Purdue University
Purdue e-Pubs

Department of Computer Science Technical Reports
Department of Computer Science

1986

Cypress Network Project Status Report

Douglas E. Comer
Purdue University, comer@cs.purdue.edu

Thomas Narten

Report Number:
86-574

Comer, Douglas E. and Narten, Thomas, "Cypress Network Project Status Report" (1986). *Department of Computer Science Technical Reports.* Paper 493.
https://docs.lib.purdue.edu/cstech/493

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries.
Please contact epubs@purdue.edu for additional information.

CYPRESS NETWORK PROJECT
STATUS REPORT

Douglas Comer
Thomas Narten

CSD-TR-574
January 1986

Cypress Network Project Status Report

Douglas Comer
Thomas Narten

Computer Science Department
Purdue University
West Lafayette, IN 47907

1. Introduction

Cypress is a leased-line based packet-switched network that forms part of the DARPA Internet. The basic technology around which Cypress is built consists of small multifunction packet switches called *implets* [1]. Point-to-point leased lines interconnect implets, and provide the fundamental communication media. Hosts using Cypress communicate using the DARPA protocol suite, popularly known as TCP/IP, which stands for Transmission Control Protocol and Internet Protocol.

Cypress implets consist of small minicomputers that run a modified version of a conventional operating system. At the lowest level, the implet functions like a store-and-forward packet switch, receiving packets over leased lines, queuing them temporarily, and forwarding them on toward their final destination. At the second level, each implet functions like an internet gateway [2], accepting packets from the local area network and routing them onto the Cypress network or vice versa. At the third level, the implet functions like a host computer capable of executing processes.

An implet is placed at each subscriber's site, where it connects to the rest of the network over serial lines and to the subscriber's machines over a Local Area Network (LAN) At present, all implets connect to an Ethernet. Using a LAN to connect the implet to user's machines keeps the systems loosely coupled, and allows many machines to connect to the implet. Furthermore, the loosely connected architecture permits implet or user hardware to be replaced without affecting the rest of the site's system. One goal of our interconnection strategy is to provide 1/5 the performance of the ARPANET at 1/10 the cost.

2. Prototype Cypress Network

We have built a prototype network, with implets located at the CSNET Coordination and Information Center (CIC), Bolt Beranek and Newman Corporation, Cambridge, Massachusetts, the University of Arizona in Tuscon, Arizona, Digital Equipment Corporation's Western Research Laboratory (DECWRL) in Palo Alto, California, and Purdue University in West Lafayette, Indiana. An additional implet, located at Boston University, will be brought online in a matter of days. Implets consist of Digital Equipment Corporation VAX 11/725 or Microvax II computers, with implets at remote sites connected to an Ethernet at the subscribers site and to implets at Purdue via 9.6Kbps leased lines. Purdue's implets connect to the DARPA Internet (see figure 1).

Implets run the 4.2BSD UNIX† operating systems (or Ultrix, Digital's version of UNIX), with kernel modifications to support Cypress. Kernel changes included modification of device drivers for efficient packet processing and the addition of IP network interface routines. Packet-switching code has been operating in the VAX 11/725 processors running UNIX since October, 1985; Microvax IIs have been operating the Cypress packet-switching code since February, 1986.

In addition to kernel modifications, several support utilities were designed and implemented. A network monitoring program allows us to monitor current implet status and network traffic from any host on the DARPA Internet. The screen oriented monitor program opens a TCP connection with server processes on the remote implet being

†UNIX is a trademark of AT&T Bell Laboratories.

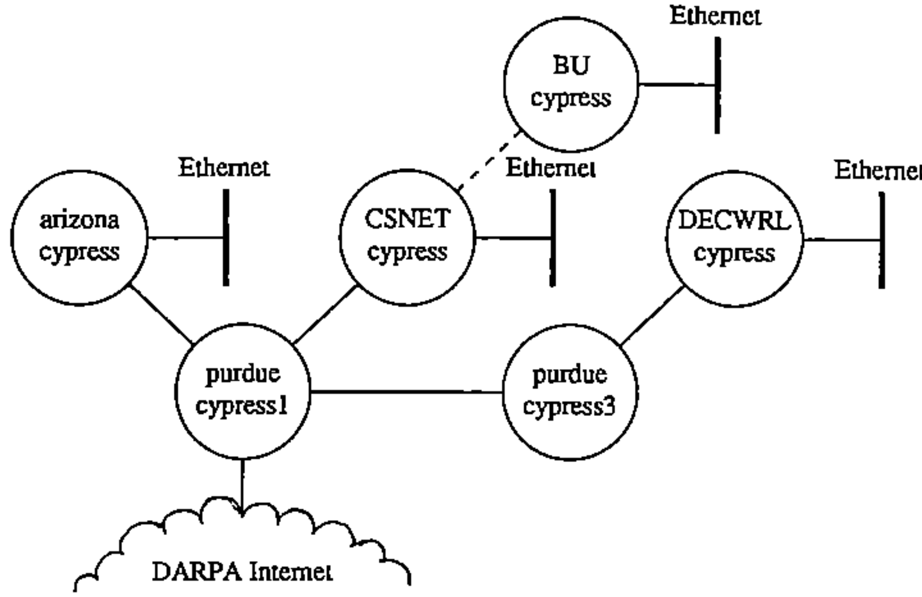

Figure 1: Current configuration of Cypress sites.

monitored. The implet returns Cypress network status information such as queue lengths, numbers and types of packets sent and received, and processor load. The information is continually updated; the monitor always displays information about the most recent state of the implet.

A continuously running error logging daemon saves on disk a record of errors and inconsistencies enountered while the network is running. Records logged include Cypress network errors such as packets that are destined to a site for which no route is known, line errors resulting in packet corruption, and software errors such as buffer allocation failure. A route generator program computes optimal routing tables for Cypress implets. As an implet boots, it opens a connection with a routing server running on the *purdue-cypress1* machine. The server computes optimal routes for that implet using a shortest-path algorithm, where the distance metric of line speed gives a good approximation of real network delays.

3. Performance

One of the most important goals of a host interconnection strategy is good performance. Using the configuration shown in figure 1, we designed experiments to measure the throughput and delay of the network. Throughput, the amount of data that can be successfully transmitted across the network between two processes during a given period of time, was measured using the file transfer utility FTP. The results are shown in table 1. Delay, refering to the amount of time between the sending of the first bit in a message and the receipt at the destination of the last bit in that message, was measured using the utility ping. Ping measures the amount of time required for a packet to travel to and return from a destination host. This so called "round trip time" is computed by measuring the amount of time required for a host to respond to an ICMP [3] echo request. Round trip times are given in table 2.

Looking at the data in table 1, we see measured data rates of approximately 805 bytes/sec. This lies within 96% the optimal value of 842 bytes/sec, which excludes protocol headers [1]. The data rate falls as the number of hops between implets increases due to interference from periodic broadcast packets; the removal from the net of all traffic except that generated by our experiments would result in similar data rates regardless of the hop count between sites. In addition, it should be noted the limiting factor in the throughput is communications line bandwidth; increasing the bandwidth results in a corresponding increase in throughput.

4. Cypress as a Production System

Currently, two of the prototype sites, DECWRL and BB&N do not use Cypress as their primary connection to the Internet. The University of Arizona, however, has been a full fledged members of the DARPA Internet since February, having their local network's address included in the routing tables of ARPANET core gateways. At any

Source	Destination	Bytes	Time (Seconds)	Data Rate (bytes/sec)
Arizona	Cypress1	237568	295	805
Cypress1	Arizona	237568	295	805
Arizona	CSNET	237568	295	805
CSNET	Arizona	237568	295	805
Arizona	Cypress2	237568	296	802
Cypress2	Arizona	238592	298	800
Arizona	DECWRL	237568	302	786
DECWRL	Arizona	237568	299	794

Table 1: Data rates as measured by FTP.

Source	Destination	Data Size	Time (msec)
Cypress1	Arizona	100	392
		200	599
		300	820
Cypress2	Arizona	100	671
		200	1094
		300	1530
DECWRL	Arizona	100	1085
		200	1704
		300	2351

Table 2: Average round trip times as measured by ping.

time, they can communicate with any host in the Internet, and any Internet host can communicate with them. Boston University, will join Cypress shortly. It should be noted that the addition of new sites can now proceed quickly. Administrative obstacles (such as the advertising of routes for behind Cypress networks) have been resolved.

5. Conclusion

We have designed a new addition to the DARPA Internet that uses leased lines to interconnect multifunction packet switches called implets. We have built and measured a prototype Cypress network. The protoype consists of 7 implets located at 5 geographically separated sites. Initial experiments have confirmed that such an interconnection strategy can yield impressive throughput performance, but round trip times can become quite long under heavy load. Further experimentation will be required to determine how long delays affect stream oriented protocols like TCP that use roundtrip times to control retransmission timers. In addition, as more sites begin using Cypress on a permanent basis, we will investigate the reliability, cost effectiveness, and user satisfaction of our system.

References

[1] Comer, D. and Narten, T., "The Cypress Multifunction Packet Switch", CSD-TR-575, Purdue University, Department of Computer Science, January, 1986.

[2] Hinden, R. and Sheltzer, A., "The DARPA Internet Gateway", RFC 823, DARPA Network Information Center, September 1982.

[3] Postel, J. (ed), "Internet Control Message Protocol", RFC 792, DARPA Network Information Center, September 1981.