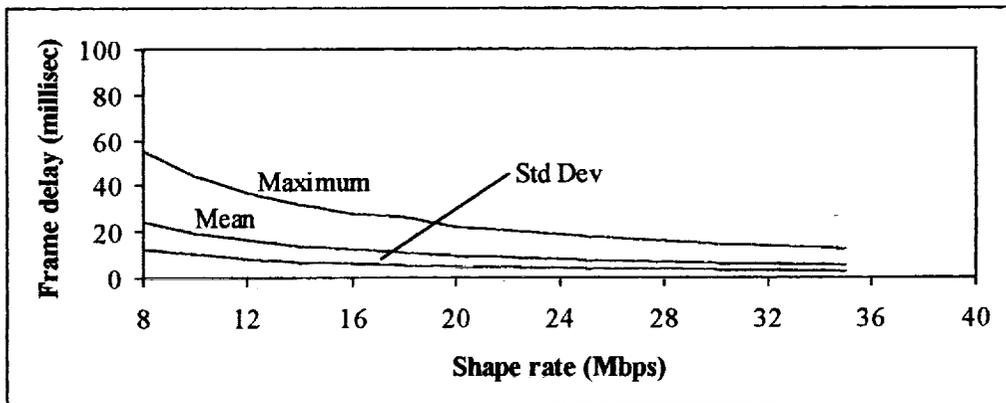


Figure 15 - Frame delay results from shaping experiment #2 for low-CoV video streams

## 8. Summary and Future Work

This work has demonstrated the use of performance modeling and the use of the CACI Predictor and Comnet tools. Modeling can be used to predict network performance and tune parameters for optimum performance. It

was shown that the current FTXS can support "a few" video streams with current load levels and traffic mixes. ATXS with its switched infrastructure offers, at least, a ten fold improvement for transport of video compared to FTXS. It was shown that the characteristics of a video stream has a direct correlation on its performance in a network. It was also shown that shaping of video streams is necessary to prevent overflow losses. Such overflow losses would both destroy video quality and affect performance of data application (i.e., by possibly inducing data traffic to also be lost). The Comnet simulation model can be used to explore parameter setting for video shaping that optimize the trade-off between video frame delay, allocated bandwidth (for video), and switch buffer utilization.

Future work should focus on fully validating the models against "real networks". This could be done by injecting, measuring, and comparing probe traffic in the real FTXS and doing the same in the model of FTXS. The developed models can be used to predict when the FTXS when "run out of room". To do this, long term trending of traffic must be done. This could be done using RMON, or some other traffic collection, probes and careful collection, archiving, and trending of traffic data over many months. In addition, the adoption of new technologies, such as streaming video, need to be fully understood. Finally, parameter settings for ATM links for transmitting shuttle launch video (e.g., to set-top MPEG-2 decoder boxes) can be optimized using the developed Comnet model. To do this, the "specs" on the actual MPEG-2 encoders to be used at NASA-KSC and sample video streams from these encoders will be required. Thus, this future work can provide useful input and validation to future set-top box projects within the AND group.

## Acknowledgments

The author would like to acknowledge Gregg Buckingham, Jane Hodges, Ray Hosler, and Judie Gilliam of the NASA-KSC University Programs Office, John Schnitzius of NASA-KSC Development Engineering, and Chris Kerios and Paul Kaun of Dynacs Engineering. The FTXS network topology was generated by Chris Knoell of Dynacs Engineering.

## References

- [1] CACI Products Company, Simulation Products, 1998. URL: <http://www.caciasl.com/>.
- [2] K. Christensen and C. Kerios, "Network Modeling and Planning Tool Specification," NETX1-0081 (Kerios, Chris, Subject 1540 Modeling and Planning Tool) EDC Purchase Requisition, August 1998.
- [3] Fiber Distributed Data Interface, Media Access Control (MAC), ANSI X3.139-1987, ISO 9314-2:1989.
- [4] P. Kaun and K. Christensen, "Crunch - A Software Program to Characterize Conversation-Pair Traffic Data from a Local Area Network (LAN) Sniffer Trace File," *Disclosure of Invention and New Technology (Including Software)*, NASA-KSC, July 28, 1999.
- [5] L. Kleinrock, Queueing Systems, Volume 1: Theory. John Wiley and Sons, New York, 1975.
- [6] KSC Communications Systems Upgrade Preliminary Engineering Report (PER). *KSC-DL-3700*, September 1996.
- [7] A. Law and D. Kelton, Simulation Modeling and Analysis. McGraw-Hill, New York, 1982.
- [8] D. Menasce, V. Almeida, and L. Dowdy, Capacity Planning and Performance Modeling. Prentice Hall, Englewood Cliffs, New Jersey, 1994.
- [9] The MPEG Home Page, 1999. URL: <http://drogo.cselt.stet.it/mpeg/>.
- [10] Network Associates, Sniffer Network Analyzer, 1999. URL: <http://www.networkassociates.com/>.
- [11] O. Rose, "Simple and Efficient Models for Variable Bit Rate MPEG Video Traffic," *Performance Evaluation*, Vol. 30, No. 1-2, pp. 69 - 85, July 1997.
- [12] Tools Page for Ken Christensen, 1999. URL: <http://www.csee.usf.edu/~christen/tools/toolpage.html>.
- [13] Traffic (Version 2.1), Knoltex Corporation, 1999. URL: <http://www.knoltex.com/>.

