Proceedings of the 1993 Winter Simulation Conference
G.W. Evans, M. Mollaghasemi, E.C. Russell, W.E. Biles (eds.)

A SIMULATION MODEL FOR ASSESSING NETWORK CAPACITY

Ankur R. Hajare
Daniel T. Wick

The MITRE Corporation
1120 National Aeronautics and Space Administration (NASA) Road One
Houston, Texas 77058, UNITED STATES OF AMERICA

ABSTRACT

NASA has procured a large scale software development facility, known as the Ground Systems Development Environment (GSDE), to support the development of the ground systems software for Space Station Freedom. The GSDE is geographically distributed and includes Local Area Networks (LANs) and T1 lines connecting users at multiple contractor facilities to a mainframe complex at the Johnson Space Center (JSC). In order to assess the adequacy of the planned GSDE network, a discrete event simulation model was developed using Network II.5. In addition to models of each Ethernet segment and each communication link, the network model included detailed models of the architecture of key components, such as routers. A utilization model was developed to estimate the traffic that would be generated by the network users. The simulation model of the GSDE network yielded component utilization, queue lengths, and other performance data. A nominal scenario was simulated, as well as five other potential scenarios that can result in higher network traffic.

1 INTRODUCTION

The Space Station *Freedom* Program requires that operational software be developed with the tools, rules, and procedures of NASA's Software Support Environment (Joyce 1989). The purpose of this support environment is to promote commonality of software development environments across the program, thereby increasing maintainability and decreasing the life-cycle costs of the developed systems.

The Ground Systems Development Environment (GSDE) (Price 1990) has been developed at JSC in response to this requirement. The components of the GSDE are geographically distributed. They include a number of workstations and servers, connected by LANs, at the facilities of two contractor organizations (referred to herein as Contractor A and Contractor B). The contractor facilities are connected by T1 communication lines to a mainframe computer complex at JSC.

The GSDE will be used by a large number of software development personnel. Proper performance of the various GSDE components is essential to maintaining the high productivity of the personnel using the system. Previous studies performed by MITRE determined the initial computer resources required for this system (Pelnik 1988, Wick and Pelnik 1991). This study assessed the capacity of the GSDE network that connects the host at JSC to the remote contractor facilities. Figure 1 illustrates the network under investigation.

The simulation study described herein differs from the numerous LAN performance studies that appear in the published literature. Most LAN performance studies focus on single segment performance. However, in this study the end-to-end performance of a network was assessed. In such a case, the performance of inter-networking devices (bridges, routers, and gateways) can be more important than the performance of the transmission medium (Boggs 1988). Under certain circumstances, inter-networking devices can become bottlenecks (Salwen 1988). Loss of packets by these devices results in error conditions and re-transmission (Quiat 1991) which deteriorates end-user response times. Hence, it is important to analyze the performance of inter-networking devices under various conditions that are encountered in an operational network.

Vendors of inter-networking devices provide performance specifications for their products. Since no standards presently exist (Salamone 1990), different metrics are reported by different vendors. Information about the conditions under which the performance data was derived is generally not provided by vendors. Since the testing methodology is not standardized either, each vendor can create tests which demonstrate their own products to be superior (Bradner 1991). Therefore, the test results that are available are seldom applicable to a real situation.

For example, the tests are performed under conditions that are not typical of what is encountered in actual network usage. Usually, tests are performed with all packets of one size that arrive at a steady rate. Consequently, the effect of differences in buffer sizes is not demonstrated. In contrast, network traffic in the real world is bursty, and buffer size does affect performance. Furthermore, most reported measurements are performed for unidirectional forwarding of all packets in a single stream with no other traffic on the network. Such test results, though not directly usable, can be used to calibrate performance models of inter-networking devices. The model can then predict performance for bursty, multiple data streams that contain a random mix of packets of various sizes.

Full scale testing of a multi-segment network for a comprehensive set of scenarios is not practical because of the large amount of test equipment and effort that would be required (Bradner 1991). Therefore, modelling is a practical alternative to assessing end-to-end performance of a large multi-segment network.

2 MODELLING TOOLS

Performance models are either analytic models or simulation models. Several analytic models have been developed for single segment LANs (Stallings 1990, Boggs 1988). However, no adequate analytic models have been reported for inter-networking devices. Analytic models are based on assumptions which convert a real world problem into one that is amenable to a closed form solution. Simulation models, on the other hand, do not require such drastic or extensive assumptions.

Analytic models usually predict only steady-state conditions, whereas simulation models demonstrate the effect of transients and the effects of initialization. For example, a typical learning bridge re-builds the address table every few minutes. Such transient conditions are best studied by means of a simulation model. Other transient conditions amenable to simulation modelling include broadcast packets creating a broadcast storm.

Simulation models can be developed using either a general purpose simulation language (such as GPSS or Simscript®) or a network modelling tool. General purpose simulation languages provide more flexibility and power but are harder to use. Network modelling tools enable quicker development of models but are relatively restricted in their capabilities. Examples of network modelling tools are Network II.5®, Lannet II.5®, Block Oriented Network Simulator™ (BONeS™), and LAN-SIM™. In addition to these commercially available tools, several large organizations, such as IBM and AT&T, have their own modelling tools for in-house use (Van Norman 1988).

The tool used for the simulation study was Network II.5, which is marketed by CACI Products, Inc. of La Jolla, California. Version 6.0 of Network II.5 was used on an IBM compatible mainframe. Network II.5 builds a discrete event simulation model from a model definition consisting of basic entities that include processing elements, storage devices, transfer devices and software modules. Each processing element has a set of instructions. Software modules, which consist of instructions, run on processing elements. These modules have fixed or probabilistic execution times. Processing elements can send messages via transfer devices to other processing elements or to storage devices. Messages queue at processing elements where they are processed by software modules. Also, software modules can queue for execution on processing elements. Network II.5 provides information on queue lengths and queueing delays, and it features scheduling mechanisms and priority disciplines. A random number generator and most of the commonly used statistical distributions are built into Network II.5. Although Network II.5 is written in Simscript II.5®, no interface is provided to user-written Simscript II.5 code. A description of Network II.5 is provided by CACI (CACI 1989).

Network II.5 contains built-in models for transfer devices that use collision, token ring, and other protocols. A specific LAN segment is, therefore, modelled by an appropriate selection of parameters. In addition to the built-in network protocols, Network II.5 provides the primitives necessary to model networking devices such as bridges, routers, gateways, communications controllers, and front-end processors.

Network II.5® does not model at the physical layer. Thus, it does not model signal propagation along with phase shift, jitter, and error conditions. Network II.5® has a fixed sized collision window for each Ethernet® segment, whereas in reality it is a function of distance. Also, the inter-frame gap is fixed for a LAN. Thus, Network II.5® cannot handle variations in Network Interface Unit (NIU) speed that result in varying inter-frame gaps (Rickert 1990). However, performance differences between NIU drivers (Freund 1991) can be incorporated into a Network II.5® model.

3 NETWORK ARCHITECTURE

The GSDE network consists of five Ethernet segments, two T1 communications links, three routers, and a Channel Attachment Unit (CAU) that is connected to a block multiplexor channel on an Amdahl 5890-300E mainframe computer system. These components and their interconnections are shown in Figure 1. All three routers in the network are of the same type. They contain multiple microprocessors and a proprietary system bus. The bus is a high-bandwidth non-standard bus that is loosely patterned after the NuBus with the addition of critical features for cache support. It is a 32 bit wide bus with a clock rate of 16 MHz. The Central Processing Unit (CPU) of the router is a Motorola 68030 microprocessor running at 30 MHz. In addition, the routers contain bit-slice RISC microprocessors running at 16 MHz. The number of these RISC microprocessors depends upon the configuration of each router. The routers have 4 megabytes of RAM on the same board as the CPU.

Each router in the network is configured differently, as described below. The GSDE router, illustrated in Figure 2, contains three boards attached to the system bus. Located on one of the boards is the CPU and memory. The other two boards hold processors and memory that are specific to the Network Interface Units (NIUs). Each of these boards has one T1 NIU and one Ethernet NIU. The two T1 NIUs of the router are connected to T1

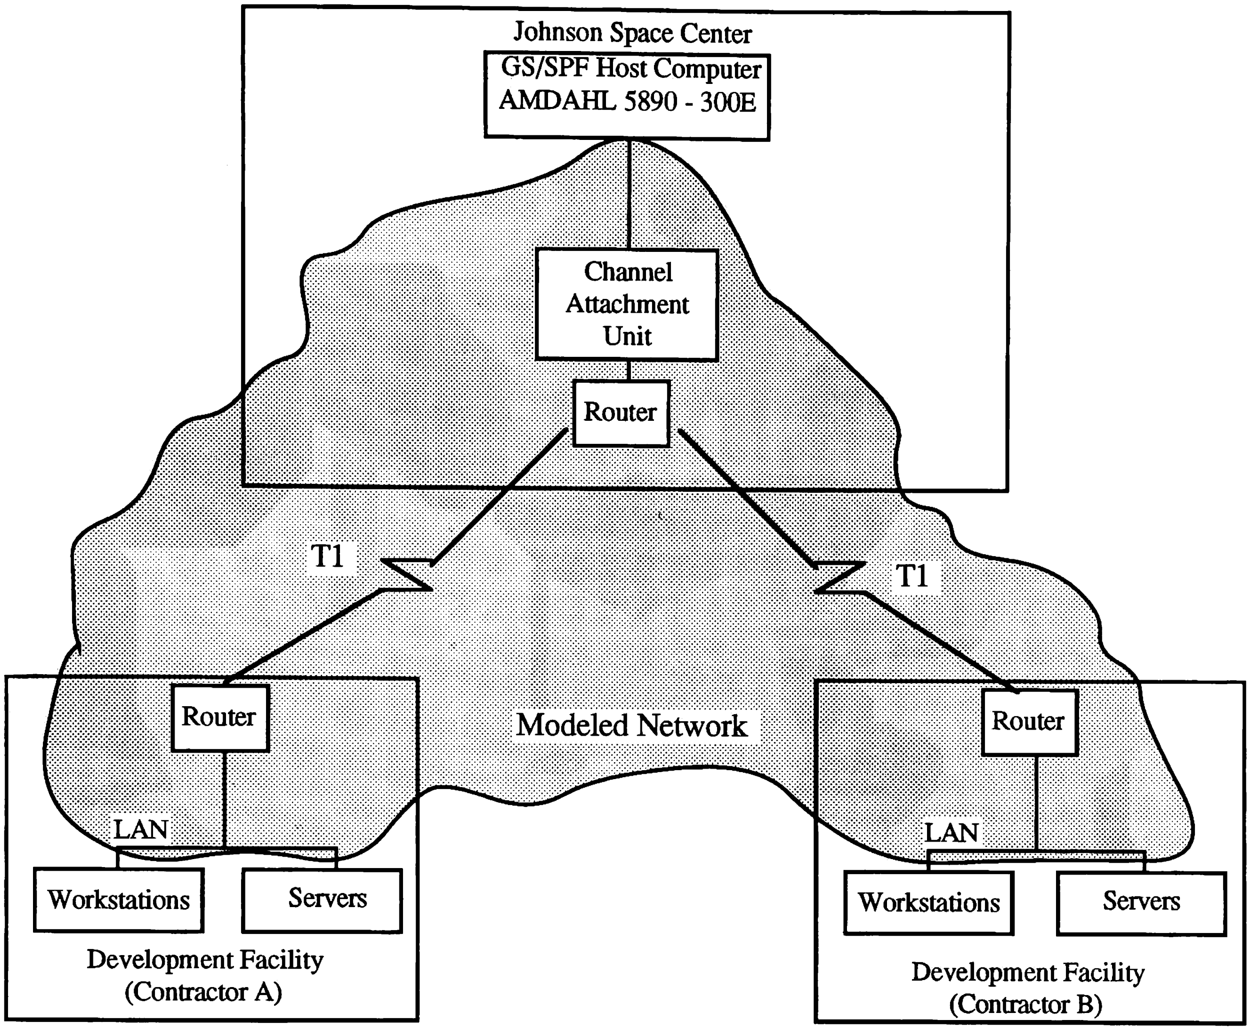

Figure 1: GSDE Network Architecture

communication lines that are, in turn, connected to T1 NIUs in routers at the development contractor sites. One of the Ethernet NIUs is connected to JSC's center-wide network. The second is connected to a CAU which is described later.

As shown in Figure 3, the Contractor A router also contains a bus and three boards. Though the CPU board is the same as in the GSDE router, the other two boards differ in their configuration. Both of these boards hold four NIUs, two for T1 and two for Ethernet. One of the T1 NIUs on the first board is not utilized at this time while the second is connected to the GDSE router via a T1 communication line. The other two T1 NIUs are connected to non-GSDE facilities at JSC. One of the four Ethernet NIUs is connected to the GSDE development facility within the Contractor A facility. The other Ethernet NIUs are connected to non-GSDE LANs. It has been reported (Edwards 1992) that the activity on these other Ethernet NIUs is insignificant and need not be considered in a performance or capacity evaluation at this time.

As shown in Figure 4, the Contractor B router also contains a bus and three boards. However, the boards differ in NUI configuration from both previous routers. One of these boards holds three NIUs, one for T1 and two for Ethernet. The other holds only two Ethernet NIUs. The T1 NIU is connected to the GDSE router via a T1 communication line. One of the four Ethernet NIUs is connected to the GSDE development facility within the Contractor A facility. The other Ethernet NIUs are connected to non-GSDE LANs. It has been reported (Lancaster 1992) that the activity on these other Ethernet NIUs is insignificant and need not be considered in a performance or capacity evaluation at this time.

The CAU acts as a high-speed interface between the mainframe host and the GSDE router. The CAU has a rated data transfer rate of 4.5 megabits per second to an Ethernet LAN and is capable of supporting three separate

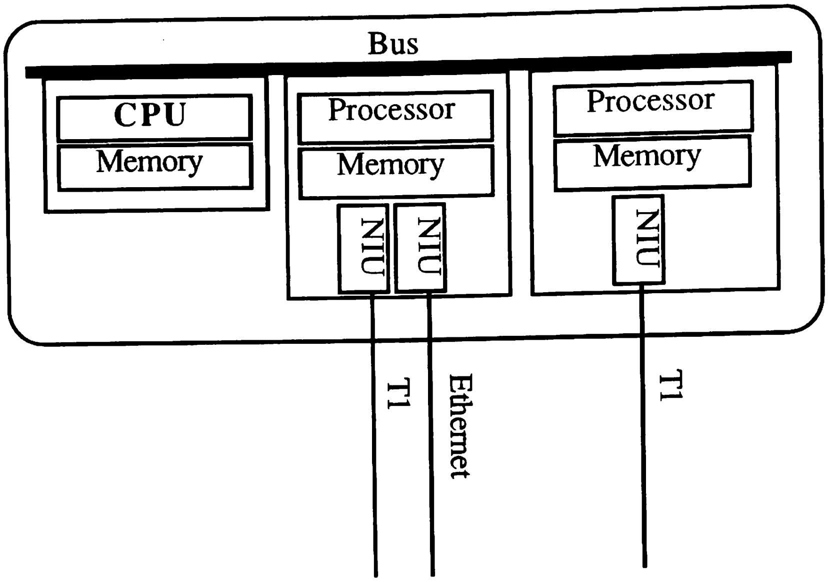

Figure 2: Configuration of GSDE Router

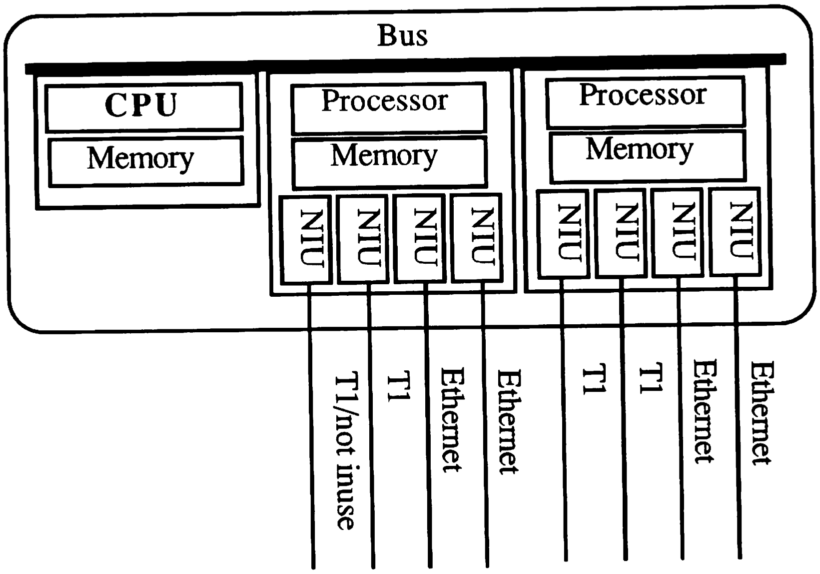

Figure 3: Configuration of Contractor A Router

Ethernet connections. The CAU is connected to a block multiplexor channel and uses three subchannels for each network connection, thus allowing three channel programs to be active in the mainframe concurrently. The CAU can transfer data to and from the mainframe channel at a 3.0 megabytes per second rate and can buffer the data within its internal memory.

The CAU is supported by software that resides on the

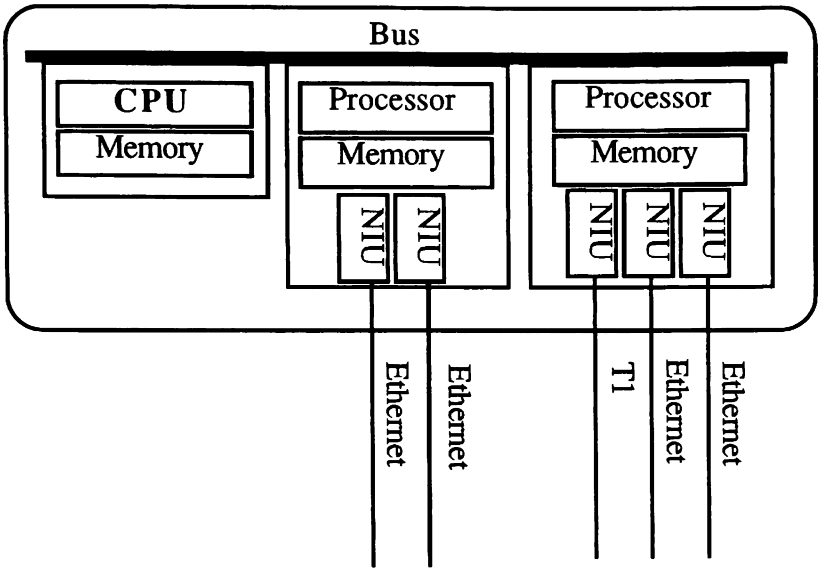

Figure 4: Configuration of Contractor B Router

mainframe host. This host resident software, which runs under the MVS operating system, makes the mainframe a TCP/IP node and provides host access to all the workstations on the GSDE network.

4 NETWORK MODEL

The integrated network model includes models of three types of components and models of their interconnections. The three component types include: transaction generators, routers, and a CAU. The physical resources that were modelled are illustrated in Figure 5.

As shown in Figure 5, the model contains four transaction generators. These transaction generators create traffic for the network model. Two of the generators represent the traffic generated by the GSDE users at the two software development sites, i.e. the traffic that flows toward the mainframe host. The traffic generator for the mainframe creates the traffic flowing from the host to the two software development sites. Finally, the fourth traffic generator creates traffic that represents electronic mail usage. The four traffic generators created data packets using a Poisson process that is built into the modelling tool. Sixteen statistical distribution functions were specified in the model in order to create the full complement of network traffic.

Network II.5 provides built-in support for network protocols, such as CSMA/CD and token passing. Thus, the Ethernet segments and T1 links in the GSDE network were modelled as transfer devices with the appropriate set of parameters. Models of the routers and the CAU were developed using Network II.5®, based on vendor provided information about the architecture and performance of each device. Given the architecture, its translation into Network II.5® terms was fairly straightforward in most cases. Buses were modelled as transfer devices, processors as processing elements, and NIUs were modelled as processing elements with buffer memory and I/O delays. The models were calibrated using reported performance data supplied by vendors, independent testing laboratories, and published results such as (Molloy 1992), (Bradner 1991), (Bradner 1992), and (Salamone 1991). Processing times for filtering and forwarding TCP/IP packets were obtained from (Hindin 1991).

The complete network model contained approximately 2000 lines of Network II.5 code. This included models of 14 processing elements, 11 transfer devices, and 108 software modules.

5 MODELLED TRAFFIC AND SCENARIOS

Currently, no network traffic measurements are available for the GSDE network. However, the Information Systems Directorate (ISD) at JSC had developed a model for network traffic generated by the typical user, based on their experience with other computing systems at JSC. The ISD traffic model was used for this simulation study. Some of the characteristics of this traffic model are shown in Table 1. As shown in the first column, there are four basic message lengths, ranging from a short message of 100 bytes (which might represent a single command entered by a user) to a large message of 50 Kbytes (which might represent downloading a code module). The ISD model represented the total traffic generated by a user, both requests to a host and host response. For the purpose of this study, the total traffic was equally distributed between the user and the host, modeling the data traffic anticipated by the current operations concept.

The contractors using the facility have provided staffing estimates for each fiscal year of their contract. Using a personnel distribution model described in (Wick and Pelnik 1991), the number of network users was estimated. The data indicates that the peak utilization year for both contractors will occur during Fiscal Year 1994 (FY94). Using the staffing estimates for FY94, six scenarios were developed for this simulation study. The six scenarios are as follows:

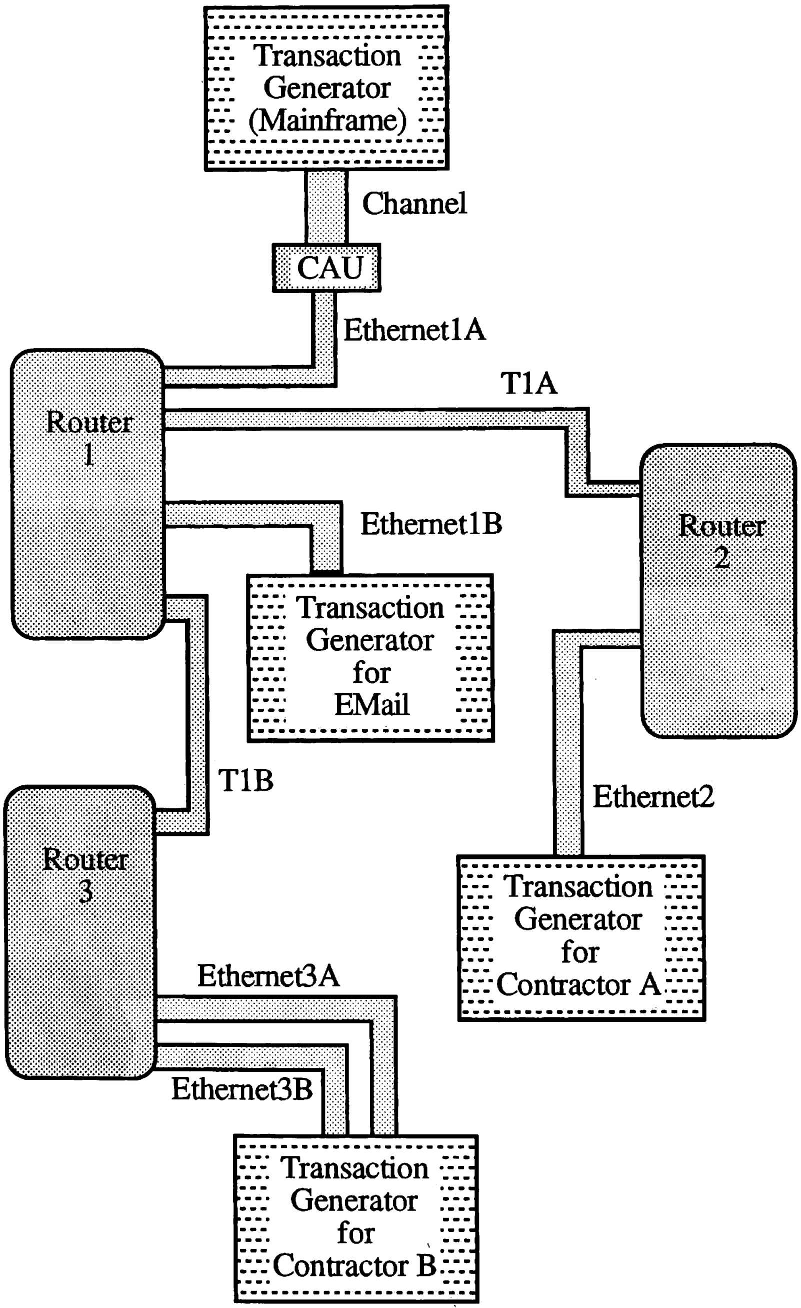

Figure 5: Configuration of Network Model

Table 1: User Traffic Assumptions per Hour

Message Size	Number of Messages	Total Traffic
100 B	190	190 KB
1.5 KB	95	143 KB
7.5KB	61	458 KB
50 KB	17	850 KB
TOTAL	363	1,469KB

1) A basic FY94 estimate.

2) A 100 % increase in the size of large transactions. This scenario demonstrates the effects of changing one assumption, specifically the message size that corresponds to modules being uploaded and downloaded.

3) Use of electronic mail in addition to the basic estimate of Scenario 1. This scenario is based on the current configuration which uses the GSDE network both as a production network and as a path for accessing electronic mail on other computer systems at JSC.

4) Burst use of electronic mail at the beginning of the work day, in addition to the basic estimate of Scenario 1. In this scenario it is assumed that all users access electronic mail during a 15 minute period at the beginning of the work day. In contrast, the previous scenario assumed that electronic mail usage was evenly distributed throughout an 8 hour work day.

5) A morning log-on burst in which all users log-on and request one code module down load within a 15 minute period. This scenario is a test of the network to cope with a burst of data traffic.

6) A 25 % increase in staffing and a 50 % increase in work per staff member. This scenario was chosen because past experience has shown that software development projects often encounter unforeseen problems resulting in a production slip. This scenario is based on an assumption that the contractors will increase the number of development staff by 25 % and that staff members will increase network traffic by 50% in order to meet schedule deadlines.

6 SIMULATION RESULTS

The simulation model in Network II.5 was developed and run on an Amdahl 5890-300E at JSC. After debugging and calibrating, the model was executed several times for each scenario using different random number streams. For four of the scenarios, each simulation run represented one hour of operation of the GSDE network. For the two scenarios depicting morning log-on burst traffic, each simulation run represented 15 minutes of operation. For the two scenarios depicting morning log-on burst traffic, each simulation run represented was run until it crashed.

The data collected from the simulation runs included queue lengths, packet transfer times, and the utilization of various resources such as processors, buses, and LANs. Due to the limited graphics capability and report generation capability of Network II.5®, it was sometimes necessary to use other software packages to analyze, format and present the data generated by the tool. The results are summarized in Table 2 which lists the utilization of the transfer devices for the first four scenarios.

The simulation demonstrated that the network is quite capable of handling the nominal data traffic expected in FY94, the year with the highest anticipated number of users. The queueing delays were minimal and the response time was excellent. As shown in Table 2, the device with the highest utilization was the T1 link to the Contractor A facility. However, the utilization was only marginally above 50% and was not a source of problems. In Scenario 2 the network was just barely capable of handling the traffic. All messages reached their destination but 2% of the messages had unacceptable delays due to long queue lengths and wait times. Long queues were formed at the T1 link to Contractor A and also within Router 1. Utilization figures for the components of this router suggested that an alternate router configuration could alleviate some congestion. However, alternate router configurations have not yet been simulated and represent an area for further work.

Scenario 3 indicated no problems. The traffic in this scenario was higher than the nominal traffic, but the network was still capable of handling it without a problem. In Scenario 4 the network had problems with a

Table 2: Transfer Device Utilization

DEVICE	SCENARIO 1	SCENARIO 2	SCENARIO 3	SCENARIO 4
T1-A	61%	91%	62%	66%
T1-B	44%	67%	45%	47%
Ethernet 1A	14%	22%	14%	14%
Ethernet 1B	0%	0%	<1%	1%
Ethernet 2	8%	13%	8%	9%
Ethernet 3A	6%	7%	6%	6%
Ethernet 3B	<1%	<1%	<1%	<1%
Channel	6%	9%	6%	6%

small percentage of messages. Although all messages eventually reached their destination, less than 1% experiences unacceptable delays. One source of the delay was the queue for the T1 link to Contractor A. Another was a bottleneck on a board within Router 1. As for Scenario 2, the simulation data suggested that a different router configuration (with an additional board) might solve the problem.

Scenarios 5 and 6 are not shown in the table because the network could not handle the traffic in these scenarios. The simulation runs were aborted due to excessive queue lengths that resulted in Network II.5 error messages. In Scenario 5 both the T1 links were overloaded, resulting in message queues that continued to grow. In Scenario 6 the T1 link to Contractor A was overloaded. The T1 link to Contractor B was capable of carrying all the traffic. Both Scenario 5 and 6 indicate the need for additional T1 lines to the contractor facilities in case the network traffic becomes as high as is depicted in these scenarios.

7 CONCLUDING REMARKS

The simulation models described in this paper indicate that the GSDE network will be capable of handling the traffic anticipated in the nominal scenario. However, the results of the study are sensitive to certain assumptions about the network traffic which, in turn, have been derived from the operations concept for the GSDE. Thus, scenarios which result in higher network traffic (arising from increased staffing or a different user profile) can result in degraded network response or even an overloaded network. Also, the high traffic of a morning log-on burst can overload the network for a short period of time.

This simulation study has been based upon certain assumptions about the traffic on the network. As GSDE use increases, the network traffic should be measured and compared with the traffic used in this study. If there is a significant difference, the simulation models described here should be run again with an updated network traffic profile.

ACKNOWLEDGEMENTS

This work was sponsored by NASA contract number NAS9-18057. Some of the performance data used in this study was provided by Scott Bradner of Harvard University and by the Enterprise Technology Center in Houston, Texas.

REFERENCES

Boggs, D. R. et al. 1988. Measured Capacity of an Ethernet: Myths and Reality. In *Proceedings of the Sigcomm '88 Symposium.* Association for Computing Machinery, New York, New York.

Bradner, S. O. 1991. Testing Multiprotocol Routers. *Data Communications* 20:2:71-86.

Bradner, S. O. 1992. Ethernet Bridges and Routers. *Data Communications* 21:3:58-69.

CACI. 1989. *Network II.5® User's Manual.* California:CACI Products Company.

Edwards, G. 1992. Personal communication.

Freund, M. 1991. Sizing Up Your Network. *LAN Technology* 7:1:17-18.

Hindin, E. M. 1991. RISC Comes to Routers. *Data Communications* 20:7:43-48.

Joyce, E. J. 1989 A System for the 21st Century. *Datamation* 35:10:51-54.

Lancaster, S. 1992. Personal communication.

Molloy, M. 1992. Bridge/Router Tests Yield Mixed Results. *Network World* 9:22:19-43.

Pelnik, T. M. 1988. A Software Support Environment Software Productions Facility Sizing Model. WP-88D0049, The MITRE Corporation, Houston, Texas.

Price, T. G. 1990. Space Station Level I/II Change Request, Ground Systems/Software Production Facility (GS/SPF) Initial Configuration. NASA Johnson Space Center, Houston, Texas.

Quiat, B. 1991. Bridging the WAN Gap: An In-Depth Look at Four T1 Remote Bridges. *Network Computing* (April 1991):18-24.

Rickert, J. R. Jr. 1990. Evaluating MAC-Layer Bridges.

Data Communications 19:5:117-122.

Salamone, S. 1990. The Need for LAN Benchmarks. *Network World* 7:37:51-74.

Salamone, S. 1991. Test Ranks Routers in Real-World Environment. *Network World* 8:21:41-52.

Salwen, H. et al. 1988. Examination of the Applicability of Router and Bridging Techniques. *IEEE Network* 2:1:77-80.

Stallings, W. 1990. *Local Networks*. 3rd ed. New York: Macmillan Publishing Company.

Van Norman, H. J. 1988. A User's Guide to Network Design Tools. *Data Communications* 17:4:43-48.

Wick, D. T. and Pelnik, T. M. 1991. A Software Support Environment Software Production Facility Sizing Model. WP-88D0049 Revision 1, The MITRE Corporation, Houston, Texas.

LIST OF ABBREVIATIONS

CAU	Channel Attachment Unit
CPU	Central Processing Unit
CSMA/CD	Carrier Sense Multiple Access/Collision Detect
NASA	National Aeronautics and Space Administration
MHz	megahertz
μs	microseconds
NIU	Network Interface Unit
LAN	Local Area Network
OSI	Open Systems Interconnection
RISC	Reduced Instruction Set Computer
sec	second
TCP/IP	Transmission Control Protocol/Internet Protocol

AUTHOR BIOGRAPHIES

ANKUR R. HAJARE is a member of the technical staff at the MITRE Corporation in Houston, Texas. His work has included simulation modelling, performance analysis, and systems engineering. He has a bachelor's and a master's degree in electrical engineering.

DANIEL T. WICK is a member of the technical staff at the MITRE Corporation in Houston, Texas. For the past eight years he has been supporting NASA in the areas of training requirements definition, systems engineering, and large scale simulator design. He holds a bachelor's degree in psychology and a master's degree in psychobiology.