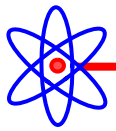
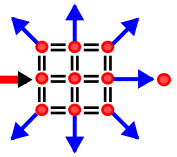

System Architectures Using Network Attached Peripherals

Rodney Van Meter

USC/Information Sciences Institute

rdv@isi.edu

http://www.isi.edu/netstation/

USC Integrated Media Systems Center
Student Council Seminar
October 30, 1997

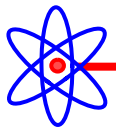

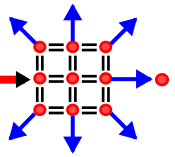

Talk Outline

- Introduction
- Network Technologies
- NAPs in Multimedia
- NAPs in Mass Storage
- Operating System Support
- Conclusion

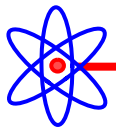

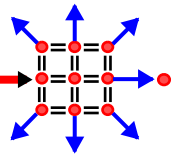

What is a Network Attached Peripheral?

Any computer peripheral attached directly to some form of network, rather than a bus.

- HiPPI frame buffers
- Fibre Channel disk drives
- ATM cameras

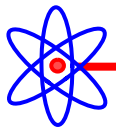

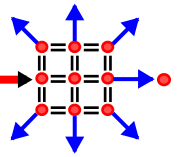

Characteristics of Network-Attached Peripherals (NAPs)

- Scalable physical interconnect
  (# of nodes, distance, etc.)
- No physically defined **owner**
- Interconnect shared w/ general-purpose traffic
- Higher latency
- Delivery subject to usual network problems
  (packet loss, out-of-order delivery, fragmentation, etc.)
- Support for **3rd party transfer**
  (direct device-to-device communication)

Present in varying degrees in different systems.

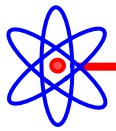

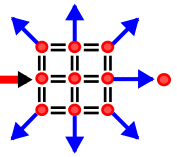

Problems Faced with NAPs

Closed, bus-centric architecture allows simplifying assumptions about resource identification, security and sharing.

- Set of resources not constrained by architecture
- Network issues of scale & heterogeneity
- Control of devices not limited to bus master
- Non-dedicated network
- Security now paramount

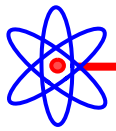

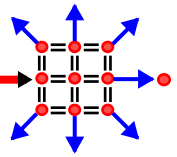

What are NAPs Good for?

- Better scaling (distance, # nodes, aggregate bandwidth)
- Simpler cabling
- Direct device-to-device communication
- Direct device-to-client comm. reduces server load

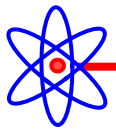

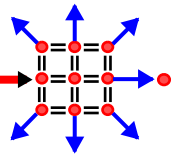

Talk Outline

- Introduction
- **Network Technologies**
- NAPs in Multimedia
- NAPs in Mass Storage
- Operating System Support
- Conclusion

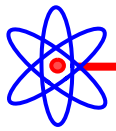

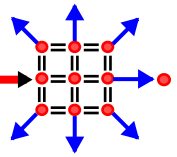

Network Technologies for NAPs

All seven layers in ISO model open to debate

- Application
- Presentation
- Session
- Transport
- Network
- Link
- Physical

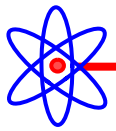

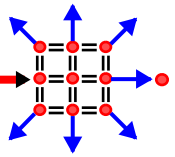

Proposed & In-Use Networks

- HiPPI 800
- HiPPI 6400
- Fibre Channel fabrics
- Fibre Channel Arbitrated Loop
- FireWire (1394)
- Gigabit ethernet
- ATM
- Serial Storage Architecture (SSA)
- Myrinet
- various others

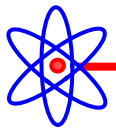

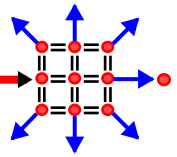

High Performance Parallel Interface (HiPPI)

- Goals: simple & fast (800 Mbps), supercomputing
- Switched or routed
- Parallel copper or serial fiber
- Phy, link layers
- IPI-3 or TCP
- Weaknesses: limited scalability

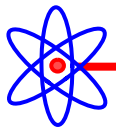
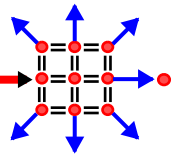

Fibre Channel

- Goals: fast, scalable, distance (ambitious)
- Serial copper coax or fiber
- 800 Mbps
- Switched fabric or arbitrated loop
- Phy, link, net, transport layers
- SCSI commands over custom transport
- Front runner for “winner”
- Weaknesses: expense, complexity;
  scalability and loop/fabric interoperability unproven
  (low pkt loss rate, in-order delivery assumptions may not hold)
- http://www.fibrechannel.org/

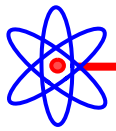
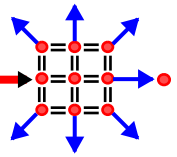

FireWire 1394

- Goals: simplicity, low cost, desktop environment
- Custom copper cables
- 100, 200, 400 Mbps
- Arbitrary physical topology, but shared/broadcast medium
- Phy, link, net, transport layers
- Very bus-like
- Weaknesses: shared low bandwidth; nothing scales
- http://www.firewire.org/

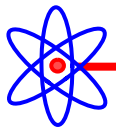

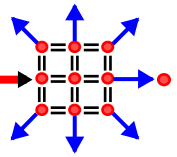

Gigabit Ethernet

- Goals: interoperability w/ ethernet switches, similar programming model
- Tweaked Fibre Channel physical
- 1 Gbps
- Phy, link layers
- Likely popular for GP traffic, can it translate to storage?
- Weaknesses: small packet size, expense, undefined storage profile

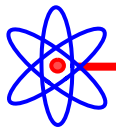

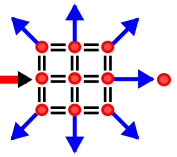

Networking Problems for NAPs

HiPPI-6400❋gigabit Ethernet❋Myrinet❋FC-AL

as I/O Nets Get Larger and More Complex:

- Media Bridging
  (Routing, Addressing)
- Congestion
- Flow Control
- Demultiplexing @ Endpoints
  (Destination Address Calculation, Control/Data Sifting, Upper Layer Protocols)
- Latency Variation
- Security
- Reliability
- Heterogeneity
  (Hosts, Traffic Types, Nets)

All Become Bigger Problems!
But...

1394❋HiPPI-800❋ATM❋SSA❋Fibre Channel

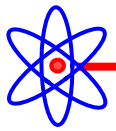

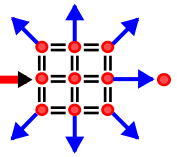

The Internet Community Has Solved Most of the Problems

- Strengths of IP: issues of scale and heterogeneity
- Weakness: Performance
- ISI's Netstation is using & promoting TCP/IP and UDP/IP
- Performance problems can be solved

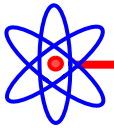

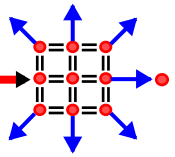

Advantages of IP

- Heterogeneous Interconnects

  Intra-Machine Room
- Wide-Area Access

  Enables Remote Mirroring and Backups
- Future Growth

  Not Media-Specific
- Lower R&D Investment in Networking

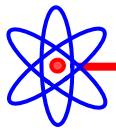

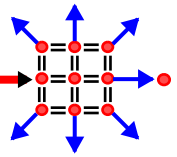

Talk Outline

- Introduction
- Network Technologies
- **NAPs in Multimedia**
- NAPs in Mass Storage
- Operating System Support
- Conclusion

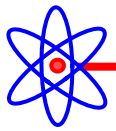

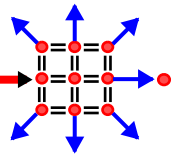

NAPs in Multimedia

Cameras, frame buffers and occasionally disk drives

- ISI’s Netstation
- MIT’s ViewStation
- Cambridge’s Desk Area Network
- HiPPI frame buffers

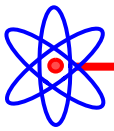

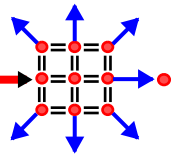

The Netstation Project

Gregory Finn (project leader),
Steve Hotz,
Rodney Van Meter,
Bruce Parham and Reza Rejaie

http://www.isi.edu/netstation/

Technologies for NAPs:

- Networking protocols
- OS paradigms
- NAP security
- Multimedia & storage

Netstation

Netstation is a system composed of **network-attached peripherals (NAPs)** created by replacing the system bus in a workstation with a gigabit network.

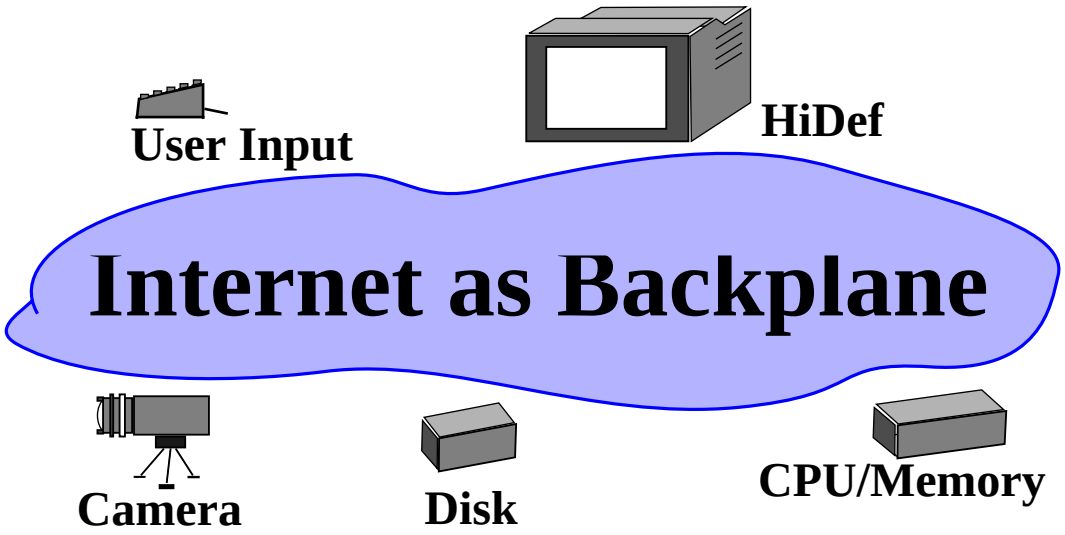

- Use Internet protocols for ubiquitous device access
- Based on ATOMIC 640 Mbps switched network

ViewStation & Desk Area Network

- Principle difference: physically-defined boundary
- ATM

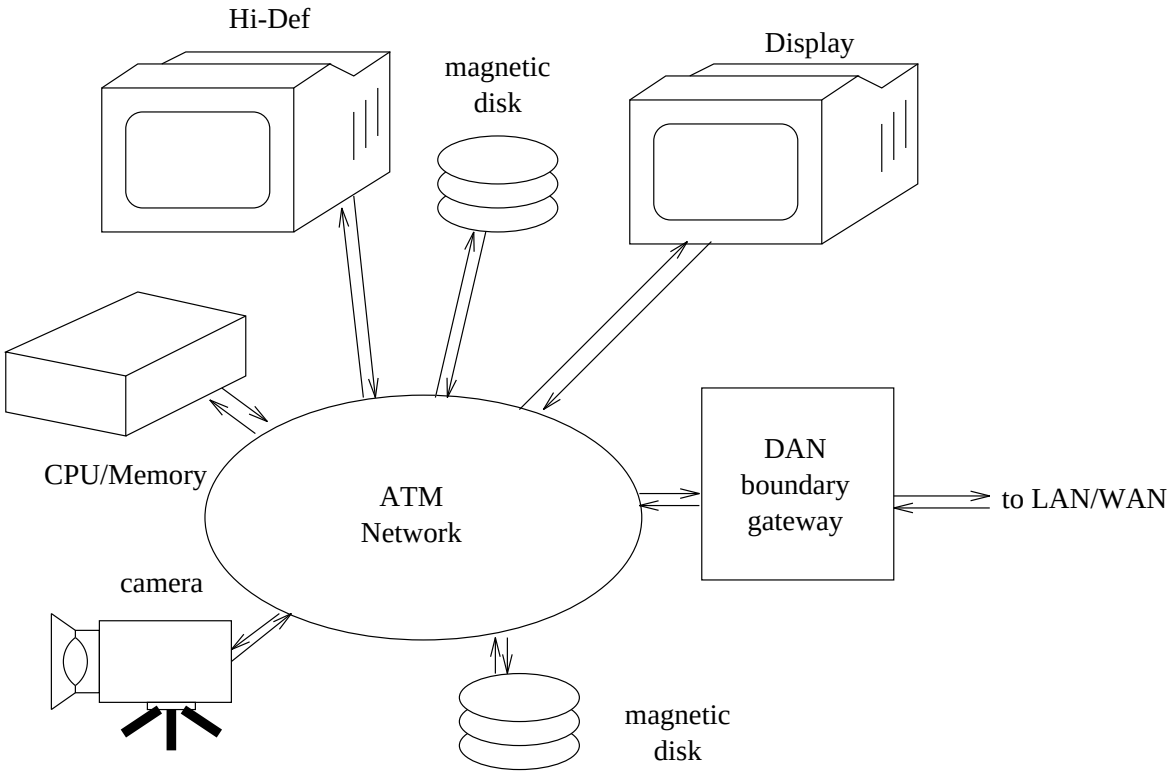

Third Party Transfer

- Direct device-to-device transfer

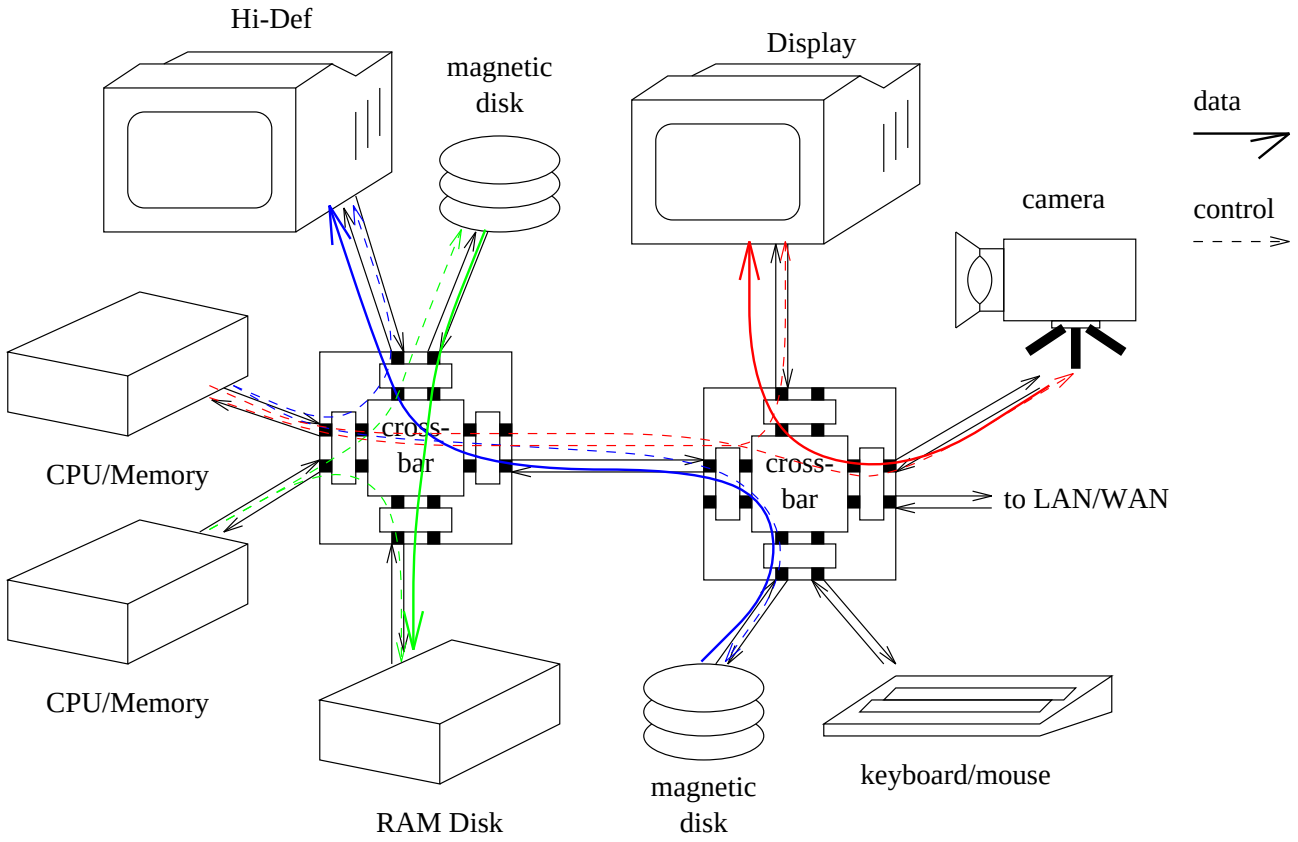

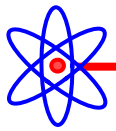

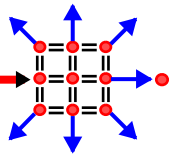

Talk Outline

- Introduction
- Network Technologies
- NAPs in Multimedia
- **NAPs in Mass Storage**
- Operating System Support
- Conclusion

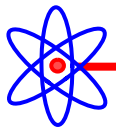

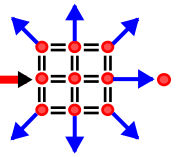

NAPs in Mass Storage

- SGI Origin 2000?
- CMU Network-Attached Secure Disk (NASD)
- LLNL's Network-Attached Peripheral (NAP) RAID
- Fibre Channel Disk Drives
- Palladio at HP Labs
- Petal/Frangipani at DEC
- Global File System from UMinn
- National Storage Industry Consortium's NASD Committee
  http://www.hpl.hp.com/SSP/NASD/

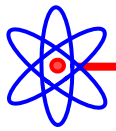
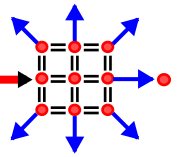

Network Disk Services

Should a drive present a SCSI (block) model,
or NFS (file) model, or something in between?

- Low-level interface easily supports other uses
  (non-Unix file systems, databases, swap space, network RAID)
- File model may distribute functionality more widely,
  scaling better
- Architectural tradeoffs are complex

CMU Network Attached Secure Disk Group

- Defined useful taxonomy
- Their disks hold “objects”, like unnamed NFS files
- File manager/name service centralized
- http://www.pdl.cs.cmu.edu/NASD/

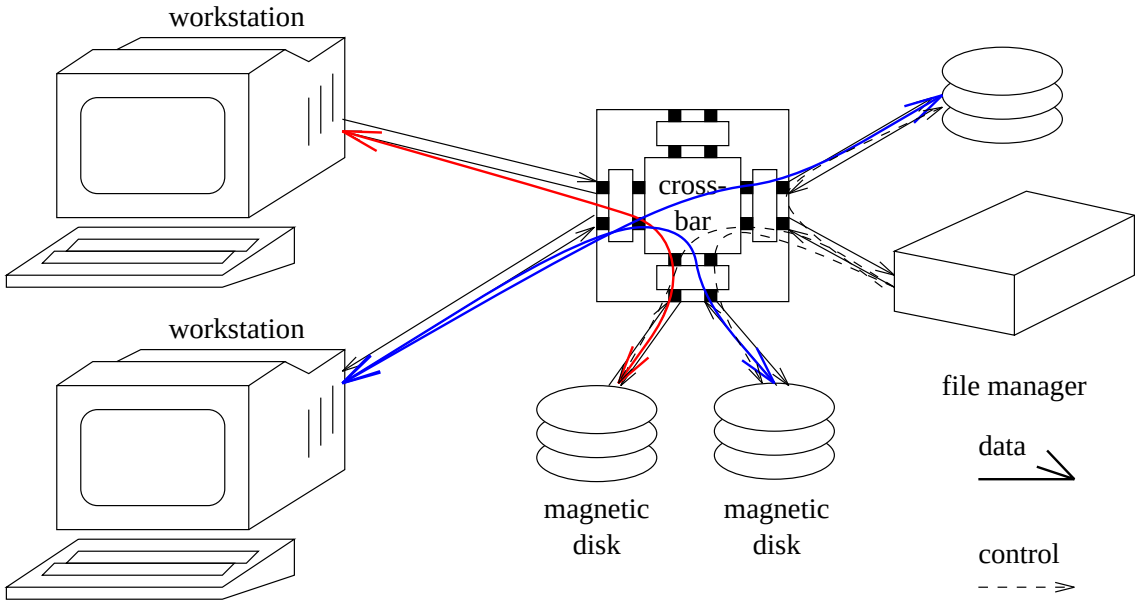

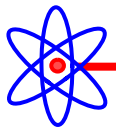

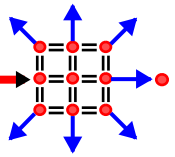

Talk Outline

- Introduction
- Network Technologies
- NAPs in Multimedia
- NAPs in Mass Storage
- **Operating System Support**
- Conclusion

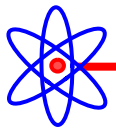

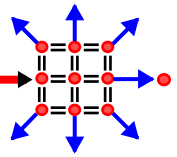

Operating System Issues with NAPs

- Resource discovery
- Concurrency/sharing
- Security
- Programming paradigms for third-party transfer

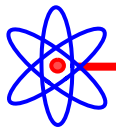
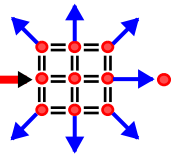

Security

- Access not physically constrained
- Cryptographic authentication required
- **Who** a request comes from is more important than **where**
- Devices don’t understand “users”
- Netstation approach: Derived Virtual Devices (DVDs)

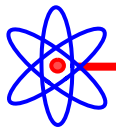

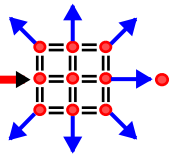

Third-Party Transfer

- **`read/write`** paradigm inadequate -- generalize to **`move(source,destination)`**
- Concurrency management
- Error handling: to partner, requestor or owner of one or both devices?
- Details: boundary conditions, blocking factors, generalized RPC formats

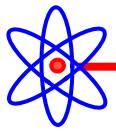

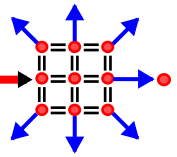

Conclusions

- Network Attached Peripherals (NAPs) allow
  new system architectures
  More scalable interconnects
  Direct device communication
- Key issues:
  Security
  Scale
  Performance
  Legacy
- “A Brief Overview of Current Work on Network Attached Peripherals”, ACM OSR Jan. ‘96 or web page below
- **`http://www.isi.edu/~rdv/`**