


Integrated Computer Control System Countdown Status Messages Simulation

C. E. Annese

October 1, 1998



This is an informal report intended primarily for internal or limited external distribution. The opinions and conclusions stated are those of the author and may or may not be those of the Laboratory.

Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161

Integrated Computer
Control System

Countdown Status Messages
Simulation

FY98 LDRD Project

C. E. Annese
October, 1998

UCRL-ID-133242

1 Introduction

1.1 Background

The Integrated Computer Control System (ICCS) CORBA-Based Simulation LDRD will guide the design of the complex ICCS framework in its application to the NIF (NIF) and thereby reduce performance risks. A simulation study was performed as a Milestone to "Develop Shot Cycle Workflow Simulation." Status messages from shot countdown activities were included in a discrete event simulation to assess on commercial components and their utilization in current designs.

1.2 Objective

The purpose of the study was to use simulation to help validate the design of status messages during countdown will not overload the network; i.e., that the computer is sized to handle the expected traffic.

1.3 Scope

The scope of the present study was to layout the framework for the Status Message Simulation (SMCS) using the discrete event simulation package, SIMPROCESS, to determine the inputs for running the model, and to obtain estimates of resource utilization. As resources become available, inputs can be updated to include additional refinements. SIMPROCESS has a "Design of Experiments" capability which can be used to evaluate scaling ICCS component utilization and resource allocation on system behavior. This has not been implemented at this time.

2 Design Description

2.1 Countdown Activities

Figure 2-1 shows the activities during a typical NIF countdown. The figure is drawn to scale, with the 'countdown' arrow representing five minutes. The four activities are: 1) plasma electrode Pockels cell (PEPC) wavefront control, 2) preamplifier module (PAM) charging, and 3) main anode (MA) Status messages will arise from each of these activities as they monitor the status on to the supervisory console for that system. The rate and size of each of the application's front-end processors (FEP) and on to the supervisory console obtained by interviews with experts and from NIF interface control document description and latencies were obtained from network designers and were scaled to represent the network interfaces for each of the four countdown activities. The network latencies relevant to the simulation will be discussed in the next sections.

2.1.1 PEPC Conditioning Status Messages

There are three PEPC FEPs for each NIF cluster of 48 laser beamlines. They are: 1) a plasma pulser / gas and vacuum (PPGV) FEP, and 3) a switch point (SP) FEP. Table 2-1 provides a table describing the worst case scenario for PEPC device monitoring. Tables 2-1 and 2-2 show the peak operation message traffic from the PM and PPGV FEPs, respectively, during the five minute pre-shot countdown. There is no message traffic from the SP FEP during countdown. The PEPC FEPs travel from the FEP through a 10 Mb/s fiber; and from the 10/100 Mb/s ethernet switch through

line to the core ethernet and out to the supervisory consoles. Figure 2-2 the ethernet connections from client to server FEP. Details of the PEPC in the first ethernet switch are shown in Figure 2-3. The packet size and byte 2-2 will provide inputs to the simulation.

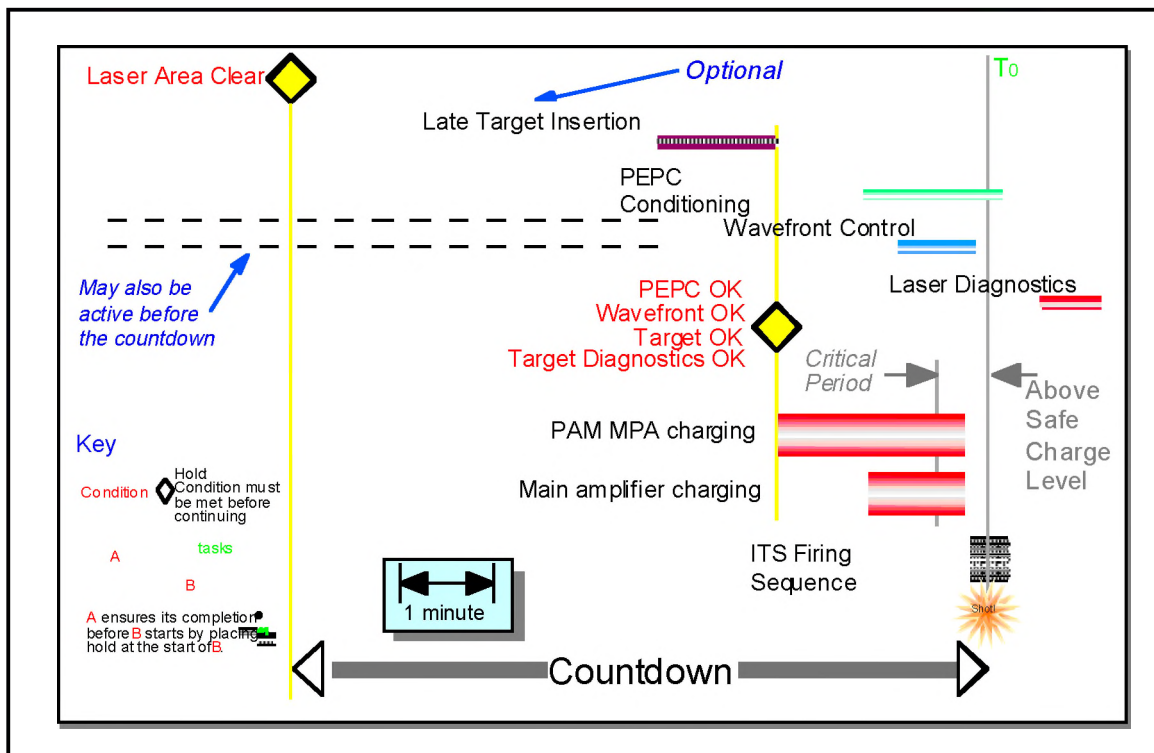


Figure 2-1. ICCS activities during countdown.

Table 2-1. Peak operation message traffic from PEPC PM FEP during countdow

PEPC FEP	Device	Device Count	Signals	Signal Device	Total Signals	Rate (Hz)	Msg / sec	Bytes / Msg	Bytes / sec
PM Solaris	plasma current	4	plasma current	1	4	0.2	1	16,000	16,000
PM Solaris	switch voltage	8	switch voltage	1	8	0.2	1	32,000	32,000
Total							2		48,000

Table 2-2. Peak operation message traffic from PEPC PPGV FEP during countd

PEPC FEP	Device	Device Count	Signal Signals	Signal Device	/Total Signals	Rate (Hz)	Msg / sec	Bytes / Msg	Bytes / sec
PPGV vxworks	simmer	48	current & voltage	2	96	0.2	1	48,000	48,000
PPGV vxworks	PEPC pressure	24	gas & vacuum pressure	2	48	1			
PPGV vxworks	PEPC gauges	84	pressure	1	84	1	1	528	528
Total								2	48,528




Figure 2-2. Typical client FEP to server FEP ethernet.

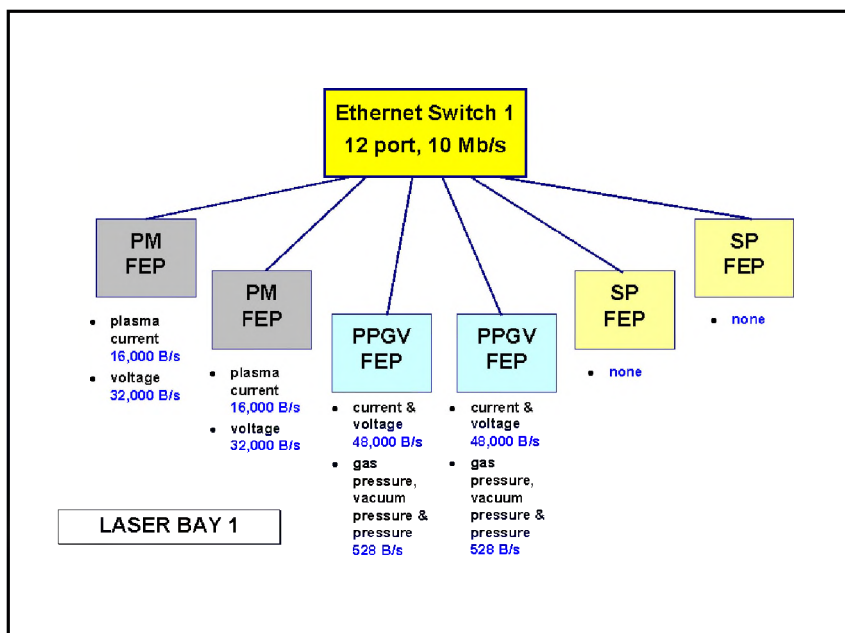


Figure 2-3. PEPC Ethernet system diagram.

2.1.2 Wavefront Control Status Messages

Wavefront control status messages during countdown were obtained from convey system designer. Figure 2-4 illustrates the size and types of messages used

2.1.3 PAM MPA Charging Status Messages

There are 48 PAM FEPs. Status message size and delivery rates were obtained with design engineers. There is one voltage status message for each PAM FEP is 4 Bytes long and is generated at once per second. Though there are four countdown of Figure 2-1 that appear to coincide with the start of the PAM (wavefront OK, target OK, and target diagnostics OK, the start of the PAM charging on receiving any of these OK signals; the charging starts two minutes prior summarizes this status message from the PAM for easy reference. As for the messages travel from the application FEP through a 10 Mb/s fiber; and from switch through a 100 Mb/s ethernet line to the core ethernet and out to th

The packet size and bytes/sec from Table 2-3 will provide inputs to the simulation the PAM Ethernet layout used in the simulation.

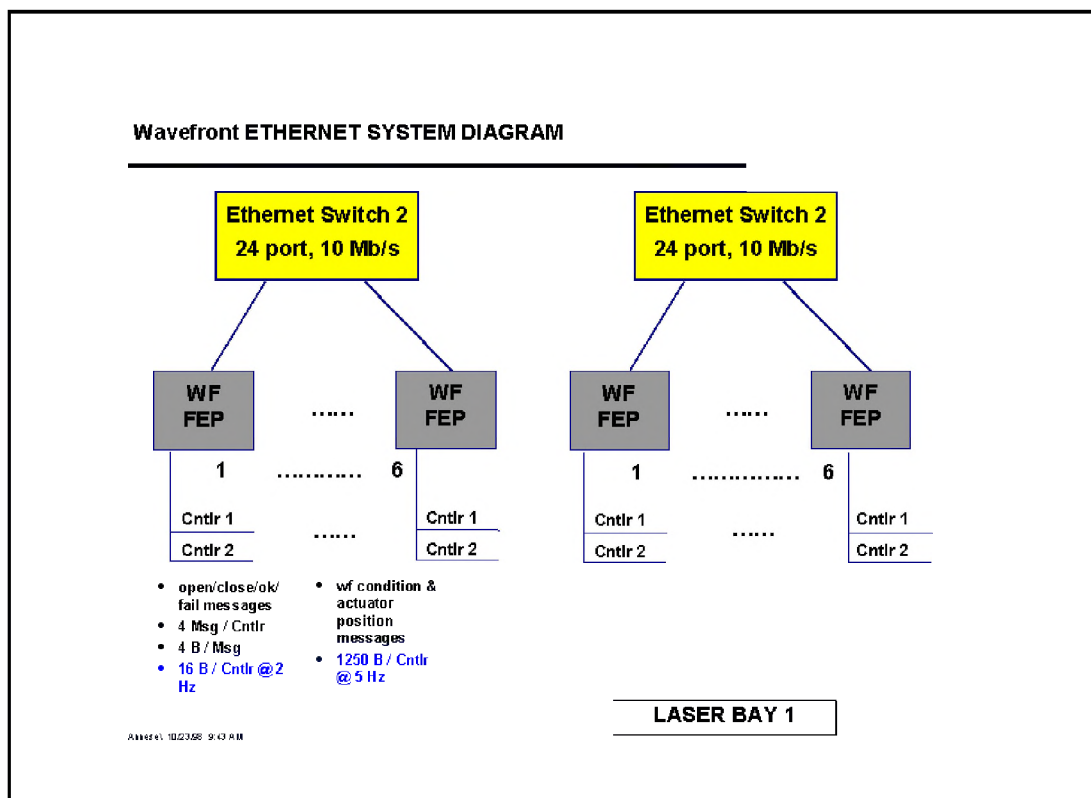


Figure 2 4. Wavefront message size and types used in the countdown simulation

Table 2-3. Peak operation message traffic from PAM FEP during five minute

PAM FEP	Device	Device Count	Signal Signals	/Total Signals	Rate (Hz)	Msg / sec	Bytes / sec	MsgBytes / sec
charging	48	voltage	1	48	1	1	4 * 48 = 192	292

2.1.4 Main Amplifier Charging Status Messages

Figure 2-6 shows the main amplifier power conditioning ethernet system diagram simulation. The main capacitors will start charging at ~77 seconds prior to ionization lamp check (PILC) capacitors will start charging at ~31 seconds status message during countdown, therefore, will consist of a floating point voltage starting at 77 seconds prior to shot, at half second intervals; the concurrent with a second floating point value for PILC capacitor voltage st to shot, also at half second intervals. The timing system will wait for per PILC modules before going into the final two seconds of countdown.

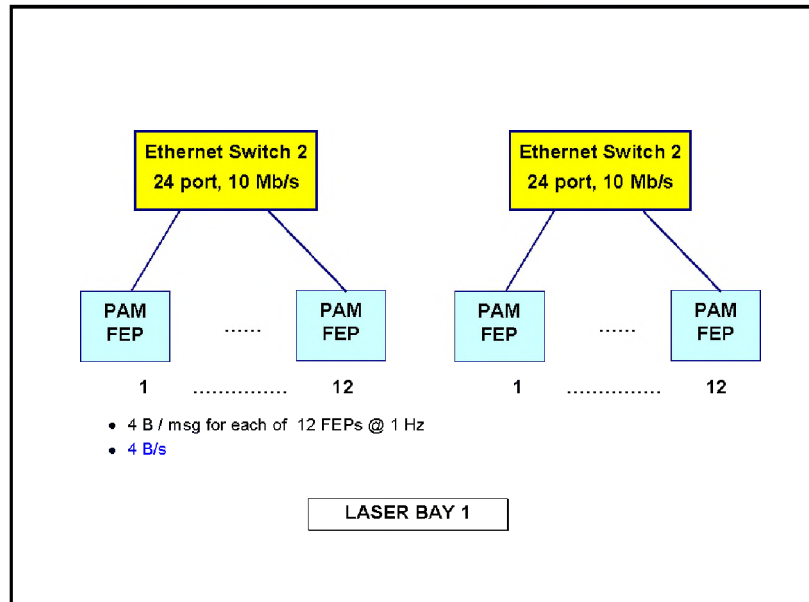


Figure 2-5. PAM Ethernet system diagram.

2.2 Transaction Time Measurements

2.2.1 FEP and Network Description

Experimental latency times for the FEPs were derived from measurements including application takes to interface with CORBA, with the transmission control protocol (TCP/IP) protocol, and the network interface card driver. Figure 2-4 shows an example of the various portions of the ICCS network from application FEPs, through Ethernet switch, and on to the supervisory FEP. These times were measured connecting a typical FEP client through Ethernet to a supervisory server. An example message is shown in Figure 2-5, often expressed as 6 hexadecimal 8-bit words separated by columns. The total message size sent is ~ 1500 bytes.

2.2.2 Message Overheads

Message sizes are increased by overhead added from CORBA, TCP, and IP. The following sections describe the size and types of information added to the message packets.

TCP Protocol¹

TCP provides timely and ordered data delivery, connection establishment, data termination, and full duplex. It adds around 20 bytes to the original message in packet switching mode, which means that data is separated into smaller packets. Sequence numbers, does not necessarily arrive at the destination in the same order at the destination in the proper order according to the sequence numbers. flow control and error checking, urgent ("out-of-band" data, data stream pause and reporting service failures). Ports 1 - 1023 are reserved for applications, ports 1024 to above can be used as long as no one else is using them. TCP lets you send data successfully. There are 16 bits for port numbers, so there are 64,000 ports available. The maximum packet size that can be sent prior to waiting for an acknowledgment is 64KB. TCP packet format is as shown in Table 2-5.

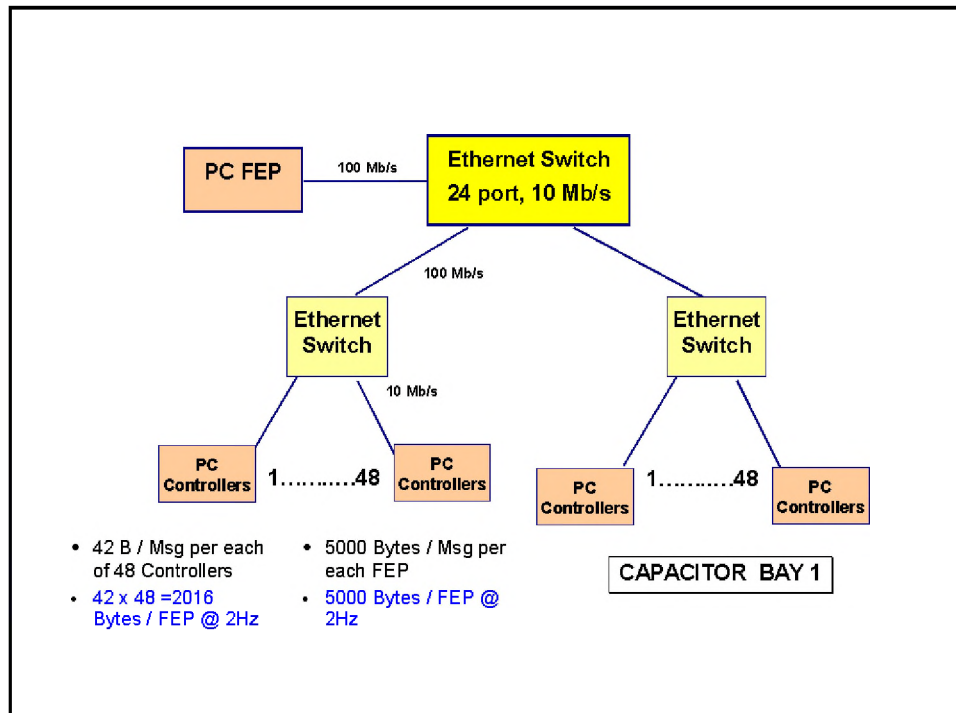


Figure 2-6. Main amplifier power conditioning Ethernet system diagram.

¹ Eugene Pinsky, Massachusetts Institute of Technology, TCP/IP Networking course notes from Boston University via satellite to LLNL, LTV 1366, 5/98, p. 4-6.

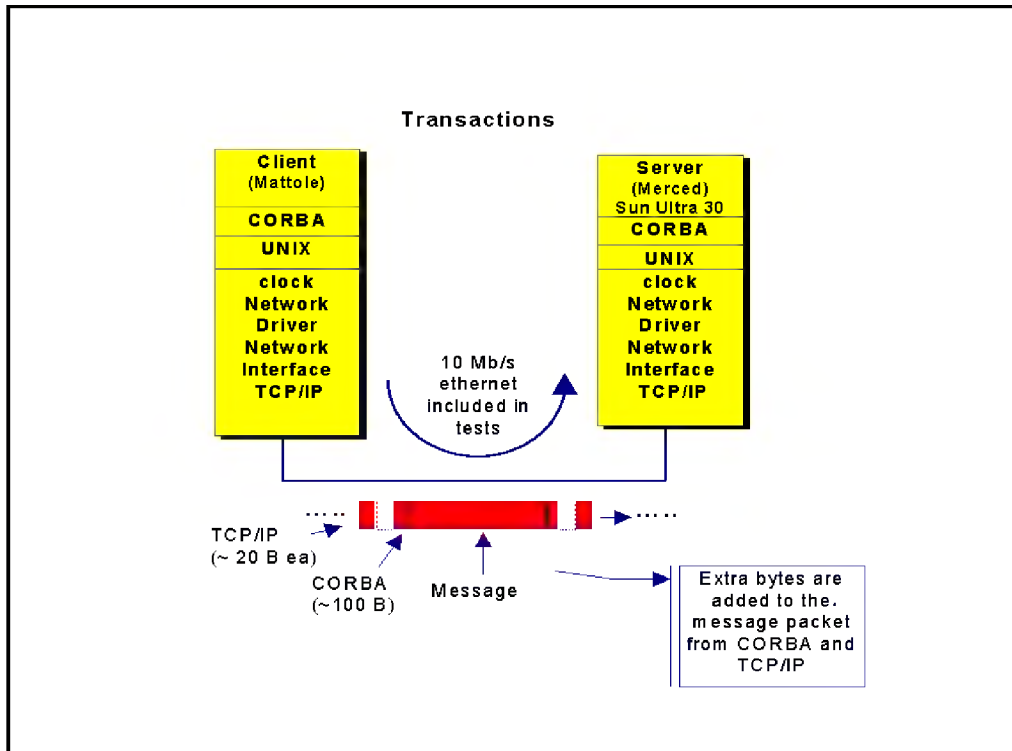


Figure 2-7. Measurements were made of transaction times between Sun Ultra 30 and 10 Mb/s ethernet in-between.

IP Protocol ¹

Internet Protocol makes "best effort" delivery, utilizes fragmentation and reassembly mechanism for end-to-end reliability, flow control or sequencing. Thus it provides the added functions and reliability. See Table 2-6 for details. The protocol identifier, such as TCP, UDP (user datagram protocol), or ICMP, for each datagram contains 32-bit IP addresses. IP adds 20 bytes in length to each message in addition to the 20+ bytes added by TCP. Common object-request broker architecture protocol adds about 100 bytes of data to each message.

CORBA

Corba adds about 100 bytes of data ² to each message.

¹ Conversation with Eric Stout, LLNL.

Table 2-4. Transactions per second from Mattole to Merced through 10Mb/s

Message Size (FPA)	Byte Size FPA * 4	Client Process Time (msec) / CPU Utilization		Server Process Time (msec)		Client Null CPU Utilization	Server Null CPU Utilization	Clock Time
2	8	.545	.44	.452	.37	.48	.25	1.233
4	16	.545	.44	.473	.38	.47	.26	1.229
8	32	.545	.43	.462	.37	.47	.28	1.254
16	64	.548	.43	.388	.30	.48	.29	1.284
32	128	.552	.41	.443	.33	.47	.28	1.346
64	256	.548	.38	.281	.19	.47	.31	1.457
128	512	.554	.33	.412	.24	.47	.29	1.700
256	1024	.566	.26	.544	.25	.46	.30	2.195
512	2,048	.634	.20	.643	.20	.47	.28	3.152
1,024	4,096	.733	.14	.932	.18	.48	.28	5.096
2,048	8,192	.917	.10	1.301	.15	.47	.31	8.934
4,096	16,384	1.128	.07	2.223	.13	.46	.31	17.060
8,192	32,768	1.600	.05	3.701	.11	.47	.30	32.758
16,384	65,536	2.583	.04	7.098	.10	.48	.26	68.134
32,768	131,072	4.578	.03	14.912	.11	.46	.31	132.143
65,536	262,144	8.769	.03	29.224	.11	.47	.30	260.339
131,072	524,288	17.384	.03	63.220	.11	.46	.30	516.352

Table 2-5. TCP adds at least 20 bytes in length (assumes 8 bits/byte) to e

0	4	8	12	16	20	24	28	32 (bits)
SOURCE PORT (APPLICATION)				DESTINATION PORT (APPLICATION)				TCP HEADER
SEQUENCE NUMBER								
ACKNOWLEDGEMENT NUMBER								
DATA OFFSET	RESERVE D	CONTROL (ACK FLAG)		WINDOW (FLOW CONTROL)				
CHECKSUM				URGENT CALL POINTER				
OPTIONS					PADDING			
DATA								

2.2.3 Experimental Measurements of Latencies

Assessment of the transactions times in Table 2-4 led to the creation of f and supervisor FEP latencies, and 10 Mb/s Ethernet latencies. An overhead transaction is attributed that leads to the following three formulations in

$$\begin{aligned}
 \text{Client FEP Latency} &= .000342 * \text{Bytes} + .54 \\
 \text{Server FEP Latency} &= .00011 * \text{Bytes} + .430 \\
 \text{10 Mb/s Ethernet Latency} &= .00087 * \text{Bytes} + .215
 \end{aligned}$$

The ratio of the formulation to the actual data is shown in Figures 2-8, 2 for client FEP, server FEP, and 10 Mb/s Ethernet latencies.

Table 2-6. IP adds at least 20 bytes in length (assumes 8 bits/byte) to ea

0	4	8	12	16	2	24	28	32 (bits)	
VERSION	HDR LENGT H	SERVICE TYPE	TOTAL LENGTH						IP HEADER
IDENTIFICATION				FLAGS	OPERATION				
TIME TO LIVE		PROTOCOL		HEADER CHECKSUM					
SOURCE INTERNET ADDRESS									
OPTIONS				PADDING					
DATA									

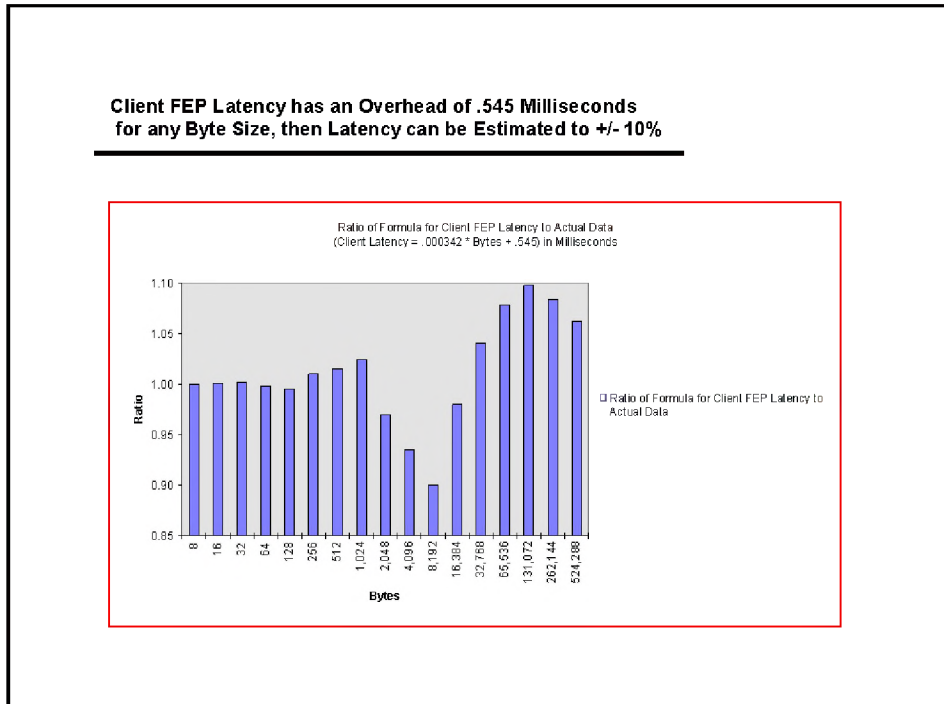


Figure 2-8. Client FEP latency can be estimated to +/- 10%.

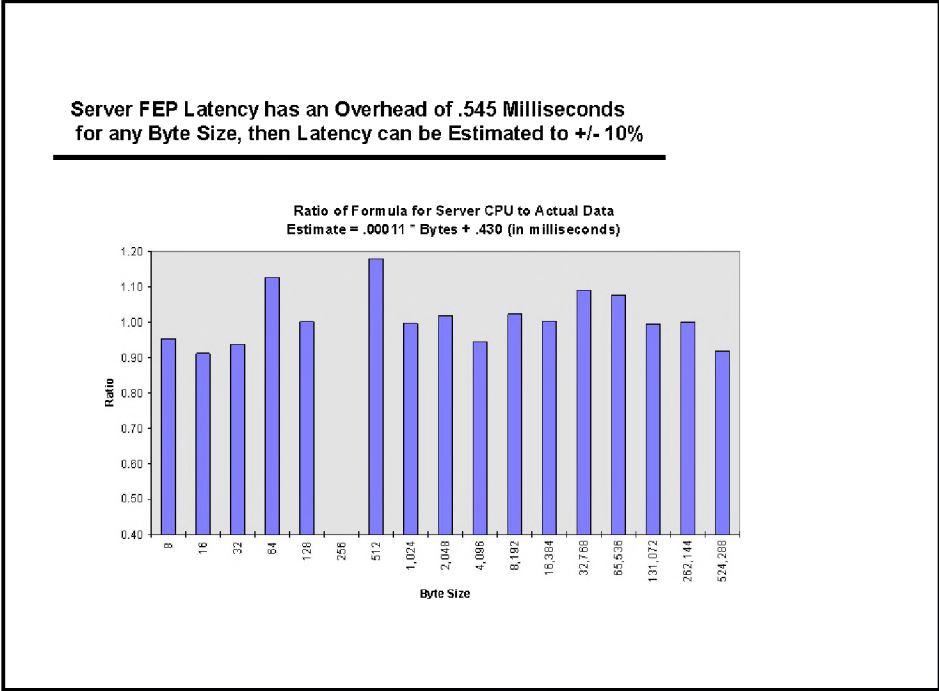


Figure 2-9. Server FEP latency can be estimated to +/- 10%.

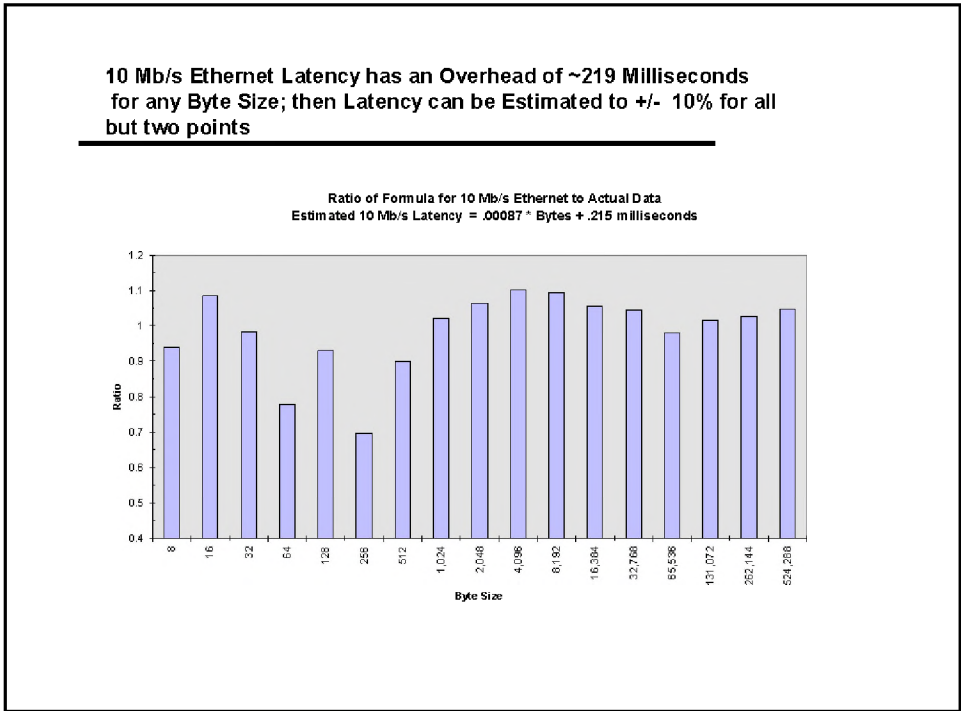


Figure 2-10. 10 Mb/s ethernet latency can be estimated to +/- 10% for all

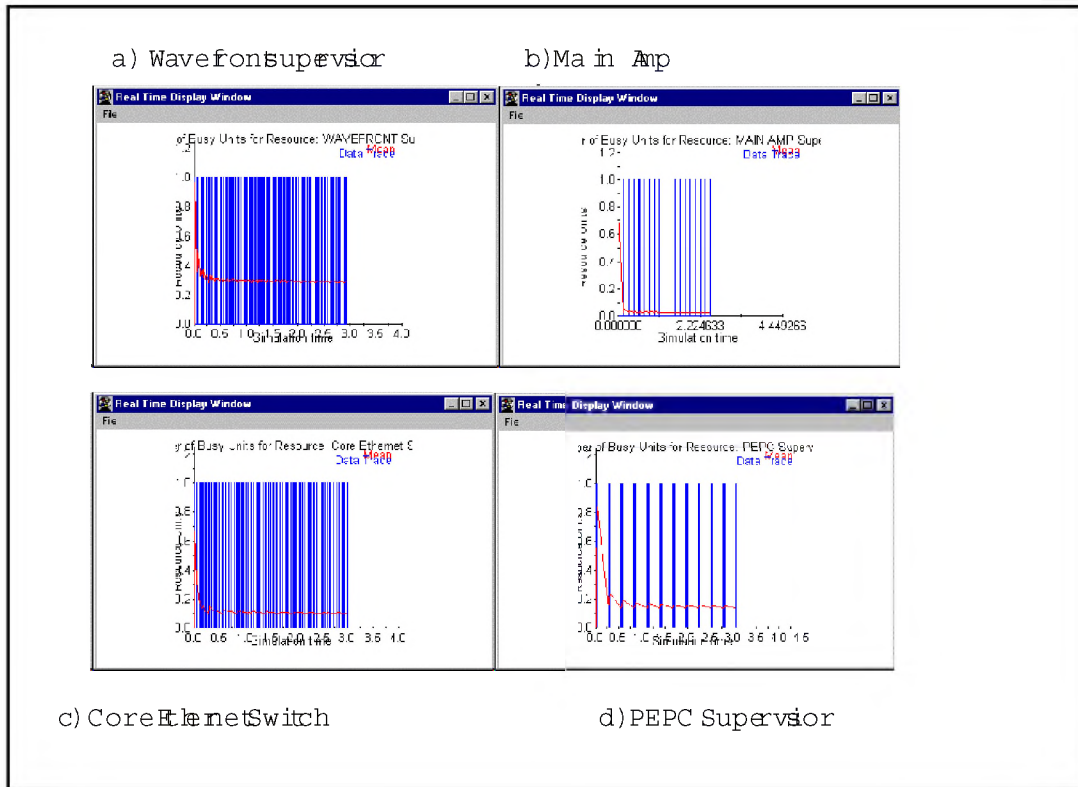


Figure 2-11. Average percent resources busy during busiest time of five-mir

The resource utilization vs simulation time is plotted in Figure 2-11 for each status messages. The x-axis represents simulation time in hours; however, it had shown that Simprocess did not work well using fractions of seconds as in the model. So, 0.28 hours on the graph, represents one second in the model.

The time that resources spent waiting for resources was also tallied and varied from milliseconds for the main amp power conditioning signals to around 80 milliseconds for current and voltage signals.

3 Discussion

To conservatively assess the demand on the system, the Simprocess model was used to represent the period when all messages occur simultaneously. Data was collected to show the average percent of resources busy. These vary from less than 0.001% for the Core Ethernet Switch to over 20% for the wavefront supervisor. They are shown in Figure 3-1.

4 Conclusions

The Simprocess status message simulation provides a good visual tool for understanding network utilization. Supervisory displays and their latencies will be included in the model when they become available. While point estimates have

³ Bar for wavefront Ethernet switch 1 also includes 10 Mb/s Ethernet plus the switch.

⁴ Conversation with B. Kettering who ran Status Monitor simulation with Simprocess and discovered that Simprocess did not work well using fractions of seconds as inputs to the model.

work, statistical distributions can easily be used in their place and a design will allow variation of parameters for optimization when loads start to look heavy.

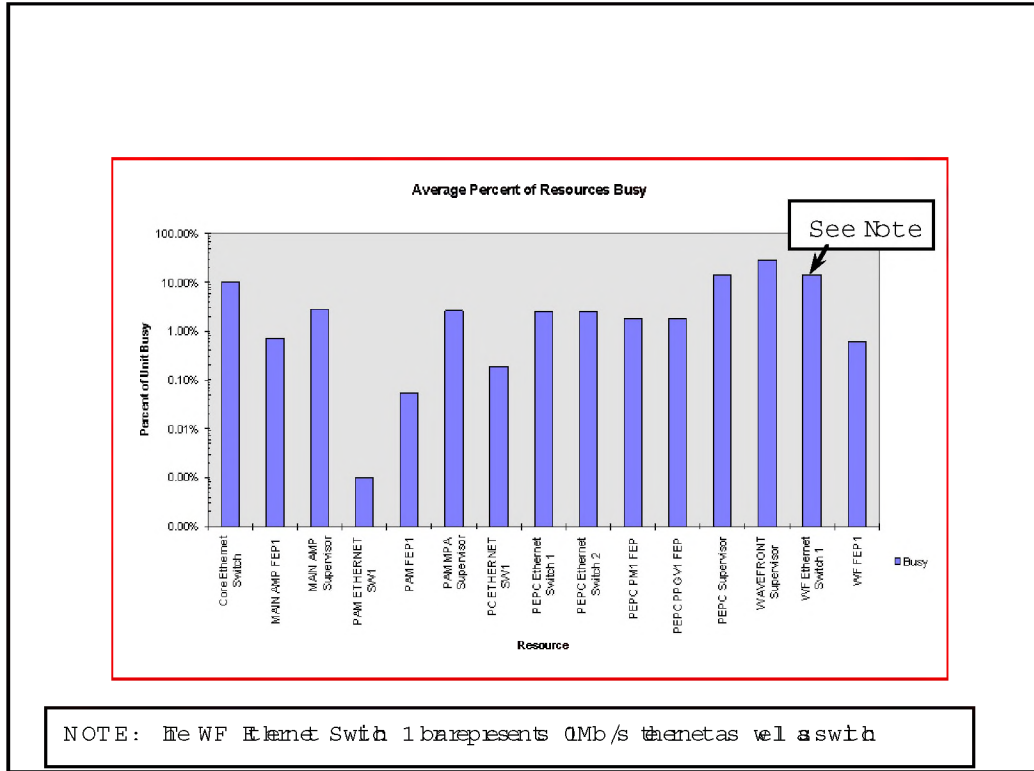


Figure 3-1. Resource utilization vs simulation time for four resources.