Porting the Tenet Real-Time Protocol Suite to a HIPPI Network

Bruce A. Mah and Domenico Ferrari
{bmah,ferrari}@CS.Berkeley.EDU

The Tenet Group
University of California at Berkeley
and
International Computer Science Institute

Fourth Gigabit Testbed Workshop
Washington, DC
16 June 1993

Credits

Support

- Corporation for National Research Initiatives
- US Department of Energy

Researchers

- Lawrence Berkeley Laboratory, Imaging Technologies Group
- University of California at Berkeley, RAID Group
- University of California at Berkeley, Tenet Group

Synopsis

The Goals

The Network

Implementation Challenges and Problems

- HIPPI and HIPPI Switches
- Sun 4 and SparcStations
- RAID-II
- Psitech HFB Frame Buffer
- Parallel Processors
- HIPPI-XUNET Adapter (HXA)

Status Report

Goals

Facilities

Provide a high-speed data path between UC Berkeley (Cory and Evans Halls) and Lawrence Berkeley Laboratory (Building 50A)

Research Issues

What kind of real-time guarantees can be supported on a HIPPI network?

What changes need to be made to the Tenet Real-Time Protocol Suite?

Performance of the Tenet Real-Time Protocol Suite at gigabit speeds

Implementation on non-workstation host architectures

Applications

Real-time transfer of video/animation

High-bandwidth transfers to and from RAID-II

The Network

High Performance Parallel Interface (HIPPI)

800 Mbps, 32-bit parallel, point-to-point links

- Optional 64-bit parallel datapath for 1.6 Gbps

Circuit Switching

- No buffers in switches
- Low latency (propagation time)
- Blocking in the network

Physical layer

- Parallel copper (< 25 m)
- Serial fiber or copper using Serial Hippi for longer distances

The Network

Initial Testing Configuration

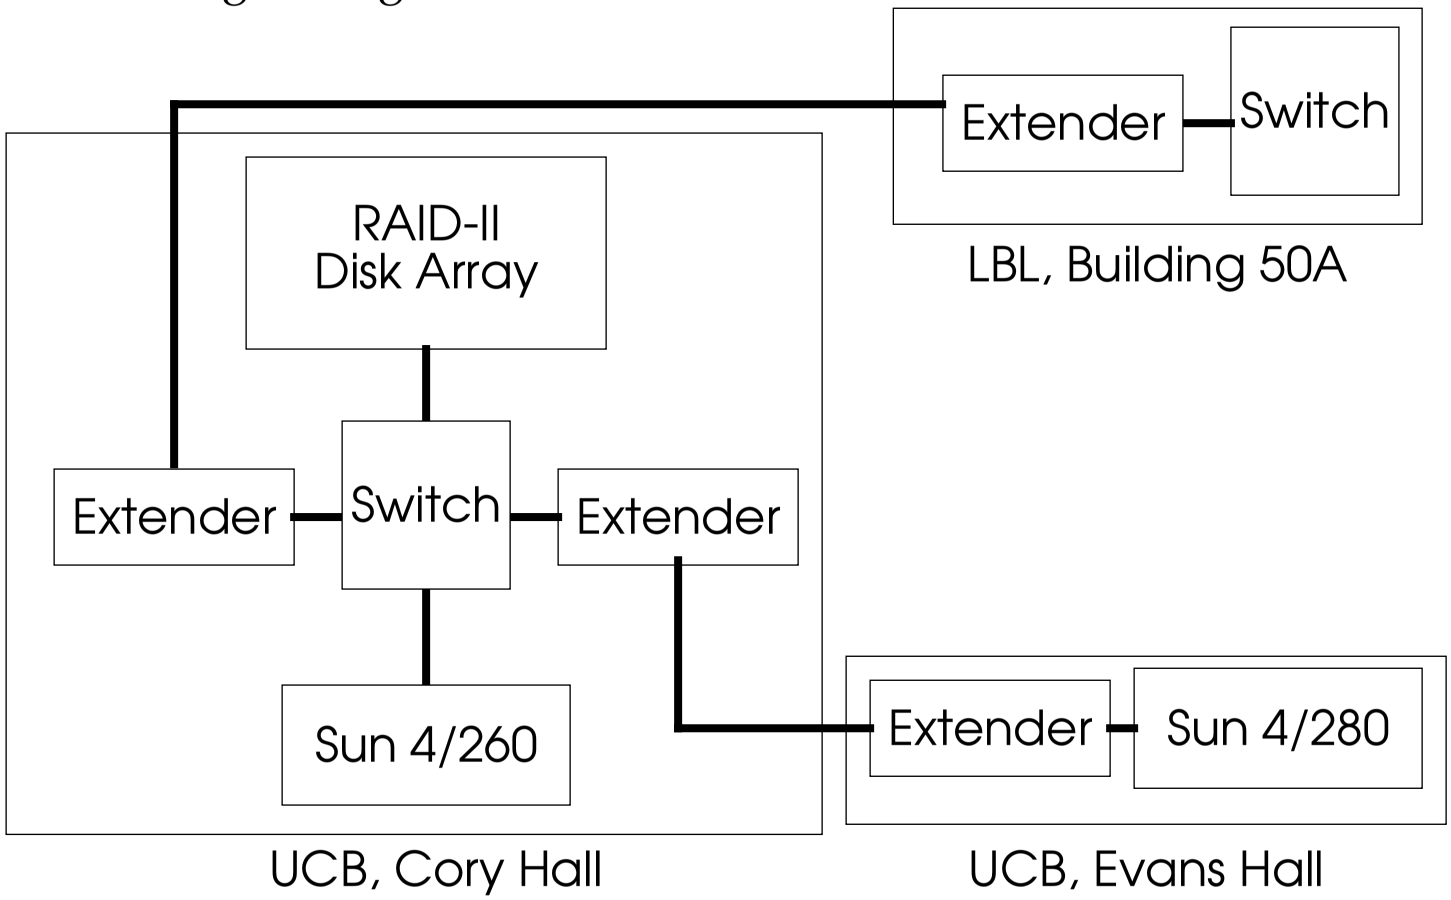

The Network

Hypothetical Configuration

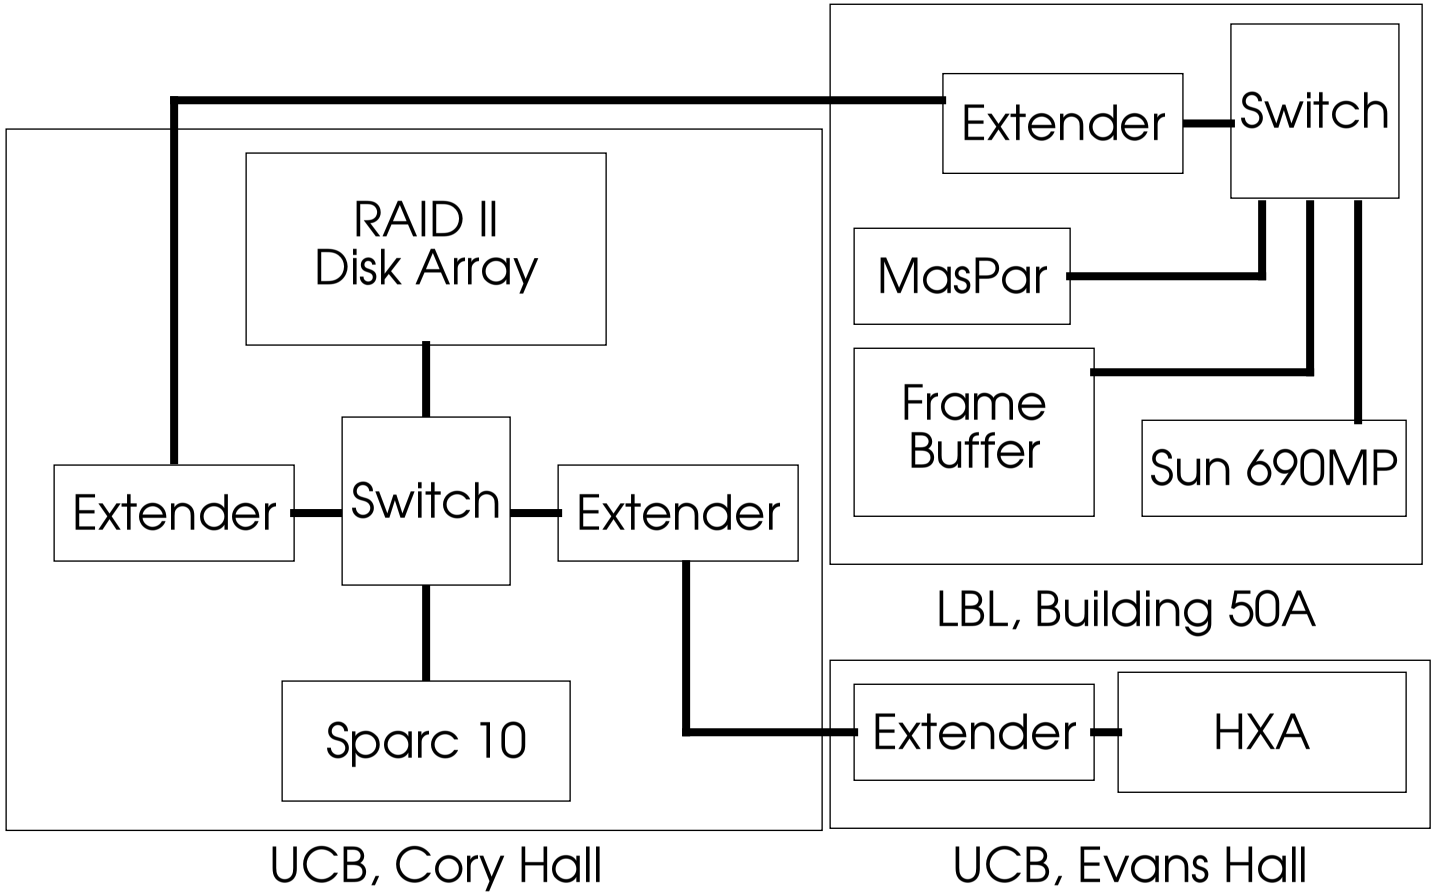

The Network

The Tenet Real-Time Protocol Suite

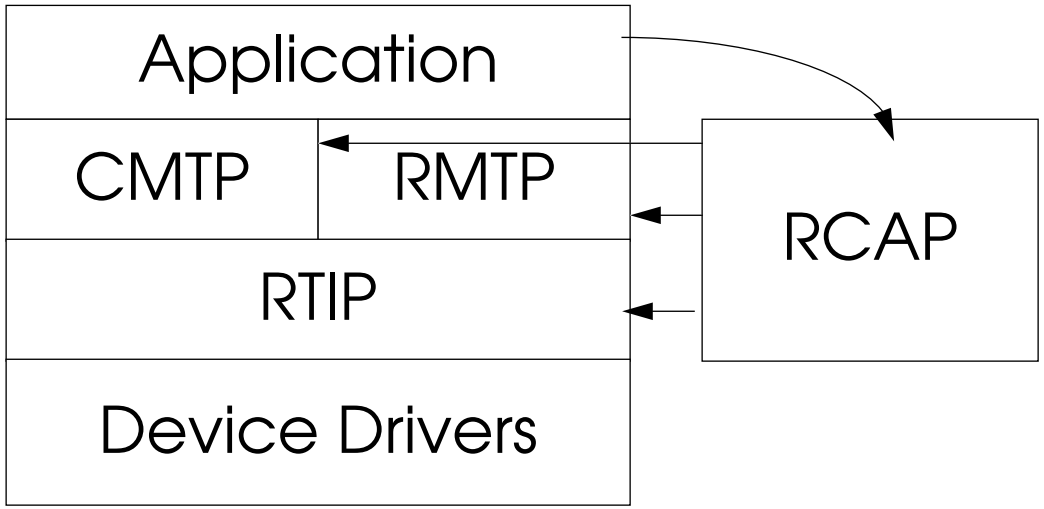

Continuous Media Transport Protocol (CMTP)

Real-Time Message Transport Protocol (RMTP)

Real-Time Internet Protocol (RTIP)

Real-Time Channel Administration Protocol (RCAP)

HIPPI

Circuit-switching technology

Blocking in network switches

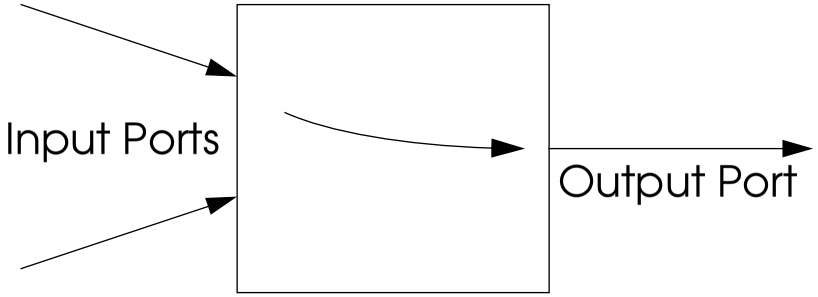

Need to hold all links on path from source to destination

"Camp on" feature allows switch to arbitrate among contending ports

Real-Time Guarantees?

Sun 4s and Sparcstations

"Normal" workstation architectures

Port Tenet Real-Time Protocol Suite to SunOS 4.1.3

- Based heavily on existing prototype Suite on DEC Ultrix 4.2A

Simply A Matter of Programming

RAID-II

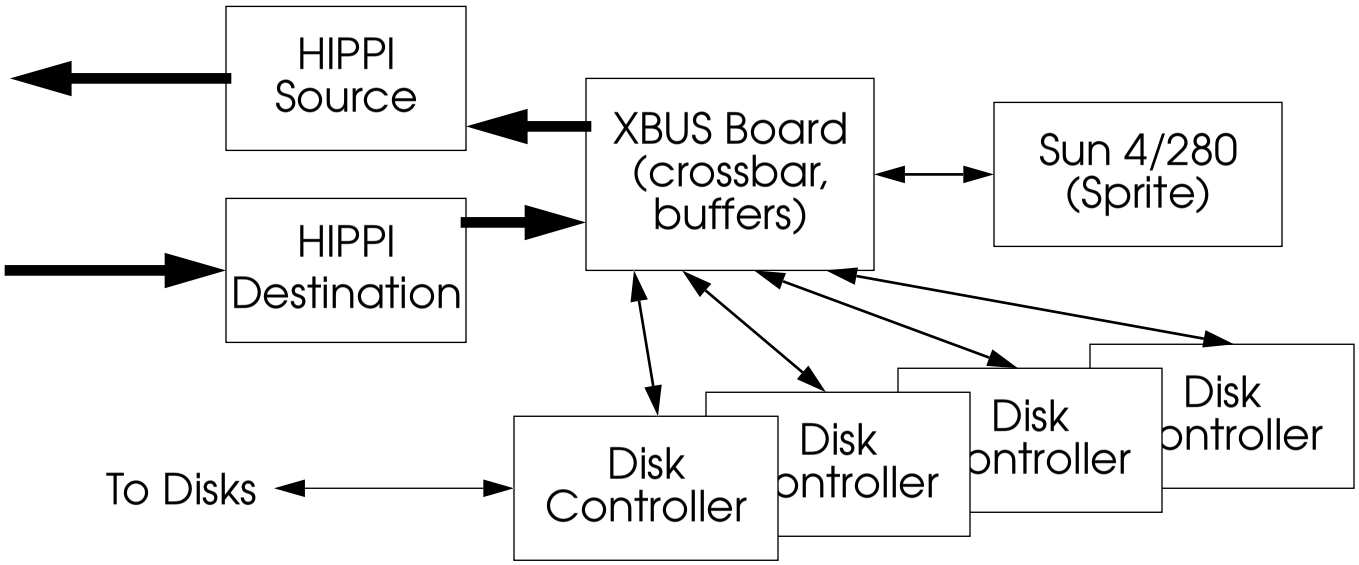

AMD 29000 processors on HIPPI adapters

- Some protocol processing on outboard processors

Network support for filesystem, LFS in Sprite kernel

- Some portion of real-time protocols in kernel
- Interface to Sprite filesystem code (LFS and inter-disk striping driver)

How can we support protocol processing on outboard processors?

PsiTech HIPPI Frame Buffer

HFB-110 supports digital video input and output.

Support for real-time communication to and from frame buffer

Implementation of Real-Time Protocol Suite to run on on-board Sparc processor

- Operating system?
- Datapath bandwidth to and from CPU

How to do protocol implementation in the absence of a traditional operating system?

Parallel Machines

University of California at Berkeley CM-5

Lawrence Berkeley Laboratory MasPar

How to make the Tenet Real-Time Protocol Suite run efficiently on parallel processors?

How can the Tenet Real-Time Protocol Suite best support the kind of network I/O required by parallel machines?

HXA

HIPPI-XUNET Adapter (HXA)

Convert HIPPI frames into ATM cells for transmission and back

XUNET 3 segment will use 622Mbps line card to connect HIPPI segments of network to XUNET 2 ATM switch at UC Berkeley

How well can the Tenet Real-Time Protocol Suite provide performance guarantees in heterogeneous internetworks?

What kind of performance can our data delivery protocols provide in such an internetwork?

Status Report

Network equipment procured

SunOS version of Real-Time Protocol Suite operational

RAID II/Sprite port of Suite under development

- RMTP/RTIP to run primarily on network interfaces
- RCAP on server CPU

"On deck":

- MasPar
- PsiTech HFB