

UNIVERSITY OF
ILLINOIS LIBRARY
AT URBANA-CHAMPAIGN
ENGINEERING


NOTICE: Return or renew all Library Materials! The Minimum Fee for each Lost Book is \$50.00.

JUL 06 1988

The person charging this material is responsible for its return to the library from which it was withdrawn on or before the Latest Date stamped below.

Theft, mutilation, and underlining of books are reasons for disciplinary action and may result in dismissal from the University.
To renew call Telephone Center, 333-8400

UNIVERSITY OF ILLINOIS LIBRARY AT URBANA-CHAMPAIGN



Digitized by the Internet Archive
in 2012 with funding from
University of Illinois Urbana-Champaign

<http://archive.org/details/networkingresear0univ>

284
32
240
ENGINEERING LIBRARY
UNIVERSITY OF ILLINOIS
URBANA, ILLINOIS

Engrin

CONFERENCE ROOM

Center for Advanced Computation

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
URBANA ILLINOIS 61801

CAC Document Number 240
CCTC-WAD Document Number 7518

Networking Research in Front Ending
and Intelligent Terminals

ENFE Final Report

September 30, 1977

The Library of the

MAY 23 1978

The person charging this material is responsible for its return to the library from which it was withdrawn on or before the **Latest Date** stamped below.

Theft, mutilation, and underlining of books are reasons for disciplinary action and may result in dismissal from the University.

UNIVERSITY OF ILLINOIS LIBRARY AT URBANA-CHAMPAIGN

ENGINEERING

JUN 6 1978

FEB 21 1982

FEB 23 REC'D

CAC Document Number 240
CCTC-WAD Document Number 7518

Networking Research in Front Ending
and Intelligent Terminals

ENFE FINAL REPORT

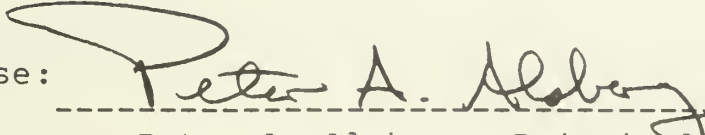
Prepared for the
Command and Control Technical Center
WWMCCS ADP Directorate
Defense Communications Agency
Washington, D.C. 20305

under contract
DCA100-76-C-0088

Center for Advanced Computation
University of Illinois at Urbana-Champaign
Urbana, Illinois 61801

September 30, 1977

Approved for Release:



Peter A. Alsberg, Principal Investigator

TABLE OF CONTENTS

	Page
SUMMARY.....	1
Background.....	1
ENFE Research Program.....	2
ENFE SOFTWARE ARCHITECTURE.....	7
General Description.....	7
Host-to-Front-End Communications.....	7
Channel Protocol.....	7
Channel Protocol module.....	8
Process-to-Service Communications.....	8
Process-to-Service Protocols.....	9
Service Structure.....	9
Host-Host Service module.....	10
Program Access Service module.....	10
Server Virtual Terminal Service module.....	10
PROTOCOL SPECIFICATIONS.....	11
Host-to-Front-End Protocol.....	11
ARPANET Host-Host Process-to-Service Protocol.....	12
Program Access Process-to-Service Protocol.....	12
Server Virtual Terminal Process-to-Service Protocol..	13
Other Process-to-Service Protocols.....	14
Telnet Data Entry Terminal Option.....	14
EXPERIMENTATION.....	16
ENFE Experiment Plan.....	16
Goals.....	16
Tools.....	16
Specific Tests.....	17
UNIX/ENFE Experimental Performance Report.....	17
Experimentation Software.....	17
Experiment Results.....	18
IMPLICATIONS OF EXPERIMENT RESULTS.....	21
Multi-Host Study.....	21
OFFLOADING STRATEGIES.....	41
Offloading the Telnet Protocol.....	41
Offloading the File Transfer Protocol.....	41
Offloading Other ARPANET Protocols.....	42

ALTERNATIVE ARCHITECTURES..... 43

 Alternative Architecture Research Plan..... 43

 State of the Art..... 43

 Research Directions..... 43

 Research Plan..... 43

SUMMARY

Background

Under contract DCA100-76-C-0088, the Center for Advanced Computation of the University of Illinois at Urbana-Champaign has investigated the capabilities of network front ends. As a part of that contract, an experimental network front end (ENFE) has been developed to interface a World-Wide Military Command and Control System (WWMCCS) H6000 to the ARPA network and to conduct experiments with the proposed ARPANET Host-to-Front-End Protocol. A total of 194.28 man-months were expended over a period of 12 months (1 October 1976 to 30 September 1977).

An experimental network front end (ENFE) was a primary contract deliverable. Delay in GFE hardware delivery significantly decreased the work that could have been accomplished. Digital Equipment Corporation delivered the ENFE development min-computer (DEC PDP-11/70) three months late. This slowed construction of critical ENFE operating system software. Associated Computer Consultants delivered the Honeywell-DEC communications link three months late, which delayed software installation, testing, and evaluation. Therefore, only a crude evaluation of ENFE capabilities is available at this time. Despite these delays, however, all of the essential work contracted for has been successfully completed. The ENFE was built, tested, and is currently operational.

ENFE Research Program

The ENFE research program was organized into two teams. The first team implemented the ENFE. The ENFE uses a DEC PDP-11/70 computer running a Unix general purpose operating system. The Unix operating system capabilities were expanded to provide more general support for Host-to-Front-End Protocol (HFP) software. Measurement software was added to Unix to support tests and experiments.

The second team was concerned with protocol issues. This team revised the HFP specifications, generated Telnet protocol options for Honeywell VIP terminals, conducted offloading studies, followed AUTODIN II protocol developments, and constructed a plan for research into alternative front-end architectures.

Individuals and small groups were drawn from both of these teams to generate an experiment plan and to carry out and analyze ENFE tests and experiments.

There were some risks associated with the ENFE research program. The state of the art in network communications is geared to machine/terminal interaction rather than machine/machine interaction. A generalized machine-to-machine protocol like HFP had never been used to facilitate communications between a large computer and a front end. Furthermore, Unix was not designed to support high speed message switching. Given the state of the art in HFP and the architecture of Unix, the Unix ENFE and HFP software were strictly experimental. The primary product of this research program was experience with the

host front-ending problem.

All work (except the multi-host study) performed under the contract has already been thoroughly documented (Tables 1 and 2). Thus, this final report will abstract those reports produced.

Table 1
Contract Deliverables

FAC Document Number	CCTC-WAD Document Number	Title	Date
220	7501	DRAFT H6000 Software Specifications	15 Nov. 1976
221	7502	DRAFT Experimental Network Front End Functional Description	15 Dec. 1976
220	7501	FINAL H6000 Software Specifications	10 Jan. 1977
221	7502	FINAL Experimental Network Front End Functional Description	15 Jan. 1977
227	7509	DRAFT Experimental Network Front End Experiment Plan	28 Mar. 1977
227	7509	FINAL Experimental Network Front End Experiment Plan	16 May 1977
232	7512	DRAFT Alternative Architecture Research Plan	16 May 1977
233	7515	DRAFT Experimental Network Front End Software Functional Description	1 July 1977
233	7515	FINAL Experimental Network Front End Software Functional Description	1 Aug. 1977

ENFE Final Report 9/30/77

CAC Document Number	CCTC-WAD Document Number	Title	Date
232	7512	FINAL Alternative Architecture Research Plan	30 Sept. 1977
239	7517	UNIX/ENFE Experimen- tal Performance Report	30 Sept. 1977
230	7511	Offloading ARPANET Protocols to a Front End	30 Sept. 1977
241	7519	ENFE Nassi-Shneider- mann Flow Charts	30 Sept. 1977
242	7520	ENFE Listings and Object Code	30 Sept. 1977
240	7518	ENFE Final Report	30 Sept. 1977

Table 2
Unscheduled Reports Delivered

AC Technical Number	CCTC-WAD Document Number	Title	Date
219	7503	Host to Front End Protocol/Version I	19 Aug. 1977
80	7504	ARPANET Host-Host Process-to-Service Protocol Specification	17 Mar. 1977
81	7505	Program Access Proc- ess-to-Service Spec- ification	10 Mar. 1977
82	7506	Server Virtual Ter- minal Process-to Ser- vice Protocol Spec- ification	17 Mar. 1977
84	7507	Illinois Inter-Proc- ess Communication Fa- cility for Unix	1 Apr. 1977
94	7514	Telnet Data Entry Terminal Option	27 June 1977

ENFE SOFTWARE ARCHITECTURE

General Description

The offloaded network software can be thought of as a set of services provided to host (H6000) processes or to users. These services allow the network and the various hosts connected to the network to be conveniently used. A complete functional description of the ENFE software architecture is contained in CAC Document No. 233. The key features are summarized below.

Host-to-Front-End Communications

A basic mechanism must be provided to support communication between host processes and front-end services. This mechanism is the Host-to-Front-End Protocol (HFP), which is defined in CAC Document 219 (ARPA Request for Comments (RFC) 710). The HFP specification distinguishes two protocol layers: the channel protocol and the process-to-service protocols.

Channel Protocol. By means of the channel protocol, logical channels are set up between host processes and the front-end services, and messages are transmitted on these channels. Provisions are made for flow control and for out-of-sequence signaling. The channel protocol defines five types of HFP Messages:

1. BEGIN, which sets up logical channels;
2. END, which terminates logical channels;
3. TRANSMIT, which transmits data;
4. SIGNAL, which provides a means for synchronizing the ends of a logical channel, for interrupting the other end, and for flushing data from the other end of the channel; and
5. EXECUTE, which provides a means for passing service-specific information "out of band;" i.e., outside the strict sequencing required for TRANSMIT Messages.

Each Message type can be either a Command (requesting that the action defined by the Message be taken) or a Response (indicating whether the action was taken and, if not, providing some explanation). The HFP specifications use the capitalized word, Message, to refer to these Message types.

Channel Protocol module. The front end contains a software module, the Channel Protocol module (CPM), which manages the logical channels and serves as a bi-directional multiplexor. The host also contains a CPM which similarly manages the other ends of the logical channels.

Process-to-Service Communications

Communications between a host process and a front-end service may be divided into three stages:

1. communications between the host process and the host CPM (described in CSC Document No.

R493700056-2-1, "Host to Front-End Processor Protocol Interface Functional Description"),

2. communications between the host CPM and the front-end CPM (described in CAC Document No. 220, "H6000 Software Specifications" and CAC Document No. 219, "Host-to-Front-End Protocol"), and
3. communications between the front-end CPM and a front-end service (described in CAC Document No. 233, "Experimental Network Front End Software Functional Description").

Process-to-Service Protocols. The process-to-service protocols specify the content, sequencing, and type of HFP Messages by which host processes communicate with front-end services. The process-to-service protocols implemented in the ENFE are:

1. ARPANET Host-Host Process-to-Service Protocol (CAC Technical Memorandum No. 80),
2. Program Access Process-to-Service Protocol (CAC Technical Memorandum No. 81), and
3. Server Virtual Terminal Process-to-Service Protocol (CAC Technical Memorandum No. 82).

Service Structure. Each front-end service implements one process-to-service protocol. All front-end services execute within their own address spaces; i.e., as user-level programs.

Each program is structured as a finite state machine that accepts two types of inputs. HFP Message inputs are generated by processes in the host requesting action from the front-end services. I/O completion event inputs are generated by the system

in response to service-initiated device I/O operations. Each input is associated with a specific HFP logical channel. The input type and current channel state determine the immediate action and next channel state. Most actions result in the transmission of data to another destination and in the generation of an HFP Response indicating the success or failure of the action.

Host-Host Service module. The ARPANET Host-Host Service (HHS) module enables programs running in the host to use the ARPANET Network Control Program (NCP) in the front end. It implements the ARPANET Host-Host process-to-service protocol.

Program Access Service module. The Program Access Service (PAS) module enables programs running in the host to execute arbitrary programs in the front end. It implements the Program Access process-to-service protocol.

Server Virtual Terminal Service module. The ARPANET Server Virtual Terminal Service (SVTS) module enables programs on the host to be accessed by terminals connected to other hosts on the ARPANET. It implements the ARPANET Server Virtual Terminal process-to-service protocol. It also implements the ARPANET Telnet protocol described in NIC Document No. 15372.

PROTOCOL SPECIFICATIONS

Host-to-Front-End Protocol

The performance of data communications tasks such as terminal handling and network protocol interpretation can impose a significant load on a host. Some of these tasks can be performed by the front end for the host. The Host-to-Front-End Protocol (HFP) defines a form of communication between the host and the front end to enable this "offloading" of services.

Thus, the HFP provides specifications for:

1. a channel protocol,
2. individual Commands and Responses,
3. the process-to-CPM interface, and
4. the service-to-CPM interface.

In addition, the HFP provides specifications for specifying process-to-service protocols.

Each HFP Message contains a HEADER carrying channel protocol information and may contain TEXT carrying process-to-service protocol information. Process-to-service protocols use HFP Messages to carry information between a process and a service module. The HFP Message types are:

1. BEGIN Command/Response,
2. END Command/Response,

3. TRANSMIT Command/Response,
4. SIGNAL Command/Response, and
5. EXECUTE Command/Response.

ARPANET Host-Host Process-to-Service Protocol

CAC Technical Memorandum No. 80 specifies a process-to-service protocol for providing ARPANET Host-Host Protocol and Initial Connection Protocol services to a process through the HFP. The Host-Host Protocol is the basic inter-process communication protocol for the ARPANET (ARPANET NIC Document 8246). The program which implements it in each host is the Network Control Program (NCP). The service described here provides an interface, through the HFP, between a process in a host and an NCP in a front end. This enables the host process to establish and use ARPANET connections.

Program Access Process-to-Service Protocol

CAC Technical Memorandum No. 81 specifies a process-to-service protocol for the execution of, or attachment to, arbitrary programs in the front end. The intent was to provide a general mechanism that would allow the host to access terminal-oriented front-end services. Examples of such services are User Mailnet and teleconferencing.

The protocol assumes that the program access service

itself is completely offloaded to the front end. The only software remaining in the host is a relay process that passes properly formatted data between a host terminal or process and the Program Access Service module in the front end. HFP TRANSMIT Commands are used for this data transmission.

Server Virtual Terminal Process-to-Service Protocol

CAC Technical Memorandum No. 82 specifies a protocol for offloading the server side of a virtual terminal protocol; e.g., Server Telnet for the ARPANET. This protocol allows some flexibility in the degree of offloading that may be achieved. Although the protocol is applicable to a general virtual terminal service, the discussion in the specification is in terms of the ARPANET Telnet Protocol, which is currently the only such protocol widely used.

The functions of the typical Server Telnet implementation include:

1. manipulating network connections,
2. negotiating Telnet options,
3. mapping between local terminal representations and network virtual terminal representations,
4. transmitting data over connections,
5. handling special control functions, and

6. interfacing remote terminals so that they appear as if they were local terminals.

The server virtual terminal process-to-service protocol uses HFP Messages to carry information between the residual part of Server Telnet in the host (the "process") and the Server Virtual Terminal Service module (the "service") in the front end.

Other Process-to-Service Protocols

Further study of the problem of offloading Telnet led us to conclude that a single, symmetric protocol should be designed to handle both User and Server Telnet services. We have constructed, but have not implemented, such a protocol (CAC Technical Memorandum No. 103, "Network Virtual Terminal Process-to-service Specification"). This protocol specification, in addition to specifications for three process-to-service protocols constructed in connection with our study of strategies for offloading the ARPANET File Transfer Protocol, is appended to CAC document No. 230, "Offloading ARPANET Protocols to a Front End."

Telnet Data Entry Terminal Option

Under the current contract, we have been tasked to provide facilities for attaching data entry terminals, specifically the 7705 VIP terminal, to the ENFE and the Telnet software. However, the Telnet protocol was originally designed to support simple, scroll-mode terminals. To get the maximum amount of usefulness out of a data entry terminal, the Telnet protocol needed to

be extended. Fortunately, the Telnet protocol has a built-in mechanism, the "option negotiation" mechanism, to allow such extension. We have therefore defined an option to support data entry terminals. In effect, this option defines the Network Virtual Data Entry Terminal. This option supports a minimal set of useful functions common to most data entry terminals and also allows a number of highly sophisticated functions to be negotiated. Details of this option may be found in "Data Entry Terminal Option," CAC Technical Memorandum No. 94 (ARPANET RFC 731).

EXPERIMENTATION

Experimental Network Front End Experiment Plan

This document (CAC Document No. 227) described our preliminary plans for the ENFE tests and experiments. The three main sections identified:

1. goals of the experimentation,
2. tools for experimentation, and
3. scenarios for specific experiments.

Goals. The goals of the experiment plan fell into three categories:

1. performance testing,
2. fine-tuning of the system, and
3. investigating benefits to be gained from design changes.

Tests were proposed to determine whether the system works as it is supposed to and is able to provide the front-ending facilities needed in the short term by the WWMCCS community. In particular, throughput and terminal support tests were given high priority.

Tools. In this document we described an optimal set of tools for thoroughly understanding the front end. The set we planned to implement included:

1. timing mechanisms for generating interrupts as specified by the experimenter and for timestamping messages,
2. artificial traffic generators (involving both software and hardware) at the three front-end

interfaces (the interfaces to the host, to the network, and to the terminals), and

3. software to collect data as the experiments are run.

Other tools described in this document such as a simulation program and queueing theory analysis are essential to a follow-on study of how the front-end design may be improved.

Specific Tests. In the last section of the report, we presented a set of scenarios for specific tests. These were necessarily tentative, particularly as to details. We also mapped out a more comprehensive program of experimentation than we expected to be able to carry out under the current contract. Tests and experiments will continue under a follow-on contract (the Phase B work).

ENIX/ENFE Experimental Performance Report

CAC Document No. 239 reports on the results of the Phase A program of testing and experimentation.

Experimentation Software. The front-end tests used a configuration of software modules similar to the standard ENFE configuration, except that local and foreign host processes were simulated by processes resident in the ENFE itself. These processes served as message generators, sending messages to each other via the standard ENFE modules. An extra copy of the Channel Protocol module was also included in the ENFE to interface the "local host" message generator to the front-end's Channel

protocol module.

In order to make accurate timing measurements, a programmable clock was attached to the PDP-11/70. System calls were implemented to enable the experiment software to utilize this clock to get clock readings and to schedule interrupts. Timestamping software was built into the front end at various points. This software inserted clock readings (timestamps) into the texts of messages as they were transmitted through the front end. In this way the progress of a message through the ENFE could be measured to a high degree of accuracy.

Small modifications were made in the standard Unix monitoring facilities to provide for monitoring of kernel-level as well as user-level processes. This allowed a detailed analysis of processor usage.

Experiment Results. A large part of the Phase A testing and evaluation task involved testing the software to make certain that it operates correctly. A certain amount of fine-tuning (more than was anticipated in the Experiment Plan) was also carried out. The Phase A measurements reported here have allowed us to identify which portions of the system need to be made more efficient and to draw broad conclusions regarding the system architecture.

The measurements reported here include:

1. timing tests of the Inter-Process Communication (IPC) primitives.
2. timing of the progress of single messages

through the front end,

3. monitoring of processor usage, and
4. saturation throughput using several different configurations.

From timing the IPC primitives, we found that it requires a minimum of five to six milliseconds to relay a message from one front-end process to another. The single-message timing measurements indicate that (except in the case of the NCP) the time a message spends being processed by the modules is a small fraction of the time required to relay the message between modules. Monitoring the processor usage by the Channel Protocol module and by the Host-Host Service shows that 85 to 90 percent of CPU time is expended in system calls. Furthermore, kernel-level monitoring shows that about half of this time is just in the overhead of making system calls. We conclude that, as long as the front-end architecture requires Unix system calls to relay messages from one module to another, no dramatic improvement in the efficiency of the front-end services can be expected.

Obtaining meaningful throughput measurements has been made difficult by the self-contained experimentation configuration, which included two passages through the NCP. In this configuration, we found that throughput is severely limited by the CP. To investigate the extent of this limitation, we have sent messages as fast as possible from the ENFE through the Urbana IMP to an 11/50 at Urbana. With this configuration, each message is handled only once by the ENFE NCP. The maximum throughput

measured to date with this configuration is about 50 kilobaud (already enough to saturate the ARPANET) when messages are sent by the 11/50 to the ENFE. However, when messages are sent from the ENFE the throughput is roughly 40 percent less. We attribute this difference to inability of the slower 11/50 to receive data rapidly.

To further investigate the factors affecting throughput, we have separately exercised the two major portions of the message path in the self-contained configuration. The front-end portion of the path (from the message generator to the Host-Host Service) has a message throughput that is four to five times greater than that of the network portion (from a message generator through the NCP to the IMP and back through the NCP to a message receiver). These results provide corroboration of the limiting effect of the NCP.

IMPLICATIONS OF EXPERIMENT RESULTS

Multi-Host Study

In this study we examine the impact of a multi-host configuration on a network front end (NFE). This impact is a function of both the capacity and the structure of the NFE. In this study, the NFE examined is the WWMCCS Phase A Experimental NFE (ENFE). Our estimates show that the WWMCCS Phase A ENFE can support a multi-host configuration. Minor protocol changes may be required.

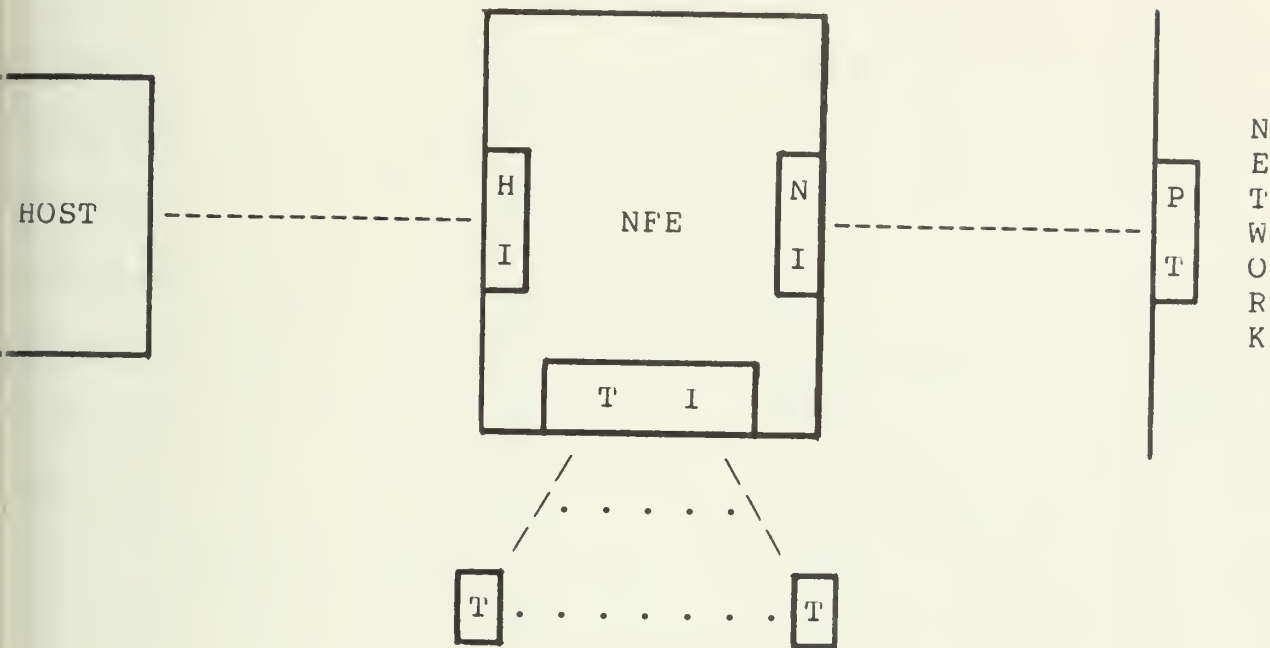
Figure 1 shows a model NFE in a single-host configuration. The NFE communicates with the host via a host interface. In the NFE this is the Asynchronous Bit Serial Interface (ABSI). The NFE communicates with the network via a network interface. In the ENFE this is the IMP-11A. The NFE communicates with terminals via terminal interfaces. In the ENFE, these are DH-11 and V-11 interfaces.

Figure 2 shows an NFE in a multi-host configuration. The difference between this and the single-host configuration is that in the multi-host configuration there are additional host interfaces (ABSI's in the ENFE).

We discuss the quantitative and qualitative effects of multi-host configurations on the NFE. We first deal with the quantitative effects. In so doing, we develop a method for determining the resources required in the NFE for a given load

imposed by the many hosts. We then deal with the qualitative effects on the NFE and the protocols that it uses.

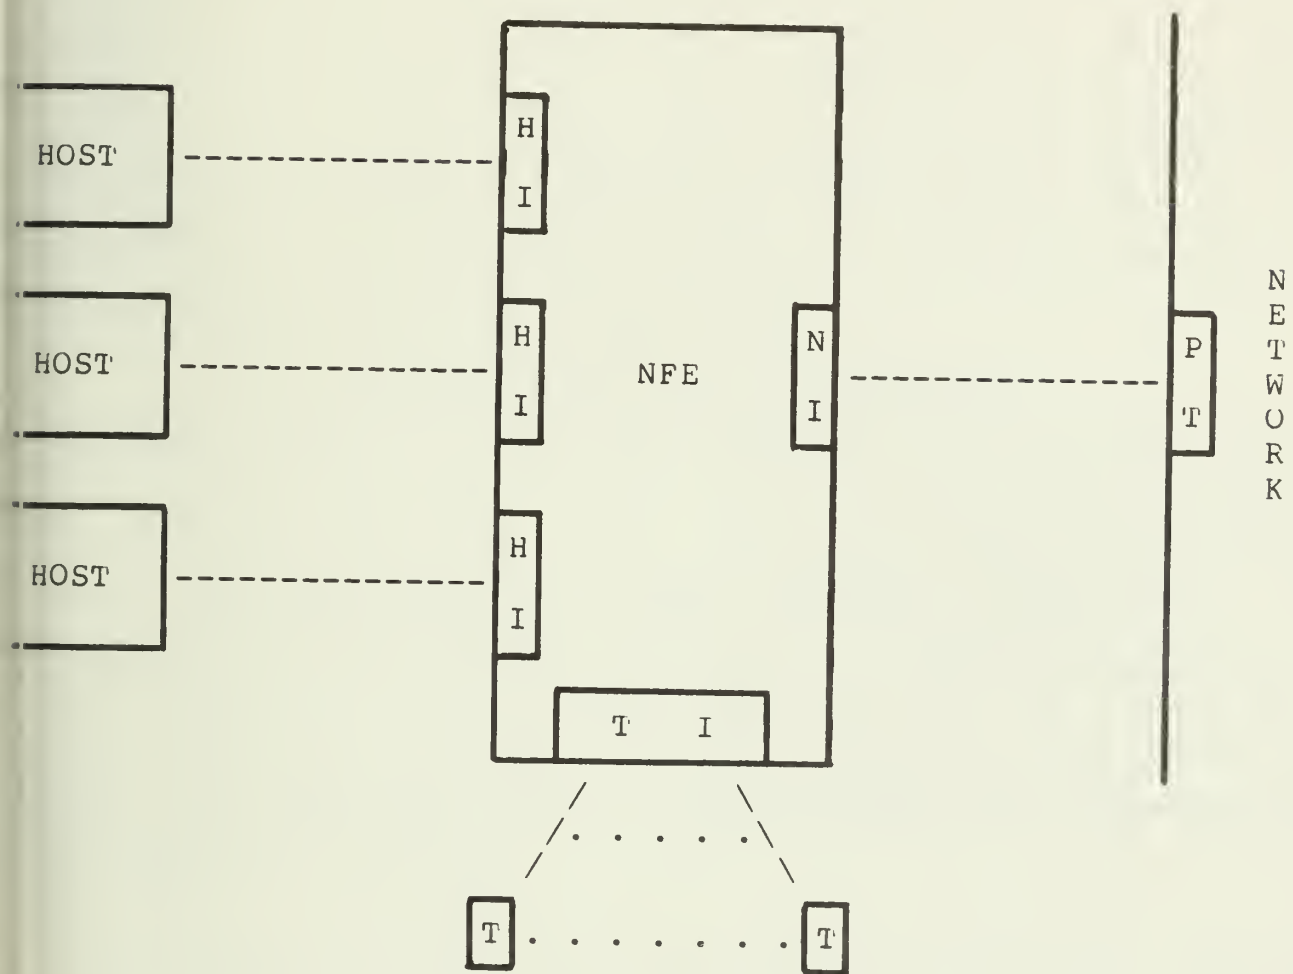
Single-Host NFE Configuration



- were:
- HI = Interface to the host
 - NI = Interface to the network
 - TI = Interfaces to terminals
 - PT = Port
 - T = Terminal

Figure 1.

Multi-Host NFE Configuration



were:

- HI = Interface to the host
- NI = Interface to the network
- TI = Interfaces to terminals
- PT = Port
- T = Terminal

Figure 2.

When we ask whether an NFE can support a given multi-host configuration, we are asking, in part, whether the resources available in the NFE are sufficient to support the load imposed by the hosts. Put another way, if R_{limit} is the amount of a given resource R which is available in the NFE and R_{req} is the total amount of the resource R which is required to support the load imposed by the hosts, we must have:

$$R_{req} \leq R_{limit}, \text{ for all resources } R.$$

To determine whether or not this condition holds, we must compute R_{req} . To do this, we break the NFE down into its component modules, and determine the resources required by each module. We then analyze the load in terms of the use of NFE modules which it entails. Finally, we obtain the total resource requirement by summing the resources required by each module for supporting the load imposed by the hosts.

Let $R_{module}[m]$ be the amount of resource R that is required by module m . Then R_{req} is computed by the summation

$$(1) \quad R_{req} = \sum_{m=1}^M R_{module}[m],$$

where M is the number of modules in the NFE. $R_{module}[m]$ will in general consist of two parts: the fixed part, $R_{fixed}[m]$, and the variable part, $R_{var}[m]$. Thus

$$(2) \quad R_{module}[m] = R_{fixed}[m] + R_{var}[m].$$

$R_{fixed}[m]$ is the amount of resource R which is required by module

independently of the load. $Rvar[m]$ is the amount of resource R which is required by module m and which varies with the load.

We can usually determine $Rfixed[m]$ directly from m by measurement. To determine $Rvar[m]$, we first determine the amount of resource R that is required by module m for each unit of load. This we call $Rload[m]$. We then determine the total load that entails the use of module m . This we call $Lmodule[m]$. Then

$$(3) \quad Rvar[m] = Rload[m] * Lmodule[m].$$

$Lmodule[m]$ will be the sum of the loads which are imposed by each of the hosts and which entail the use of module m . We let $L[h,m]$ be the load which is imposed by host h and which entails the use of module m . If H is the number of hosts,

$$(4) \quad Lmodule[m] = \sum_{h=1}^H L[h,m],$$

$$(5) \quad Rvar[m] = Rload[m] * \sum_{h=1}^H L[h,m],$$

$$(6) \quad Rmodule[m] = Rfixed[m] + (Rload[m] * \sum_{h=1}^H L[h,m]), \text{ and}$$

$$(7) \quad R_{req} = \sum_{m=1}^M (R_{fixed}[m] + (R_{load}[m] * \sum_{h=1}^H L[h,m])).$$

We note that the model we have developed employs a linear relationship between load and resource consumption. In the real world, such a relationship does not always hold, particularly with respect to time-limited resources such as processor time.

We now discuss this method in more detail. In an NFE, the resources we consider are of two kinds: space-limited resources and time-limited resources. The space-limited resource we consider is primary memory. The time-limited resources we consider are processor time and interface bandwidth. We will discuss the computation of R_{req} for primary memory in some detail. We will discuss the other two resources only insofar as their treatment differs from that of primary memory.

We first consider the primary memory resource C . Let C_{limit} be the total available memory in the NFE. We want to compute C_{req} , the total amount of memory required to support the load imposed by the hosts. Then we can determine whether

$$C_{req} \leq C_{limit}.$$

If this relation holds, then the memory in the NFE is sufficient to support the load imposed by the hosts.

Consulting equation (4), we see that the first step in computing C_{req} requires the determination of $L[h,m]$ for each host h

nd for each module m. This information can be conveniently represented by a matrix:

$$\begin{matrix} L[1,1] & L[1,2] & . & . & . & L[1,M] \\ L[2,1] & L[2,2] & . & . & . & L[2,M] \\ . & . & & & & . \\ . & . & & & & . \\ . & . & & & & . \\ L[H,1] & L[H,2] & . & . & . & L[H,M] \end{matrix}$$

ach row represents the load which is imposed by a given host. ach column represents the load which entails a given module. We ote that the sum of column m is Lmodule[m].

To determine the L[h,m], we need two sets of information hich involve the services which the NFE performs for the hosts. e let U[h,s] be the number of simultaneous uses of service s by ost h. This information can also be represented by a matrix:

$$\begin{matrix} U[1,1] & U[1,2] & . & . & . & U[1,S] \\ U[2,1] & U[2,2] & . & . & . & U[2,S] \\ . & . & & & & . \\ . & . & & & & . \\ . & . & & & & . \\ U[H,1] & U[H,2] & . & . & . & U[H,S] \end{matrix}$$

ere S is the number of different services performed by the NFE. ach row represents the use of services by a given host. Each olumn represents the use of a given service across all hosts.

We let E[s,m] be the number of unit loads imposed on module m for each use of service s. Again the information can be rpresented by a matrix:

$$\begin{matrix} E[1,1] & E[1,2] & . & . & . & E[1,M] \\ E[2,1] & E[2,2] & . & . & . & E[2,M] \\ . & . & & & & . \\ . & . & & & & . \\ . & . & & & & . \\ E[S,1] & E[S,2] & . & . & . & E[S,M] \end{matrix}$$

Each row represents the load on the modules of the NFE for each use of a given service. Each column represents the load on a given module when each service is used once.

Then

$$(8) \quad L[h,m] = \sum_{s=1}^S U[h,s] * E[s,m].$$

That is,

$$L = UE.$$

Consulting equation (5), we see that the next step in computing Creg requires the determination of Cload[m] for each module m. This is simply the amount of memory required by module m for each instance of whatever module m does. If a copy of module m is required for each instance of whatever module m does, then Cload[m] is the total memory used by each copy of module m. If the copies of module m all share the same reentrant program but each has a separate copy of the data space, then Cload[m] is the size of the data space for a copy of module m. If there is only one copy of module m which allocates table space and buffer space for each instance of whatever module m does, then Cload[m] is the size of the table space and buffer space allocated for each instance of whatever module m does.

Consulting equations (6) and (7), we see that the final steps in computing Creg require the determination of Cfixed[m]

for each module m . This is simply the amount of memory required by module m independent of how many instances there are of whatever module m does. If a copy of module m is required for each instance of whatever module m does, then $C_{fixed}[m]$ is zero. If the copies of module m all share the same reentrant program but each has a separate copy of the data space, then $C_{fixed}[m]$ is the size of the reentrant program. If there is only one copy of module m which allocates table space and buffer space for each instance of whatever module m does, then $C_{fixed}[m]$ is the size of module m when no such table space and buffer space is allocated.

The computation of C_{reg} thus consists of the following 9 steps:

- 1) construction of the matrix U ,
- 2) construction of the matrix E ,
- 3) computation of the matrix $L = UE$,
- 4) computation of $L_{module}[m]$ using equation (4),
- 5) determination of $C_{load}[m]$,
- 6) computation of $C_{var}[m]$ using equation (5),
- 7) determination of $C_{fixed}[m]$,
- 8) computation of $C_{module}[m]$ using equation (6), and
- 9) computation of C_{reg} using equation (7).

We note that steps 1, 2, 3, and 4 are independent of the resource under consideration. Therefore these steps need be performed only once for all resources.

If $C_{limit} \geq C_{reg}$, the total memory in the NFE will be adequate. But if the NFE is implemented on a mini-computer, we may

also have to take into account the limitation on the amount of memory that can be addressed by each module.

We let $Caddr$ be the address limit imposed on each module by the structure of the NFE hardware. Then the condition that must hold is

$$Caddr \geq Cmodule[m], \text{ for all modules } m.$$

This requires no additional computation, since the $Cmodule[m]$ were computed in step 7 above.

We now turn our attention to the time-limited resources: processor time and interface bandwidth.

Let us first consider the processor time resource P . This resource will be treated differently from primary memory in three ways:

- 1) We are more likely to be interested in the fractional utilization of P than in absolute measures in seconds. That is, P_{limit} will be 1 and $P_{load}[m]$, $P_{fixed}[m]$, $P_{module}[m]$, and P_{req} will be expressed in fractions of the available processor time.
- 2) For most modules m , $P_{fixed}[m]$ will be very small or zero. Exceptions will be those modules which implement communication protocols that require continuous activity for maintaining communication, even when there is no data to be sent or received.
- 3) Our simple linear model of the relation between load and resource use may have to be replaced by a more sophisticated model such as queueing theory.

The last resource we consider is the interface bandwidth resource B . This resource will be treated differently from pri-

ary memory in the three ways mentioned above under processor time. Further, interfaces are specialized resources as opposed to universal resources such as memory space and processor time. This has two consequences:

- 1) Use of each interface will probably be limited to a single module.
- 2) The computations for each interface, or kind of interface, will have to be made separately.

We now apply our modeling method to a concrete example using numerical data. We set ourselves the task of computing C_{req} for the ENFE in a hypothetical multi-host configuration. Our purpose is to illustrate the application of the method. Therefore we will oversimplify wherever it suits our convenience.

We employ the 9-step method discussed above.

In step 1, we construct the matrix U . We recall that $U[h,s]$ is the number of simultaneous uses of service s by host h . We must therefore define the services performed by the ENFE. There are five:

- 1) \underline{HH} performs ARPANET Host-Host Protocol interpretation for each of the local hosts.
- 2) \underline{SVTn} serves as intermediary between the local hosts on one hand and remote terminals connected via the ARPANET on the other.
- 3) \underline{SVTt} serves as intermediary between the local hosts on one hand and local terminals attached to the ENFE on the other.
- 4) \underline{UVTh} serves as intermediary between terminals attached to the local hosts on one hand and remote hosts connected via the ARPANET on the other.
- 5) \underline{UVTt} serves as intermediary between local termi-

nals attached to the ENFE on one hand and remote hosts connected via the ARPANET on the other.

e note that the service UVTt does not involve the local hosts. t only involves local terminals attached to the ENFE. To account for the ENFE resources consumed by this service, we introduce a "phantom host" T. Let there be 3 local hosts, and let the use of ENFE services be as shown in Figure 3. This then is the matrix U.

		service				
		HH	SVTn	SVTt	UVTh	UVTt
h	1	10	5	5	10	-
o	2	12	6	7	4	-
s	3	9	3	11	6	-
t	T	-	-	-	-	15

Figure 3. U[h,s]

In step 2, we construct the matrix E. We recall that E[h,s] is the number of unit loads imposed on module m for each use of service s. We must therefore examine the structure of the ENFE to identify the modules in the ENFE, determine what constitutes a unit load for each module, and determine how many unit loads are imposed on each module for each use of each service.

There are 9 modules in the ENFE that are pertinent to this discussion:

- 1) CPM is the Host-to-Front-End Protocol (HFP) channel protocol module. Its unit load is a logical channel.
- 2) HHS is the HFP service module that enables the local hosts to access the ARPANET NCP. Its unit load is a logical channel and the associated du-

plex ARPANET connection.

- 3) NCPD is the portion of the ARPANET NCP which is implemented as a daemon process. Its unit load is a duplex ARPANET connection.
- 4) NCPK is the portion of the ARPANET NCP which is implemented as part of the Unix operating system kernel. Its unit load is a simplex ARPANET connection.
- 5) SVTS is the HFP service module that enables the local hosts to access both remote terminals connected via the ARPANET and local terminals attached to the ENFE. Its unit load is a logical channel and the associated terminal.
- 6) UTH is the program that enables terminals attached to the local hosts to access remote hosts connected to the ARPANET. Its unit load is a terminal and the associated ARPANET connection.
- 7) UTT is the program that enables terminals attached to the ENFE to access remote hosts connected via the ARPANET. Its unit load is a terminal and the associated ARPANET connection.
- 8) TD is the Unix terminal device driver. It enables terminals attached to the ENFE to access other modules in the ENFE. Its unit load is a terminal.
- 9) PA is the HFP service module that enables the local hosts to access programs in the ENFE (such as UTH). Its unit load is a logical channel and the associated program.

Given the functions performed by each of these modules, the matrix E is as shown in Figure 4.

		module								
		CPM	HHS	NCPD	NCPK	SVTS	UTH	UTT	TD	PA
s	HH	1	1	1	2	0	0	0	0	0
e	SVTn	1	0	1	2	1	0	0	0	0
r	SVTt	1	0	0	0	1	0	0	1	0
v	UVTh	1	0	1	2	0	1	0	0	1
i	UVTt	0	0	1	2	0	0	1	1	0
c										
e										

Figure 4. E[s,m]

In step 3, we compute the matrix $L = UE$. This is shown in Figure 5. We now have the load imposed on each module by each host. For example, host 1 imposes 30 unit loads on the CPM.

		module								
		CPM	HHS	NCPD	NCPK	SVTS	UTH	UTT	TD	PA
h	1	30	10	25	50	10	10	0	5	10
o	2	29	12	22	44	13	4	0	7	4
s	3	29	9	18	36	14	6	0	11	6
t	T	0	0	15	30	0	0	15	15	0

Figure 5. $L[h,s]$

In step 4, we compute $L_{module}[m]$ by summing the columns of L . For example, the sum for the CPM column of L is:

$$L_{module}[CPM] = 30 + 29 + 29 + 0 = 88.$$

We now have the total load imposed on each module by all hosts.

In step 5, we determine $C_{load}[m]$ for each module m . For example, by examining the ENFE program listings, we find that for each logical channel, the CPM requires 14 bytes of table space. Therefore, $C_{load}[m]$ for the CPM is 14 bytes.

In step 6, we compute $C_{var}[m]$ as the product of $L_{module}[m]$ and $C_{load}[m]$. For example, for the CPM we have:

$$C_{var}[CPM] = 88 * 14 = 1232 \text{ bytes.}$$

In step 7, we determine $C_{fixed}[m]$. For example, by examining the memory maps of the ENFE modules, we find that the CPM requires 12522 bytes, exclusive of the memory required per logical

channel. We note that the modules NCPK and TD are actually implemented as parts of the Unix operating system. To account for the memory used by the operating system, we introduce another module called UNIX. We include the Cfixed[m] for NCPK and TD in the Cfixed[m] for UNIX.

In step 8, we compute Cmodule[m] as the sum of Cvar[m] and Cfixed[m]. For example, for the CPM we have:

$$Cmodule[CPM] = 1232 + 12522 = 13754 \text{ bytes.}$$

We now have the total memory required by each module under the load that is defined by matrix U.

The results of steps 4 through 8 are shown in Figure 6.

Module m	Lmodule[m]	Cload[m]	Cvar[m]	Cfixed[m]	Cmodule[m]
CPM	88	14	1232	12522	13754
HHS	31	32	992	10982	11974
NCPD	80	58	4640	17864	22504
NCPK	160	202	32320	0	32320
SVTS	37	56	2072	13162	15234
UTH	20	1294	25880	2974	28854
UTT	15	3976	59640	7183	66823
TD	38	192	7296	0	7296
PA	20	32	640	13522	14162
UNIX	0	0	0	71452	71452

Figure 6.

Finally, in step 9, we compute Creg as the sum of the Cmodule[m]. We find that

$$Creg = 284373 \text{ bytes.}$$

For a PDP-11/70

$$Climit = 2^{**22} = 4194304 \text{ bytes.}$$

Hence we have, for the load defined by matrix U,

Creq < Climit.

All of the foregoing deals with the quantitative aspects of deciding whether an NFE can support a multi-host configuration. But there are also qualitative aspects to this question. These qualitative aspects deal with the structure of the NFE hardware and software and with the structure of the protocols that the NFE uses. To facilitate our discussion, we will consider the ENFE and the protocols that it uses.

We first consider the structure of the ENFE hardware. The hardware basis of the ENFE is the PDP-11/70 computer. A multi-host configuration will require the addition of ABSI's to the ENFE hardware. The UNIBUS structure of the PDP-11/70 makes the addition of ABSI's relatively easy.

We next consider the structure of the ENFE software. The software must take into account the existence of multiple ABSI's and of multiple hosts connected through them. There are two effects which must be considered:

- 1) the effect on the Unix operating system, and
- 2) the effect on the CPM and on the service modules.

The existence of multiple ABSI's presents no problem in the Unix operating system. The structure of the I/O system permits multiple devices of the same kind to be driven by a single re-entrant device driver without confusion. It will permit the CPM to determine on which ABSI a given message arrived. Therefore the CPM will be able to determine from which host the message was

received. It will also permit the CPM to direct a message to the proper ABSI, hence to the proper host.

The existence of multiple ABSI's, and of multiple hosts connected through them, may have some effect on the structure of the CPM and of the service modules. This effect can be very small if a minor change is made to the HFP. We first discuss the problem.

The current version of the CPM and of the service modules uses logical channel numbers to identify the separate logical communications channels between the CPM and the service modules. These are the same logical channel numbers that are used in communication between the CPM in the host and the CPM in the ENFE. If there are multiple hosts, and if no change is made in the current policy for assigning logical channel numbers, confusion will result in the ENFE when two or more hosts use the same logical channel number.

One solution is to add a host field to the state tables in the CPM and in the service modules. A host field would also have to be added to all communications between the CPM and the service modules. This host field would be used to distinguish logical channels with the same logical channel numbers but from different hosts. This solution would require substantial alteration of the CPM and the service modules.

Another solution is to allot to each of the hosts a disjoint subset of the logical channel name space. This would require that the CPM check the logical channel number as it receives each


message. This would ensure that the logical channel number in the message matches the host from which it was received. The CPM would also have to use the logical channel number to direct each outgoing message to the proper ABSI. This solution requires relatively minor changes to the CPM. It requires no change at all to the service modules.

We last of all consider the structure of the protocols that the ENFE uses. Two protocols could affect, or could be affected by, a multi-host configuration: the HFP and the ARPANET Host-Host Protocol.

The only change which might be required in the HFP is in the way that logical channel numbers are assigned. Currently the (single) host may attempt to establish a logical channel using any 28-bit number whose high-order bit is zero. In a multi-host configuration, additional high-order bits could be used for identifying which logical channels belong to which hosts. This is not a significant change to the structure of the HFP or to its implementation.

The ARPANET Host-Host Protocol assumes a one-to-one correspondence between network ports, hosts, and NCP's. This means that confusion could result if a multi-host configuration used a single network interface as shown in Figure 2. It might be necessary to have a network interface for each host as shown in Figure 7. This situation should be avoided in the design of future networks such as AUTODIN II.

Multi-Host NFE Configuration
(Multiple Network Interfaces)



where:

HI	=	Interface to the host
NI	=	Interface to the network
TI	=	Interfaces to terminals
PT	=	Port
T	=	Terminal

Figure 7.

OFFLOADING STRATEGIES

In CAC Document No. 230, "Offloading ARPANET Protocols to a Front End," we presented a broad survey of offloading strategies for the most important ARPA Network protocols. This survey should be useful as a basis both for the future design of expanded front ends and for quantitative studies to determine optimal offloading strategies.

Offloading the Telnet Protocol. We discussed in detail the trade-offs involved in different degrees of offloading and provided a brief analysis of the potential for offloading each of the Telnet options. We also discussed which process-to-service protocols were needed to implement the various schemes. The symmetry of this protocol allowed us to develop a new process-to-service protocol (the network virtual terminal PSP) designed to efficiently implement an intermediate level of Telnet offloading. (The maximum offloading strategy adopted for the ENFE was implemented with two separate PSP's - one for the user side and one for the server side).

Offloading the File Transfer Protocol. We identified two major aspects of the File Transfer Protocol (FTP) that were candidates for offloading: the data transfer process and the book-keeping and marker handling required for restarting a transfer that has aborted. Since FTP makes use of the Telnet protocol, the Telnet functions can also be offloaded. Considering only the User FTP, we found that there are eight possible offloading schemes that differ from one another in major ways.

With regard to the offloading of Server FTP, we confined

ourselves to examining the individual FTP Commands, classifying them as to whether they must be handled in the host or can be handled in the front end.

To make specific the schemes for offloading FTP, we designed three new process-to-service protocols: a User FTP PSP, a Server FTP PSP, and a File Access PSP. The latter provides a general facility for transferring files between host and front end.

Offloading Other ARPANET Protocols. Although our major contractual obligation was to study the offloading of FTP (and Telnet as a natural conjunct of this), we also looked briefly at three other protocols: Remote Job Entry (RJE), Teleconferencing, and Network Graphics.

ALTERNATIVE ARCHITECTURES

Alternative Architecture Research Plan

We developed a research plan (CAC Document No. 232) leading to the specification of a network front end (NFE) designed to meet WWMCCS needs through the 1980's. In CAC Document 232, we briefly reviewed the current state of the art as it affects NFE development, identified some promising directions for research, and presented a detailed plan for conducting the research.

State of the Art. We concluded that the NFE must be modular, efficient, and multi-level secure if it is to meet WWMCCS needs. Three groups of systems were considered relevant to NFE development:

1. existing network access systems,
2. secure systems, and
3. high-bandwidth communications systems.

Examination of these three groups revealed that there does not exist, nor will there exist in the near future, a system which will meet WWMCCS needs.

Research Directions. We presented some promising ideas for research which might lead to solutions to NFE problems. Two alternative hardware architectures were presented for solving the bandwidth problem. An alternative software architecture, the Hub System, was presented which may solve the problems of producing a modular, efficient, and multi-level secure system.

Research Plan. We presented a research plan with four phases:

1. the preparation phase,

2. the research phase,
3. the prototype phase, and
4. the specification phase.

The preparation phase will produce a set of technical constraints for the design of the WWMCCS NFE and will select a set of design features to be studied in the research phase. The selection will be performed through mathematical modeling using the technical constraints. The research phase will design and construct a Research NFE that will be used for evaluating architectural concepts through an iterative process of implementation, testing, and measurement. The prototype phase will design and construct a Prototype NFE which will serve as the basis for specifying the WWMCCS NFE. The specification phase will develop the WWMCCS NFE specification.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
REPORT NUMBER CAC Document Number 240 CCTC-WAD Document Number 7518	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
TITLE (and Subtitle) Networking Research in Front Ending - Final Report	5. TYPE OF REPORT & PERIOD COVERED Research	
	6. PERFORMING ORG. REPORT NUMBER CAC #240	
AUTHOR(●)	8. CONTRACT OR GRANT NUMBER(●) DCA100-76-C-0088	
PERFORMING ORGANIZATION NAME AND ADDRESS Center for Advanced Computation University of Illinois at Urbana-Champaign Urbana, Illinois 61801	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
CONTROLLING OFFICE NAME AND ADDRESS Command and Control Technical Center WWMCCS ADP Directorate, 11440 Isaac Newton Sq., N Reston, Virginia 22090	12. REPORT DATE September 30, 1977	
	13. NUMBER OF PAGES 44	
MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report) UNCLASSIFIED	
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
6. DISTRIBUTION STATEMENT (of this Report) Copies may be requested from the address given in (11) above.		
7. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) No restriction on distribution.		
8. SUPPLEMENTARY NOTES None.		
9. KEY WORDS (Continue on reverse side if necessary and identify by block number) Network front end Network protocol		
ABSTRACT (Continue on reverse side if necessary and identify by block number) The CAC has been engaged in an investigation of the benefits to be gained by employing a network front end. A DEC PDP-11/70 was used as front end for connecting a Honeywell 6000 host to the ARPANET. All work (except the multi-host study) performed under the contract has already been thoroughly documented. Thus, this final report abstracts those reports produced.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)



UNIVERSITY OF ILLINOIS-URBANA



3 0112 005438004