

# High-Performance Network and Channel Based Storage

RANDY H. KATZ, SENIOR MEMBER, IEEE

*Invited Paper*

*In the traditional mainframe-centered view of a computer system, storage devices are coupled to the system through complex hardware subsystems called I/O channels. With the dramatic shift toward workstation-based computing, and its associated client/server model of computation, storage facilities are now found attached to file servers and distributed throughout the network. In this paper, we discuss the underlying technology trends that are leading to high-performance network-based storage, namely advances in networks, storage devices, and I/O controller and server architectures. We review several commercial systems and research prototypes that are leading to a new approach to high-performance computing based on network-attached storage.*

## I. INTRODUCTION

The traditional mainframe-centered model of computing can be characterized by small numbers of large-scale mainframe computers, with shared storage devices attached via I/O channel hardware. Today, we are experiencing a major paradigm shift away from centralized mainframes to a distributed model of computation based on workstations and file servers connected via high-performance networks.

What makes this new paradigm possible is the rapid development and acceptance of the client/server model of computation. The client/server model is a message-based protocol in which clients make requests of service providers, which are called servers. Perhaps the most successful application of this concept is the widespread use of file servers in networks of computer workstations and personal computers. Even a high-end workstation has rather limited capabilities for data storage. A distinguished machine on the network, customized either by hardware, software, or both, provides a file service. It

Manuscript received October 1, 1991; revised March 24, 1992. This work was supported by the Defense Advanced Research Projects Agency and the National Aeronautics and Space Administration under contract NAG2-591 (Diskless Supercomputers: High Performance I/O for the TerOp Technology Base). Additional support was provided by the State of California MICRO Program in conjunction with industrial matching support provided by DEC, Emulex, Exabyte, IBM, NCR, and Storage Technology corporations.

The author is with the Computer Science Division, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720.

IEEE Log Number 9203670.

accepts network messages from client machines containing open/close/read/write file requests and processes these, transmitting the requested data back and forth across the network.

This is in contrast to the pure distributed storage model, in which the files are dispersed among the storage on workstations rather than centralized in a server. The advantages of a distributed organization are that resources are placed near where they are needed, leading to better performance, and that the environment can be more autonomous because individual machines continue to perform useful work even in the face of network failures. While this has been the more popular approach over the last few years, there has emerged a growing awareness of the advantages of the centralized view. That is, every user sees the same file system, independent of the machine they are currently using. The view of storage is pervasive and transparent. Further, it is much easier to administer a centralized system, to provide software updates and archival backups. The resulting organization combines distributed processing power with a centralized view of storage.

Admittedly, centralized storage also has its weaknesses. A server or network failure renders the client workstations unusable and the network represents the critical performance bottleneck. A highly tuned remote file system on a 10 megabit (Mbit) per second Ethernet can provide perhaps 500K bytes per second to remote client applications. Sixty 8K byte I/O's per second would fully utilize this bandwidth. Obtaining the right balance of workstations to servers depends on their relative processing power, the amount of memory dedicated to file caches on workstations and servers, the available network bandwidth, and the I/O bandwidth of the server. It is interesting to note that today's servers are not I/O limited: the Ethernet bandwidth can be fully utilized by the I/O bandwidth of only two magnetic disks!

Meanwhile, other technology developments in processors, networks, and storage systems are affecting the relationship between clients and servers. It is well known that processor performance, as measured in MIPS ratings,

0018-9219/92\$03.00 © 1992 IEEE

is increasing at an astonishing rate, doubling on the order of once every 18 months to two years. The newest generation of RISC processors has performance in the 50 to 60 MIPS range. For example, a recent workstation announced by the Hewlett-Packard Corporation, the HP 9000/730, has been rated at 72 SPECMarks (1 SPECMark is roughly the processing power of a single Digital Equipment Corporation VAX 11/780 on a particular benchmark set). Powerful shared memory multiprocessor systems, now available from companies such as Silicon Graphics and Solborne, provide well over 100 MIPS performance. One of Amdahl's famous laws equated one MIPS of processing power with one megabit of I/O per second. Obviously such processing rates far exceed anything that can be delivered by existing server, network, or storage architectures.

Unlike processor power, network technology evolves at a slower rate, but when it advances, it does so in order of magnitude steps. In the last decade we have advanced from 3 Mbit/second Ethernet to 10 Mbit/second Ethernet. We are now on the verge of a new generation of network technology, based on fiber-optic interconnect, called FDDI. This technology promises 100 Mbits per second, and at least initially, it will move the server bottleneck from the network to the server CPU or its storage system. With more powerful processors available on the horizon, the performance challenge is very likely to be in the storage system, where a typical magnetic disk can service 30 8K byte I/O's per second and can sustain a data rate in the range of 1 to 3 Mbytes per second. And even faster networks and interconnects, in the gigabit range, are now commercially available and will become more widespread as their costs begin to drop [1].

To keep up with the advances in processors and networks, storage systems are also experiencing rapid improvements. Magnetic disks have been doubling in storage capacity once every three years. As disk form factors shrink from 14 inch to 3.5 inch and below, the disks can be made to spin faster, thus increasing the sequential transfer rate. Unfortunately, the random I/O rate is improving only very slowly, owing to mechanically limited positioning delays. Since I/O and data rates are primarily disk actuator limited, a new storage system approach called disk arrays addresses this problem by replacing a small number of large-format disks by a very large number of small-format disks. Disk arrays maintain the high capacity of the storage system, while enormously increasing the system's disk actuators and thus the aggregate I/O and data rate.

The confluence of developments in processors, networks, and storage offers the possibility of extending the client/server model so effectively used in workstation environments to higher performance environments, which integrate supercomputer, near supercomputers, workstations, and storage services on a very high performance network. The technology is rapidly reaching the point where it is possible to think in terms of diskless supercomputers in much the same way as we think about diskless workstations. Thus, the network is emerging as the future "backplane" of high-performance systems. The challenge is to develop

the new hardware and software architectures that will be suitable for this world of network-based storage.

The emphasis of this paper is on the integration of storage and network services, and the challenges of managing the complex storage hierarchy of the future: file caches, on-line disk storage, near-line data libraries, and off-line archives. We specifically ignore existing mainframe I/O architectures, as these are well described elsewhere (for example, in [2]). The rest of this paper is organized as follows. In the next three sections, we will review the recent advances in interconnect, storage devices, and distributed software, to better understand the underlying changes in network, storage, and software technologies. Section V contains detailed case studies of commercially available high-performance networks, storage servers, and file servers, as well as a prototype high-performance network-attached I/O controller being developed at the University of California, Berkeley. Our summary, conclusions, and suggestions for future research are found in Section VI.

## II. INTERCONNECT TRENDS

### A. Networks, Channels, and Backplanes

Interconnect is a generic term for the "glue" that interfaces the components of a computer system. Interconnect consists of high-speed hardware interfaces and the associated logical protocols. The former consists of physical wires or control registers. The latter may be interpreted by either hardware or software. From the viewpoint of the storage system, interconnect can be classified as high-speed networks, processor-to-storage channels, or system backplanes that provide ports to a memory system through direct memory access techniques.

Networks, channels, and backplanes differ in terms of the interconnection distances they can support, the bandwidth and latencies they can achieve, and the fundamental assumptions about the inherent unreliability of data transmission. While no statement we can make is universally true, in general, backplanes can be characterized by parallel wide data paths and centralized arbitration, and are oriented toward read/write "memory mapped" operations. That is, access to control registers is treated identically to memory word access. Networks, on the other hand, provide serial data, distributed arbitration, and support more message-oriented protocols. The latter require a more complex handshake, usually involving the exchange of high-level request and acknowledgment messages. Channels fall between the two extremes, consisting of wide data paths of medium distance and often incorporating simplified versions of networklike protocols.

These considerations are summarized in Table 1. Networks typically span more than 1 km, sustain 10 Mbit/second (Ethernet) to 100 Mbit/second (FDDI) and beyond, experience latencies measured in several milliseconds (ms), and the network medium itself is considered to be inherently unreliable. Networks include extensive data integrity features within their protocols,

**Table 1** Comparison of Network, Channel, and Backplane Attributes

	Network	Channel	Backplane
Distance	>1000 m	10-100 m	1 m
Bandwidth	10-100 Mb/s	40-1000 Mb/s	320-1000+ Mb/s
Latency	high (>ms)	medium	low (<μs)
Reliability	low Extensive CRC	medium Byte Parity	high Byte Parity

The comparison is based upon the interconnection distance, transmission bandwidth, transmission latency, inherent reliability, and typical techniques for improving data integrity.

including CRC checksums at the packet and message levels, and the explicit acknowledgment of received packets.

Channels span small 10's of meters, transmit at anywhere from 4.5 Mbytes/second (IBM channel interfaces) to 100 Mbytes/second (HiPPI channels), incur latencies of under 100 μs per transfer, and have medium reliability. Byte parity at the individual transfer word is usually supported, although packet-level check-summing might also be supported.

Backplanes are about 1 m in length, transfer from 40 (VME) to over 100 (FutureBus) MBytes/second, incur sub μs latencies, and the interconnect is considered to be highly reliable. Backplanes typically support byte parity, although some backplanes (unfortunately) dispense with parity altogether.

In the remainder of this section, we will look at each of the three kinds of interconnect, network, channel, and backplane, in more detail.

### B. Communications Networks and Network Controllers

An excellent overview of networking technology can be found in [3]. For a futuristic view, see [4] and [5]. The decade of the 1980's has seen a slow maturation of network technology, but the 1990's promise much more rapid developments. Today, 10 Mbit/second Ethernets are pervasive, with many environments advancing to the next generation of 100 Mbit/second networks based on the FDDI (Fiber Distributed Data Interface) standard [6]. FDDI provides higher bandwidth, longer distances, and reduced error rates, largely because of the introduction of fiber optics for data transmission. Unfortunately cost, especially for replacing the existing copper wire network with fiber, coupled with disappointing transmission latencies, has slowed the acceptance of these higher speed networks. The latency problems have more to do with FDDI's protocols, which are based on a token passing arbitration scheme, than anything intrinsic in fiber-optic technology.

A network system is decomposed into multiple protocol layers, from the application interface down to the method of physical communication of bits on the network. Figure 1 summarizes the popular seven-layer ISO protocol model. The physical and link levels are closely tied to the

<b>Application</b>	Detailed information about the data being exchanged
<b>Presentation</b>	Data representation
<b>Session</b>	Management of connections between programs
<b>Transport</b>	Delivery of packet sequences
<b>Network</b>	Format of individual packets
<b>Link</b>	Access to and control of transmission medium
<b>Physical</b>	Medium of transmission

Fig. 1. Seven-layer ISO protocol model. The physical layer describes the actual transmission medium, be it coax cable, fiber optics, or a parallel backplane. The link layer describes how stations gain access to the medium. This layer deals with the protocols for arbitrating for and obtaining grant permission to the media. The network layer defines the format of data packets to be transmitted over the media, including destination and sender information as well as any check sums. The transport layer is responsible for the reliable delivery of packets. The session layer establishes communication between the sending program and the receiving program. The presentation layer determines the detailed formats of the data embedded within packets. The application layer has the responsibility of understanding how these data should be interpreted within an applications context.

underlying transport medium, and deal with the physical attachment to the network and the method of acquiring access to it. The network, transport, and session levels focus on the detailed formats of communications packets and the methods for transmitting them from one program to another. The presentation and applications layers define the formats of the data embedded within the packets and the application-specific semantics of that data.

A number of performance measurements of network transmission services point out that the significant overhead is not protocol interpretation (approximately 10% of instructions are spent in interpreting the network headers). The culprits are memory system overheads arising from data movement and operating system overheads related to context switches and data copying [7]-[10]. We will see this again and again in the sections to follow.

The network controller is the collection of hardware and firmware that implements the interface between the network and the host processor. It is typically implemented on a small printed circuit board, and contains its own processor, memory mapped control registers, interface to the network, and small memory to hold messages being transmitted and received. The on-board processor, usually in conjunction with VLSI components within the network interface, implements the physical and link-level protocols of the network.

The interaction between the network controller and the host's memory is depicted in Fig. 2. Lists of blocks containing packets to be sent and packets that have been received are maintained in the host processor's memory. The locations of buffers for these blocks are made known to the network controller, and it will copy packets to and from the request/receive block areas using direct memory access (DMA) techniques. This means that the copy of data across the peripheral bus is under the control of the network controller, and does not require the intervention of the host processor. The controller will interrupt the host whenever a message has been received or sent.

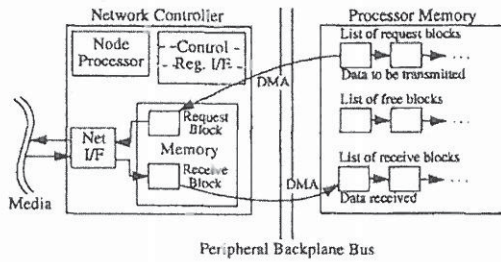


Fig. 2. Network controller/processor memory interaction. The figure describes the interaction between the network controller and the memory of the network node. The controller contains an on-board microprocessor, various memory-mapped control registers through which service requests can be made and status checked, a physical interface to the network media, and a buffer memory to hold request and receive blocks. These contain network messages to be transmitted or which have been received respectively. A list of pending requests and messages already received resides in the host processor's memory. Direct memory operations (DMA's), under the control of the node processor, copy these blocks to and from this memory.

While this presents a particularly clean interface between the network controller and the operating system, it points out some of the intrinsic memory system latencies that reduce network performance. Consider a message that will be transmitted to the network. First the contents of the message are created within a user application. A call to the operating system results in a process switch and a data copy from the user's address space to the operating system's area. A protocol-specific network header is then appended to the data to form a packaged network message. This must be copied one more time, to place the message into a request block that can be accessed by the network controller. The final copy is the DMA operation that moves the message within the request block to memory within the network controller.

Data integrity is the aspect of system reliability concerned with the transmission of correct data and the explicit flagging of incorrect data. An overriding consideration of network protocols is their concern with reliable transmission. Because of the distances involved and the complexity of the transmission path, network transmission is inherently lossy. The solution is to append check-sum protection bits to all network packets and to include explicit acknowledgment as part of the network protocols. For example, if the check sum computed at the receiving end does not match the transmitted check sum, the receiver sends a negative acknowledgment to the sender.

### C. Channel Architectures

Channels provide the logical and physical pathways between I/O controllers and storage devices. They are medium-distance interconnect that carry signals in parallel, usually with some parity technique to provide data integrity. In this subsection, we will describe three alternative channel organizations that characterize the opposite ends of the performance spectrum: SCSI (small computer system

interface), HIPPI (high-performance parallel interface), and FCS (fibre channel standard).

1) *Small Computer System Interface* SCSI is the channel interface most frequently encountered in small form factor (5.25 in diameter and smaller) disk drives, as well as a wide variety of peripherals such as tape drives, optical disk readers, and image scanners. SCSI treats peripheral devices in a largely device-independent fashion. For example, a disk drive is viewed as a linear byte stream; its detailed structure in terms of sectors, tracks, and cylinders is not visible through the SCSI interface. A SCSI channel can support up to eight devices sharing a common bus with an 8-bit-wide data path. In SCSI terminology, the I/O controller counts as one of these devices, and is called the host bus adapter (HBA). Burst transfers at 4 to 5 Mbytes/s are widely available today. In SCSI terminology, a device that requests service from another device is called the master or the initiator. The device that is providing the service is called the slave or the target.

SCSI provides a high-level message-based protocol for communications between initiators and targets. While this makes it possible to mix widely different kinds on devices on the same channel, it does lead to relatively high overheads. The protocol has been designed to allow initiators to manage multiple simultaneous operations. Targets are intelligent in the sense that they explicitly notify the initiator when they are ready to transmit data or when they need to throttle a transfer.

It is worthwhile to examine the SCSI protocol in some detail, to clearly distinguish what it does from the kinds of messages exchanged on a computer network. The SCSI protocol proceeds in a series of phases, which we summarize below:

- **Bus Free:** No device currently has the bus allocated.
- **Arbitration:** Initiators arbitrate for access to the bus. A device's physical address determines its priority.
- **Selection:** The initiator informs the target that it will participate in an I/O operation.
- **Reselection:** The target informs the initiator that an outstanding operation is to be resumed. For example, an operation could have been previously suspended because the I/O device had to obtain more data.
- **Command:** Command bytes are written to the target by the initiator. The target begins executing the operation.
- **Data Transfer:** The protocol supports two forms of the data transfer phase, *Data In* and *Data Out*. The former refers to the movement of data from the target to the initiator. In the latter, data move from the initiator to the target.
- **Message:** The message phase also comes in two forms, *Message In* and *Message Out*. *Message In* consists of several alternatives. *Identify* identifies the reselected target. *Save Data Pointer* saves the place in the current data transfer if the target is about to disconnect. *Restore Data Pointer* restores this pointer. *Disconnect* notifies the initiator that the target is about to give up the data bus. *Command Complete* occurs when the target tells the initiator that the operation has completed. *Message*

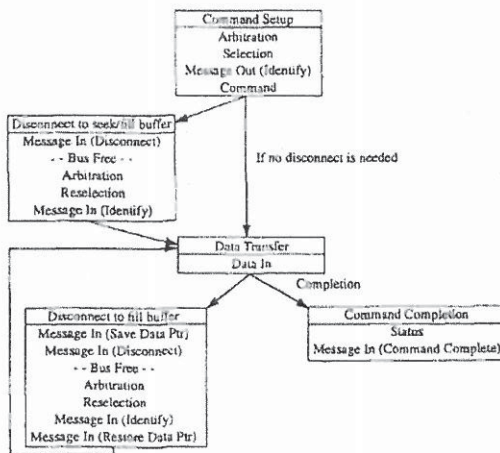


Fig. 3. SCSI phase transitions on a read. The basic phase sequencing for a read (from disk) operation is shown. First the initiator sets up the read command and sends it to the I/O device. The target device disconnects from the SCSI bus to perform a seek and to begin to fill its internal buffer. It then transfers the data to the initiator. This may be interspersed with additional disconnects, as the transfer gets ahead of the internal buffering. A command complete message terminates the operation. This figure is adapted from [40].

*Out* has just one form: *Identify*. This is used to identify the requesting initiator and its intended target.

- **Status:** Just before command completion, the target sends a status message to the initiator.

To better understand the sequencing among the phases, see Fig. 3. This illustrates the phase transitions for a typical SCSI read operation. The sequencing of an I/O operation actually begins when the host's operating system establishes data and status blocks within its memory. Next, it issues an I/O command to the HBA, passing it pointers to command, status, and data blocks, as well as the SCSI address of the target device. These are staged from host memory to device-specific queues within the HBA's memory using direct memory access techniques.

Now the I/O operation can begin in earnest. The HBA arbitrates for and wins control of the SCSI bus. It then indicates the target device it wishes to communicate with during the selection phase. The target responds by identifying itself during a following message out phase. Now the actual command, such as "read a sequence of bytes," is transmitted to the device.

We assume that the target device is a disk. If the disk must first seek before it can obtain the requested data, it will disconnect from the bus. It sends a disconnect message to the initiator, which in turn gives up the bus. Note that the HBA can communicate with other devices on the SCSI channel, initiating additional I/O operations. Now the device will seek to the appropriate track and will begin to fill its internal buffer with data. At this point, it needs to reestablish communications with the HBA. The device now arbitrates for and wins control of the bus. It next enters

the reselection phase, and identifies itself to the initiator to reestablish communications.

The data transfer phase can now begin. Data are transferred one byte at a time using a simple request/acknowledgment protocol between the target and the initiator. This continues until the need for a disconnect arises again, such as when the target's buffer is emptied, or perhaps the command has completed. If it is the first case, the data pointer must first be saved within the HBA, so we can restart the transfer at a later time. Once the data transfer pointer has been saved, the target sequences through a disconnect, as described above.

When the disk is once again ready to transfer, it re-arbitrates for the bus and identifies the initiator with which to reconnect. This is followed by a restore data pointer message to reestablish the current position within the data transfer. The data transfer phase can now continue where it left off.

The command completion phase is entered once the data transfer is finished. The target device sends a status message to the initiator, describing any errors that may have been encountered during the operation. The final command completion message completes the I/O operation.

The SCSI protocol specification is currently undergoing a major revision for higher performance. In the so-called SCSI-1, the basic clock rate on the channel is 10 MHz. In the new SCSI-2, "fast SCSI" increases the clock rate to 20 MHz, doubling the channel's bandwidth from 5 Mbyte/s to 10 Mbyte/s. Recently announced high-performance disk drives support fast SCSI. The revised specification also supports an alternative method of doubling the channel bandwidth, called wide SCSI. This provides a 16-bit data path on the channel rather than SCSI-1's 8-bit width. By combining wide and fast SCSI-2, the channel bandwidth quadruples to 20 Mbyte/s. Some manufacturers of high-performance disk controllers have begun to use SCSI-2 to interface their controllers to a computer host.

2) *High-Performance Parallel Interface* The high performance parallel interface, HIPPI, was originally developed at the Los Alamos National Laboratory in the mid 1980's as a high-speed unidirectional (simplex) point-to-point interface between supercomputers [11]. Thus, two-way communications requires two HIPPI channels, one for commands and write data (the *write channel*) and one for status and read data (the *read channel*). Data are transmitted at a nominal rate of 800 Mbit/s (32-bit-wide data path) or 1600 Mbit/s (64-bit-wide data path) in each direction.

The physical interface of the HIPPI channel was standardized in the late 1980's. Its data transfer protocol was designed to be extremely simple and fast. The source of the transfer must first assert a request signal to gain access to the channel. A connection signal grants the channel to the source. However, the source cannot send until the destination asserts ready. This provides a simple flow control mechanism.

The minimum unit of data transfer is the *burst*. A burst consists of 1 to 256 words (the width is determined by the physical width of the channel; for a 32-bit channel, a

burst is 1024 bytes), sent as a continuous stream of words, one per clock period. A burst is in progress as long as the channel's burst signal is asserted. When the burst signal goes unasserted, a CRC (cyclic redundancy check) word computed over the transmitted data words is sent down the channel. Because of the way the protocol is defined, when the destination asserts ready, it means that it must be able to accept a complete burst.

Unfortunately, the upper level protocol (ULP) for performing operations over the channel is still under discussion within the standardization committees. To illustrate the concepts involved in using HIPPI as an interface to storage devices, we restrict our description to the proposal to layer the IPI-3 device generic command set on top of HIPPI, put forward by Maximum Strategies and IBM Corporation [12].

A logical unit of data, sent from a source to a destination, is called a packet. A packet is a sequence of bursts. A special channel signal delineates the start of a new packet. Packets consist of a header, a ULP (upper layer protocol) data set, and fill. The ULP data consist of a command/response field and read/write data field.

Packets fall into three types: command, response, or data-only. A command packet can contain a header burst with an IPI-3 device command, such as read or write, followed by multiple data bursts if the command is a write. A response packet is similar. It contains an IPI-3 response within a header burst, followed by data bursts if the response is a read transfer notification. Data-only packets contain header bursts without command or response fields.

Consider a read operation over a HIPPI channel using the IPI-3 protocol. On the write channel, the slave peripheral device receives a header burst containing a valid read command from the master host processor. This causes the slave to initiate its read operation. When data are available, the slave must gain access to the read channel. When the master is ready to receive, the slave will transmit its response packet. If the response packet contains a transfer notification status, this indicates that the slave is ready to transmit a stream of data. The master will pulse a ready signal to receive subsequent data bursts.

The original HIPPI specification limits the interconnect distance to 25 m over twisted pair copper cable. A bit serial version of HIPPI has recently been proposed [13] which will make it possible to support gigabit/s data transfers over a distance of up to 10 km using optical fiber. McFarland *et al.* [14] describes a VLSI chip set developed by the Hewlett-Packard Corporation to support serial HIPPI interconnections.

3) *Fibre Channel Standard* The Fibre Channel Standard is a rapidly emerging specification for high-performance bit-serial point-to-point communications over optical fiber [15]. Much like a HIPPI channel, it has been designed to support high-speed computer or storage device to computer communications, albeit over a bit-serial connection. But unlike HIPPI, FCS purposefully blurs the distinction between networks and channels. FCS has been designed as a multilayer series of protocols to make it suitable as the basis of a high-speed network. A crucial aspect of the standard

is its definition of the concept of a switching "fabric," a network formed from switched high-speed links that can be used to communicate between user nodes.

FCS supports three levels of network service: dedicated connections, multiplexed connections, and datagrams. A dedicated connection guarantees sequential delivery of data frames at the full bandwidth of the interconnect, with end-to-end flow control. This level of service is essentially the same as that provided by HIPPI.

At the next level of service, multiplexed connections allow multiple data transfers to time-share the use of a given fiber channel link within the fabric on a frame or multiple-frame basis. End-to-end confirmation is optional and sequential delivery of frames is no longer guaranteed. Flow control is handled on a hop-by-hop basis by the internal switches of the fabric.

FCS's datagram service involves the exchange of a single frame at a time from source to destination through the switching fabric. The transport provided by the FCS does not notify that the datagrams have arrived. This is left to higher level protocols to build on top of the datagram mechanism.

It should be noted that fabrics of switched HIPPI links can also be constructed [16], but these are severely limited in terms of interconnect distance, ability to multiplex the links, and support for network broadcast or multicast. Because it has been designed from the outset as the basis of a high-speed network, with planned support for the existing popular storage device protocols (such as HIPPI, SCSI, and IPI), many experts believe that FCS will be the predominant gigabit/s storage interconnect of the mid 1990's.

#### D. Backplane Architecture

Backplanes are designed to interconnect processors, memory, and peripheral controllers (such as network and disk controllers). They are relatively wide, but short distance. The short distances make it possible to use fast, centralized arbitration techniques and to perform data transfers at a higher clock rate. Backplane protocols make use of addresses and read/write operations, rather than the more message-oriented protocols to be found on networks and channels.

Table 2 gives some of the metrics of three popular backplane buses: VME, FutureBus, and Multibus II (this table is adapted from [2]). For purposes of comparison, we include the same metrics for the SCSI-I channel specification. The two numbers reported for SCSI bandwidth represent the synchronous transfer rate (5 Mbyte/s) and the asynchronous transfer rate (1.5 Mbyte/s) respectively. The table includes the width of the interconnect (including control and data signals), whether the address and data lines are multiplexed, the data width, whether the transfer size is a single or multiple word, the number of bus masters supported, whether split transactions are supported (these are networklike request and acknowledgment messages), the clocking scheme, the interconnect's bandwidth under a variety of assumptions (single versus multiple word transfers, 0 ns access time memories versus

Table 2 Comparisons of Popular Backplane and Channel Interconnects

Metric	VME	FutureBus	MultiBus	SCSI-I
Bus width (signals)	128	96	96	25
Address/data multiplexed?	no	yes	yes	no
Data width	16-32	32	32	8
Xfer size	single/multiple	single/multiple	single/multiple	single/multiple
Number of bus masters	multiple	multiple	multiple	multiple
Split transactions	no	optional	optional	optional
Clocking	async	async	sync	either
Bandwidth, single word (0 ns mem)	25	37	20	5, 1.5
Bandwidth, single word (150 ns mem)	12.9	15.5	10	5, 1.5
Bandwidth, multiple word (0 ns mem)	27.9	95.2	40	5, 1.5
Bandwidth, multiple word (150 ns mem)	13.6	20.8	13.3	5, 1.5
Maximum no. devices	21	20	21	7
Max bus length	0.5 m	0.5 m	0.5 m	25 m
Standard	IEEE 1014	IEEE 896	ANSI/IEEE 1296	ANSI X3.131

150 ns access time), the maximum number of controllers or devices per bus, the maximum bus length, and the relevant ANSI or IEEE standard that defines the interconnect.

The most dramatic differences are in the interconnect width and the maximum bus width. In general, channel interconnects are narrow and long distance while backplanes are wide but short distance.

However, some of the distinctions are being to blur. The SCSI channel has many of the attributes of a bus, FutureBus has certain aspects that make it behave more like a channel than a bus, and nobody could describe a 64-bit HIPPI channel as being narrow! For example, let's consider FutureBus in a little more detail. The bus supports distributed arbitration, asynchronous signaling (that is, no global clocks), single source/multiple destination "broadcast" messages, and request/acknowledge split bus transactions [17]. The latter are very much like SCSI disconnect/reconnect phases. A host issues a read request message to a memory or I/O controller, and then detaches from the bus. Later on, the memory sends a response message to the host, containing the requested data.

In addition, there are several evolving serial bus and device interfaces that further blur the distinction between the kinds of interconnects we have been discussing. Among these proposals are the IEEE Standard P1394 Serial Bus [18] and the Scalable Coherent Interface [19]. Both seek to define high speed, medium distance bit-serial interconnects

that can be used to lash together boxes of hardware containing more traditional parallel backplanes. These are still backplane interconnects, however, because the protocols are based on memory-oriented read/write operations rather than the exchange of network messages.

### III. STORAGE TRENDS

#### A. The Storage Hierarchy

1) *Concept of Storage Hierarchy:* The storage hierarchy is traditionally modeled as a pyramid, with a small amount of expensive, fast storage at the pinnacle and larger capacity, lower cost, and lower performance storage as we move toward the base. In general, there are order of magnitude differences in capacity, access time, and cost among the layers of the hierarchy. For example, main memory is measured in megabytes, costing approximately \$50/Mbyte, and can be accessed in small numbers of microseconds. Secondary storage, usually implemented by magnetic disk, is measured in gigabytes, costs below \$5/Mbyte, and is accessed in tens of milliseconds. These costs are "order of magnitude" for workstations and personal computers; the costs for mainframes and supercomputers are typically much higher. The operating system can create the illusion of a large fast memory by judiciously staging data among the levels. However, the organization of the storage hierarchy must adapt as magnetic and optical recording methods continue to improve and as new storage devices become available.




Fig. 4. Typical storage hierarchy, circa 1980. The microsecond access is provided by the file cache, a small number of bytes stored in semiconductor memory. Medium capacity, denominated in several hundred megabytes with tens of millisecond access, is provided by disk. Tape provides unlimited capacity, but access is restricted to tens of seconds to minutes.

Figure 4 depicts the storage hierarchy of a typical mini-computer of 1980. (It should be noted that large mainframe and supercomputer storage hierarchies were more complex than what is depicted here.) A small file cache (or buffer), allocated by the operating system from the machine's semiconductor memory, provides the fastest but most expensive access. The job of the cache is to hold data likely to be accessed in the near future, because it is near data recently accessed (spatial locality) or because it has recently been accessed itself (temporal locality). Prefetching is a strategy that accesses larger chunks of file data than requested by an application, in the hope that it will soon access spatially local data.

Either a buffer or a cache can be used to decouple application accesses in small units from the larger units needed to efficiently utilize secondary storage devices. It is not efficient to amortize the millisecond latency cost to access secondary storage for a small number of bytes. Accesses in the range of 512 to 8192 bytes are more appropriate. The primary distinction between application memory and a cache is the latter's ability to keep resident certain data. For example, frequently accessed file directories can be held in a cache, thus avoiding slow accesses to the lower levels of the hierarchy.

Secondary storage is provided by magnetic disk. Data are recorded on concentric tracks on stacked platters, which have been coated with magnetic materials. The same track position across the platters is called a cylinder. A mechanical actuator positions the read-write heads to the desired recording track, while a motor rotates the platters containing the data under the heads.

Tertiary storage, provided primarily for archive/backup, is implemented by magnetic tape. A spool of magnetic tape is drawn across the read-write mechanism in a sequential fashion. A good rule of thumb for a unit of tertiary storage media, such as a tape spool, is that it should have as much capacity as the secondary storage devices it is meant to back up. As disk devices continue to improve in capacity, tertiary storage media are driven to keep pace.

In 1980, a typical machine of this class would have one to

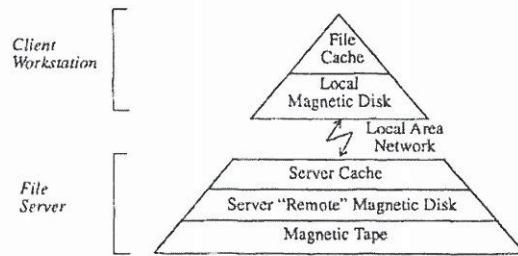


Fig. 5. Typical storage hierarchy, circa 1990. The file cache has become substantially larger, and may be partially duplicated at the client in addition to the server. Secondary storage is split between local and remote disk. Tape continues to provide the third level of storage.

two megabytes of semiconductor memory, of which only a few thousand bytes might be allocated for input/output buffers or file system caches. The secondary storage level might include a few hundred megabytes of magnetic disk. The tape storage level is limited only by the amount of shelf space in the machine room.

2) *Evolution of the Storage Hierarchy:* Figure 5 shows the storage hierarchy distributed across a workstation/server environment of today. Most of the semiconductor memory in the server can be dedicated to the cache function because a server does not host conventional user applications. The file system "metadata," that is, the data structures describing how logical files are mapped onto physical disk blocks, can be held in fast semiconductor memory. This represents much of the active portion of the file system. Thus, disk latency can be avoided while servicing user requests.

The critical challenge for workstation/server environments is the added latency of network communications. These are comparable to those of magnetic disk, and are measured in small tens of milliseconds. The figure shows one possible solution, which places small high-performance disks in the workstation, with larger potentially slower disks at the server.

If most accesses can be serviced by the local disks, the network latencies can be avoided altogether, improving client performance and responsiveness. However, there are several choices for how to partition the file system between the clients and the servers. Each of these partitionings represents a different trade-off between system cost, the number of clients per server, and the ease of managing the clients' files.

A *swapful* client allocates the virtual memory swap space and temporary files to its local disk. The operating system's files and user files remain on the server. This reduces some of the network traffic to the server, leaving the issues of system management relatively uncomplicated. For example, in this configuration, the local disk does not need to be backed up. However, executing an operating system command still requires an access to the remote server.

A *dataless* client adds the operating system's files to the client's local disk. This further reduces the client's demand






Fig. 6. Typical storage hierarchy, circa 1995. Conventional disks have been replaced by disk arrays, a method of obtaining much higher I/O bandwidth by striping data across multiple disks. A new level of storage, "near line," emerges between disk and tape. It provides very high capacity, but at access times measured in seconds.

on the server, thus making it possible for a single server and network to support more clients. While it is still not necessary to back up the local disk, the system is more difficult to administer. For example, system updates must now be distributed to all of the workstations.

A *diskfull* client places all but some frequently shared files on the client. This yields the lowest demands on the server, but represents the biggest problems for system management. Now the personal files on the local disk need to be backed up, leading to significant network traffic during backup operations.

An alternative approach leverages the lower cost semiconductor memory to make feasible large file caches (approximately 25% of the available memory within the workstation) in the client workstation. These "client" caches provide an effective way to circumvent network latencies, if the network protocols allow file writes to be decoupled from communications with the server (see the discussion of NFS protocols in the next section). The approach, called diskless clients, has been used with great success in the Sprite Network Operating System [20], where they report an ability to support five to ten times as many clients per server as more conventional client/server organizations.

Figure 6 depicts one possible scenario for the storage hierarchy of 1995. Three major technical innovations shape the organization: disk arrays, near-line storage subsystems based on optical disk or automated tape libraries, and network distribution. We concentrate on disk arrays and near-line storage system technology in the remainder of this subsection. Network distribution is covered in Section IV.

### B. Storage Technology

1) *Disk Arrays*: Because of the rapidly decreasing form factor of magnetic disks, it is becoming attractive to replace a small number of large disk drives with very many small drives. The resulting secondary storage system can have much higher capacity since small format drives traditionally obtain the highest areal densities. And since the performance of both large and small disk drives is limited by mechanical delays, it is no surprise that performance can be dramatically improved if the data to be accessed are

spread across many disk actuators. Disk arrays provide a method of organizing many disk drives to appear logically as a very reliable single drive of high capacity and high performance [21].

Disk array organizations are organized into a multilevel taxonomy. Here, we concentrate on the two most prevalent RAID organizations: RAID level 3 and RAID level 5. Each of these spreads data across  $N$  data disks and an  $N + 1$ st redundancy disk. The group of  $N + 1$  disks is called a stripe set. In a RAID level 3 organization, data are interleaved in large blocks (for example, a track or cylinder) across all of the disks within a stripe set. The redundancy disk contains a parity bit computed bitwise across the rows of bits on the associated data disks. If a disk should fail, its contents can be reconstructed simply by examining the surviving  $N$  disks and restoring the sense of the parity computed across the bit rows. Suppose that the redundancy disk contained odd parity before the failure. If, after a disk failure, the examination of a bit row yields even parity, then the failed disk must have had a 1 in that bit row. Similarly, if the row has odd parity, then the missing bit must have been a 0. RAID level 3 organizations are read and written in full stripe units, simultaneously accessing all disks in the stripe. The organization is most suitable for high-bandwidth applications such as image processing and scientific computing.

If RAID level 3 is organized for high bandwidth, then RAID level 5 is organized for high I/O rate. The basic organization is the same: a stripe set of  $N$  data disks and one redundancy disk. However, data are accessed in smaller units, thus making it possible to support multiple simultaneous accesses. Consider a data write operation to a single disk sector. This requires the parity redundancy to be updated as well. We accomplish this by first determining the bit changes to the data sector and then invert exactly these bits in the associated parity sector. Thus a logical single sector write may involve four physical disk accesses: read old data, read old parity, write new data, and write new parity. Since placing all parity sectors on a single disk drive would limit the array to a single write operation at a time, the parity sectors are actually interleaved across all disks of the stripe. A RAID level 5 can perform  $N + 1$  simultaneous reads and simultaneous writes (in the best case).

2) *Near-Line Storage*: A comparable revolution has taken place in tertiary storage: the arrival of near-line storage systems. These provide relatively rapid access to enormous amounts of data, frequently stored on removable, easy to handle optical disk or magnetic tape media. This is accomplished by storing the high-capacity media on shelves that can be accessed by robotic media "pickers." When a file needs to be accessed, special file management software identifies where it can be found within the tape or optical disk library. The picker exchanges the currently loaded media with the one containing the file to be accessed. This is accomplished within a small number of seconds, without any intervention by human operators. By carefully exploiting caching techniques, in particular, using the secondary storage devices as a cache for the near-line store,

the very large storage capacity of a tertiary storage system can appear to have access times comparable to magnetic disks at a fraction of the cost. We describe the underlying storage technologies next.

3) *Optical Disk Technology for Near-Line Storage:* Optically recorded disks have long been thought to be ideal for filling the near line level of the storage hierarchy [22]. They combine improved storage capacity (2 Gbytes per platter surface originally to over 6 Gbytes per side today) with access times that are approximately a factor of 10 slower than conventional magnetic disks (several hundred milliseconds). The first generation of optical disks were written once, but could be read many times, leading to the term "WORM" to describe the technology. The disk is written by a laser beam. When it is turned on, it records data in the form of pits or bubbles in a writing layer within the disk. The data are read back by detecting the variations of reflectivity of the disk surface.

The write-once nature of optical storage actually makes it better suited for an archival medium than near-line storage, since it is impossible to accidentally overwrite data once they have been written. A problem has been its relatively slow transfer rate, 100K – 200K bytes per second. Newer generations of optical drives now exceed one megabyte per second transfers.

Magneto-optical technologies, based on a combination of optical and magnetic recording techniques, have recently led to the availability of erasable optical drives. The disk is made of a material that becomes more sensitive to magnetic fields at high temperatures. A laser beam is used to selectively heat up the disk surface, and once heated, a small magnetic field is used to record on the surface. Optical techniques are used for reading the disk, by detecting how the laser beam is deflected by different magnetizations of the disk surface. Read transfer rates are comparable to those of conventional magnetic disks. Access times are still slower than a magnetic disk owing to the more massive read/write mechanism holding the laser optics, which takes longer to position than the equivalent low mass magnetic read/write head assembly. The write transfer rate is worse in optical disk systems because (1) the disk surface must first be erased before new data can be recorded, and (2) the written data must be reread to verify that they were written correctly to the disk surface. Thus, a write operation could require three disk revolutions before it completes. (The study [23] details the trends and technology challenges for future optical disk technologies).

Nevertheless, as the form factor and price of optical drives continue to decrease, optical disk libraries are becoming more pervasive. Sony's recent announcement of a consumer-oriented recordable music compact disk could lead to dramatic reductions in the cost of optical disk technology. As an example of an inexpensive optical disk system, let's examine the Hewlett-Packard Series 6300 Model 20GB/A Optical Disk Library System [24]. Based on 5.25 in rewritable optical disk technology, the system provides two optical drives, 32 read/write optical disk cartridges (approximately 600 Mbytes per cartridge), and a

robotic disk changer that can move cartridges to and from the drives, all in a desk side unit the size of a three-drawer filing cabinet. The optical cartridges can be exchanged in 7 s, and require a 4 s load time and 2.4 s spin-up time. The unload and spin-down times are 2.8 and 0.8 s respectively. An average seek time requires 95 ms. The drives can sustain 680 Kbytes/s transfers on reads and 340 Kbytes/s transfers on writes.

The Kodak Optical Disk System 6800 Automated Disk Library is characteristic of the high end [25]. The system can be configured with 50 to 150 optical disk platters, and one to three optical disk drives. It is capable of storing from 340 Gbytes to 1020 Gbytes (3.4 Gbytes for each side of a 14 in platter). The average disk change time is 6.5 s. The optical disk surface is organized into five bands of varying capacity, with a certain number of tracking windows per band. The drives can sustain 1 Mbyte/s transfers, with 100 ms access times for data anywhere within the current band to 700 ms for data anywhere on the surface.

4) *Magnetic Tape Technology for Near-Line Storage* The sequential nature of the access to magnetic tape has traditionally dictated that it be used as the medium for archive. However, the success of automated tape libraries from Storage Technology Corporation has demonstrated that tape can be used to implement a near-line storage system. The most pervasive magnetic tape technology available today is based on the IBM 3480 half-inch tape cartridge, storing 200 Mbytes and providing transfer rates of 3 Mbytes per second. A second generation technology recently introduced doubles the tape capacity and transfer rate.

However, there has been an enormous increase in tape capacity, driven primarily by helical scan recording methods. In a conventional tape recording system, the tape is pulled across stationary read/write recording heads. Recorded data tracks run in parallel along the length of the tape. On the other hand, helical scan methods slowly move the tape past a rapidly rotating head assembly to achieve a very high tape to head speed. The tape is wrapped at an angle around a rotor assembly, yielding densely packed recording tracks running diagonally across the tape. The technology is based on the same tape transport mechanisms developed for video cassette recorders in the VHS and 8 mm tape formats and the newer digital audio tape (DAT) systems.

Each of these systems provides a very high storage capacity in a small, easy-to-handle cartridge. The small form factor make these tapes particularly attractive as the basis for automated data libraries. Tape systems from Exabyte, based on the 8 mm video tape format, can store 2.3 Gbytes and transfer at approximately 250 Kbytes per second. A second generation system now available doubles both the capacity and the transfer rate. A tape library system based on a 19 in rack can hold up to four tape readers and over one hundred 8 mm cartridges, thus providing a storage capacity of 250–500 GB's [26].

DAT tape provides smaller capacity and bandwidth than 8 mm, but enjoys certain other advantages [27]. Low-cost tape readers in the 3.5 in form factor, the size of a personal computer floppy disk drive, are readily available.

This makes possible the construction of tape libraries with a higher ratio of tape readers to tape media, increasing the aggregate bandwidth to the near-line storage system. In addition, the DAT tape formats support subindex fields which can be searched at a speed 200 times greater than the normal read/write speed. A given file can be found on a DAT tape in an average search time of only 20 s, compared with over 10 min for the 8 mm format.

VHS-based tape systems can transfer up to 4 Mbytes/s and can hold upto 15 Gbytes per cartridge. Tape robotics in use for the broadcast industry have been adapted to provide a near-line storage function. Helical scan techniques are not limited to consumer applications, but have also been applied for certain instrument recording applications, such as satellite telemetry, which require high capacity and high bandwidth. These tape systems are called DD1 and DD2. A single tape cartridge can hold up to 150 Gbytes, and can transfer at a rate of up to 40 Mbyte/s. However, such systems are very expensive, and a good rule of thumb is that the tape recorder will cost \$100 000 for each 10 Mbytes/s of recording bandwidth it can support.

5) *Optical Tape Technology for Near-Line Storage* A recording technology that appears to be very promising is optical tape [28]. The recording medium is called digital paper, a material constructed from an optically sensitive layer that has been coated onto a substrate similar to magnetic tape. The basic recording technique is similar to write-once optical disk storage: a laser beam writes pits in the digital paper to indicate the presence (or absence) of a bit. Since the pits have lower reflectivity than the unwritten tape, a reflected laser beam can be used to detect their presence. One 12 inch by 2400 ft reel can hold 1 TB of data, can be read or written at the rate of 3 Mbyte/s, and can be accessed in a remarkable average time of 28 s.

Two companies are developing tape readers for digital paper: the CREO Corporation and the LaserTape Corporation. CREO makes use of a 12 inch tape reel and a unique laser scanner array to read and write multiple tracks 32 bits at a time [29]. The system is rather expensive, selling for over \$200 000. LaserTape places digital paper in a conventional 3480 tape cartridge (50 Gbyte capacity and 3 Mbyte/s transfer rate), and replaces a 3480 tape unit's magnetic read/write heads with an inertialess laser beam scanner. The scanner operates by using a high-frequency radio signal of known frequency to vibrate a crystal, which is then transferred to a laser beam to steer it to the desired read/write location. A 3480 tape reader can be "retrofitted" for approximately \$20 000. Existing tape library robotics for the 3480 cartridge form factor can be adapted to LaserTape without changes.

6) *Summary:* Table 3 summarizes the relevant metrics of the alternative storage technologies, with a special emphasis on helical scan tapes. The metrics displayed are the capacity, bits per inch (BPI), tracks per inch (TPI), areal density (BPI\*TPI in millions of bits per square inch), data transfer rate (Kbytes per second sustained transfers), and average positioning times. The latter is especially important for evaluating near-line storage media. An access time




Fig. 7. I/O data flow. In response to a read operation, data move from the device, to the embedded controller, to the I/O controller, to operating system buffers, and finally to the application across a variety of different interfaces.

measured in a small number of seconds begins to make tape technology attractive for near-line storage applications, since the robotic access times tend to dominate the time it takes to pick, load, and access data on near-line storage media.

### C. Storage Controller Architecture

1) *I/O Data Flow:* Figure 7 shows the various interfaces across which a typical I/O request must flow. The actual flow of data starts at the I/O device. In the following discussion, we will assume that the device is an intelligent magnetic disk for something like a SCSI interface and that we are considering a read operation. The mechanical portion of the disk drive is called the head/disk assembly, or HDA. The control and interface to the outside world is provided by an embedded controller.

Data move across a bit-serial interface from the disk signal processing electronics to track buffers associated with the embedded controller. The amount of memory associated with the track buffers varies from 32 Kbytes to 256 Kbytes. Since the typical track on today's small form factor disks is in the range of 32 K - 64 Kbytes, a typical embedded controller can buffer more than one track.

The interface between the embedded controller and the host is provided by an I/O controller. We called such a controller a host bus adapter, or HBA, in subsection II-C. It couples the host peripheral bus to the disk channel interface. Data are staged into buffers within the HBA, from which they are copied out via direct memory access techniques to the host's memory. The typical size of I/O controller buffers is in the range of 1 to 4 Mbytes.

The host's memory is coupled to the processor via a high-speed cache memory. The connection to the I/O controllers is through a slower speed peripheral bus. Direct memory access operations copy data from the controller's buffers to operating system buffers in main memory. Before the data can be used by the application, they may need to be copied once again, to stage them into a portion of the memory address space that is accessible to the application. Note that the same memory and operating system overheads that limit network performance also affect I/O performance. This is critically important in file and storage servers, where both the I/O and network traffic must be routed through the memory system bottleneck.

If the host is actually a file server, the size of the operating system's buffers may be quite large, perhaps as large as

Table 3 Relevant Metrics for Alternative Storage Technologies

Technology	Capacity (MB)	BPI	TPI	BPI*TPI (million)	Data Transfer (Kbyte/s)	Access Time
<i>Conventional Tape</i>						
Reel-to-reel (1/2 in)	140	6250	18	0.11	549	minutes
Cartridge (1/4 in)	150	12000	104	1.25	92	minutes
IBM 3480 (1/2 in)	200	22860	38.1	0.87	3000	seconds
<i>Helical Scan Tape</i>						
VHS (1/2 in)	15000	not specified	not specified	unknown	4000	minutes
Video (8 mm)	4600	43200	1638	70.56	492	minutes
DAT (4 mm)	1300	61000	1870	114.07	183	20 s
<i>Optical Tape</i>						
CREO (35 mm)	1 TB	9336000	24	224	3000	28 s
<i>Magnetic Disk</i>						
Seagate Elite-1 (5.25 in)	1350	33528	1880	63.01	3000-4000	16 ms
IBM 3390 (10.5 in)	3800	27940	2235	62.44	4250	20 ms
Floppy Disk (3.5 in)	2	17434	135	2.35	92	1 s
<i>Optical Disk</i>						
CD ROM (3.5 in)	540	27600	15875	438.15	183	1 s
Sony MO (5.25 in)	640	24130	18796	453.54	87.5	100 ms
Kodak (14 in)	3200	21000	14111	296.33	1000	100's ms

128 Mbytes. In addition, the data flow must be extended to include transfers across the network interconnect into the application's address space on the client. A detailed examination of the operating system management of the I/O path will be left until Section IV.

2) *Internal Organization of I/O Controller:* Figure 8 shows the internal organization of a typical high-performance host bus adapter I/O controller. Interestingly enough, it is not very different in its internal architecture from the network controller of Figure 2. Usually implemented on a single printed circuit board, the controller contains a microprocessor, a modest amount of memory dedicated to buffers and run-time data structures, a ROM to hold the controller firmware, a DMA/peripheral bus interface, and an I/O channel interface.

The system interface is also similar to the network controller described previously. Request blocks containing I/O commands and data are organized into as a linked list in the host memory. The host writes to a memory-mapped command register within the I/O controller to initiate an operation. Using DMA techniques, the controller fetches the request blocks into its own memory. The on-board microprocessor unpackages the I/O commands and write data, and sends these over the I/O channel interface. Status and read data are repackaged into response blocks that are copied back to reserved buffers in the host memory. The host can choose whether the I/O controller will interrupt the host whenever an operation has been completed.

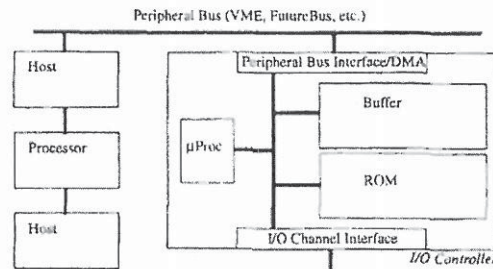


Fig. 8. Internal organization of an I/O controller. The I/O controller couples a peripheral bus to the I/O channel via a buffer memory. Hardware in the peripheral interface implements direct memory access between this buffer memory and the host's memory. The I/O channel interface implements the handshaking protocols with the I/O devices. The microprocessor is the "traffic cop," coordinating the actions of the two interfaces.

The controller of Fig. 8 is notable because of its support for direct memory access. Some lower performance controllers require that commands and data be written a word (or half word) at a time to memory-mapped controller registers over the peripheral bus. Since a typical command block can be 16 to 32 bytes in length, simply down-loading a command may take tens of microseconds, requiring a good deal of host processor intervention.

In implementing a high-performance file service on a network, a critical relationship exists between the network

and I/O controller architectures. The network interface and the I/O controller must be coupled by a high-performance interconnect and memory system. This key observation provides the motivation for several of the systems reviewed in Section V, especially the prototype being developed at U. C. Berkeley described in subsection V-F.

#### IV. SOFTWARE TRENDS

##### A. Network File Systems

One of the most important software developments over the past decade has been the rapid development of the concept of remote file services. In a location transparent manner, these systems provide a client with the ability to access remote files without the need to resort to special naming conventions or special methods for access.

It is important to distinguish between the related concepts of block server and file server. A block server (sometimes called a network disk) provides the client with a physical device interface over a network. The block server supports read and write requests to disk blocks, albeit to a disk attached to a remote machine. A file server supports a higher level interface, providing the complete file abstraction to the client. The interface supports file creation, logical reads and writes, file deletion, etc. In a file server, file system related functions are centralized and performed by the server. In a block server, these functions must be handled by the clients, and if the disks are to be shared across machines, this requires distributed coordination among them.

The most ubiquitous file system model is based on that of UNIX, and so we begin our discussion with its structure. A file is uniquely named within a hierarchical name space based on directories. As far as the user is concerned, a file is nothing more than an uninterpreted stream of bytes. The file system provides operations for positioning within the file for the purpose of reading and writing bytes. Internally, the file system keeps track of the mapping between the file's logical byte stream and their physical placement within disk blocks through a data structure called an inode. The inode is "metadata," that is, data about data, and contains information such as the device containing the file, a list of the physical disk blocks containing the file's data, and pointers to additional disk blocks (called indirect blocks) should the file be large enough to exceed the mapping space of a single inode.

From the operating system perspective, tracing an I/O request from the application to disk proceeds as follows. The application program must make a system call, such as read or write, to request service from the operating system. This is handled by the UNIX system call layer, which in turn calls the file system to handle the request in detail. Within the file system are block I/O routines which handle read or write requests. These call a particular disk driver to schedule the actual disk transfers. The software layers are shown in Fig. 9.

The figure shows the software architecture for a file system on the same machine as the client application. The




Fig. 9. Software layers in the UNIX file system. The UNIX system call layer dispatches a read or write request to the file system, which in turn calls a block I/O routine. This calls a specific device driver to handle the scheduling of the I/O request.




Fig. 10. Software layers in the UNIX file system extended for NFS. The VFS interface allows requests to be mapped transparently among local file systems and remote file systems.




Fig. 11. Path of an NFS client to server request. A request to access a remote file is handled by the client's VFS, which maps the request through the NFS layer into RPC calls to the server over the network. At the server end, the requests are presented to the VFS, this time to be mapped into calls on the server's local UNIX file system.

major innovation of SUN's network file system, or NFS, is its ability to map remote file systems into the directory structure of the client's machine. That is, it is transparent to the user whether the referenced file is available locally or is being accessed over the network. This is accomplished through the new abstraction of a virtual file system, or VFS [30]. The VFS interface allows file system requests to be dispatched to the local file system, or sent to a remote server across the network. The generic software layers are shown in Fig. 10 and the path through the software taken between the client and the server is shown in Fig. 11.

The access to the remote machine is implemented via a synchronous remote procedure call (RPC) mechanism. This is a communications abstraction that behaves much like a conventional procedure call, except that the procedure being invoked may be on a remote "server" machine. Since the

RPC is synchronous, the client must wait or block until the server has completed the call and returned the requested data or status.

The NFS protocol is a collection of procedure calls and parameters built on top of such a RPC mechanism. One of the key design decisions of NFS is to make this protocol stateless. This means that each procedure call is completely self-describing; the server keeps track of no past requests. This choice was made to drastically reduce the complexity of recovery. In the event of a server crash, the client simply retries its request until it is successfully serviced. As far as the client is concerned, there is no difference between a crashed server and one that is merely slow. The server need not perform any recovery processing. Contrast this with a "stateful" protocol, in which both servers and clients must be able to detect and recover from crashes.

However, the stateless protocol has significant implications for file system performance. In order to be stateless, the server must commit any modified user data and file system metadata to stable storage before returning results. This implies that file writes cause the affected data blocks, inodes, and indirect blocks to be written from in-memory caches to disk. In addition, housekeeping operations such as file creation, file removal, and modifications to file attributes must all be performed synchronously to disk.

Some controversy surrounds the real source of bottlenecks in NFS performance. Network protocol overheads and server processing are possible culprits. However, it has now become clear that the real problem is the stateless nature of the NFS protocol, and its associated forced disk writes. By making file system operations synchronous with disk, the performance of the file system is (overly) coupled to the performance of the disk system [31].

### B. File Server Architecture

In this subsection, we examine the flow of a network-based I/O request as it arrives at the network interface, through the file server's hardware and software, to the storage devices and back again to the network. Our goal is to bring together the discussions of network interface, I/O controller, and network file system processing of an I/O request, initiated by a client on the network.

Figure 12 shows the hardware/software architecture of a conventional workstation-based file server. A data read request arrives at the Ethernet controller. The network messages are copied from the network controller to the server's primary memory. Control passes through the software levels of the network driver and protocol interpretation to process the request. At the file system level, to avoid unnecessary disk accesses, the server's primary memory is interrogated to determine if the requested data have already been cached from disk.

If the request cannot be satisfied from the file cache, the file system will issue a request to the disk controller. The retrieved data are then staged by the disk controller from the I/O device to the primary memory along the backplane bus. Usually it must be copied (at least) one more time, into templates for the response network messages. The software




Fig. 12. Conventional file server architecture. An NSF I/O request arrives at the Ethernet interface of the server. The request is passed through to the network driver, the protocol processing software, and the file system. The request may be satisfied by data cached in the primary memory; if not, the data must be accessed from disk. At this point, the process is reversed to send the requested data back over the network. This figure is adapted from [37].

path returns through the file system, protocol processing, and network drivers. The network response messages are transmitted from the memory out through the network interface.

There are two key problems with this architecture. First, there is the long instruction path associated with processing a network-based I/O request. Second, as we have already seen, the memory system and the backplane bus form a serious performance bottleneck. Data must flow from disk to memory to network, passing through the memory and along the backplane several times. In general, the architecture has not been specialized for fast processing between the network and disk interfaces. We will examine some approaches that address this limitation in Section V.

### C. Mass Storage System Reference Model

Supercomputer users have long had to deal with the problem that high-performance machines do not come with scalable I/O systems. As a result, each of the major supercomputer centers has been forced to develop its own mass storage system, a network-based storage organization in which files are staged from the back-end storage server, usually from a near-line subsystem, to the front-end supercomputer.

The mass storage system (MSS) reference model was developed by the managers of these supercomputer centers, to promote more interoperability among mass storage systems and influence vendors to build such systems to a "standard" [32]. The purpose of the reference model is to provide a framework within which standard interfaces can be defined. They begin with the underlying premise that the storage system will be distributed over heterogeneous machines potentially running different operating systems. The model firmly endorses the client/server model of computation.

The MSS reference model defines six elements of the mass storage system: name server, bitfile client, bitfile server, storage server, physical volume repository, and bitfile mover (see Fig. 13). Bit files are the model's terminology for uninterpreted bit data streams. There are different ways to assign these elements to underlying hardware. For example, the name server and bitfile server may run a single mass storage control processor, or they may run on independent communicating machines.




Fig. 13. Elements of the mass storage system reference model. The figure shows the interactions among the elements of the model. Command flows are shown in light lines while data flows are in heavy lines. The reference model clearly distinguishes among the software functions of name service, mapping of logical files onto physical devices, management of the physical media, and the transfer of files between the storage system and clients. The figure has been adapted from [32].

An application's request for I/O service begins with a conversation with the name server. The name service maps a user readable file name into an internally recognized and unique bitfile ID. The client's requests for data are now sent to the bitfile server, identifying the desired files through their ID's. The bitfile server maps these into requests to the storage server, handling the logical aspects of file storage and retrieval, such as directories and descriptor tables. The storage server handles the physical aspects of file storage, and manages the physical data volumes. It may request the physical data repository to mount volumes if they are currently off-line. Storage servers may be specialized for the kinds of volumes they need to manage. For example, one storage server may be specialized for tape handling while another manages disk. The bitfile mover is responsible for moving data between the storage server and the client, usually over a network. It provides the components and protocols for high-speed data transfer.

The MSS reference model has been incorporated into at least one commercial product: the Unitree file management system sold by General Atomics, Inc. This is a UNIX-based hierarchical storage management system, based on software originally developed at the Lawrence Livermore National Laboratory.

## V. CASE STUDIES

In this section, we look at a variety of commercial architectures and research prototypes for high-performance networks, file servers, and storage servers. Within these systems, we will see a common concern for providing high bandwidth between network interfaces and I/O device controllers.

### A. Ultranet

1) *General Organization:* The UltraNetwork is a hub-based multihop network capable of achieving up to 1 Gbit/s transmission rates. Its most frequent application is as a local area network for interconnecting workstations, storage servers, and supercomputers.




Fig. 14. UltraNet configuration. The network interconnection topology is formed by hubs connected by optical serial links. The maximum link speed is 250 Mb/s; higher transmission bandwidth is obtained by interleaving across multiple links. Host-based adapters plug into computer backplanes while adapters for channels such as HiPPI reside within the hub.

Figure 14 depicts a typical Ultranet configuration. The hubs provide the basis of the high-speed interconnect, by providing special hardware and software for routing incoming network packets to output connections. Hubs are physically connected by serial links, which consist of two unidirectional connections, one for each direction. If optical fiber is chosen for the links, data can be transmitted at rates of up to 250 Mbit/s and distances to 4 km. The Gbit transmission rate is achieved by interleaving transmissions across four serial links. The point-to-point links are terminated by link adapters within the hubs, special hardware that routes the transmissions among input and output serial links. These are described in more detail below.

Computers are connected to the network in two different ways: through host adapters and hub-resident adapters. A host-based adapter is similar to the network controller described in Fig. 2, and resides within the host computer's backplane. This kind of interface is appropriate for machines with industry standard backplanes, such as workstations and mini-supercomputers. In these kinds of clients, processors and I/O controllers, including the network interface, are treated as equals with respect to memory access. The adapter contains an on-board microprocessor and can perform its own direct memory accesses, just like any other peripheral controller.

A different approach is needed for mainframe and supercomputers, since these classes of machines connect to peripherals through special channel interfaces rather than standard backplanes. I/O devices are not peers, but are treated as slaves by the processor. The hub-resident adapters place the network interface to the Ultranet within the hub itself. These provide a standard channel interface to the computer, such as HiPPI or the IBM block multiplexer interface.

2) *UltraNet Hub Organization:* The heart of the UltraNet hub is a 64-bit wide (plus 8 parity bits), high-bandwidth backplane called the UltraBus. Its maximum bandwidth is 125 Mbyte/s. The serial links from other hubs and host-based adapters are interfaced to the UltraBus through link multiplexers, which in turn are controlled by the link adapters. The link adapters route the serial data to the parallel interface of the UltraBus. Physically, it is a bus, but logically, the interconnect is treated more like a local area network. Packets are written to the bus by the source link




Fig. 15. Internal organization of the UltraNet hub. The serial links, one for each direction, connect to the link multiplexers. Each link mux can handle up to four serial link pairs. The link adapter interfaces the link multiplexers to the wide, fast UltraNet bus. The link adapter has enough intelligence to do its own routing of network traffic among the link multiplexers it manages. With a well-configured hub, little traffic should need access to the UltraBus.

adapter and are intercepted by the destination link adapter. If the output link is controlled by the same link adapter as the input link, the transfer can be accomplished without access to the UltraBus. Figure 15 illustrates the internal organization of the hub.

The link adapter contains a protocol processor and two modules that interface to the link multiplexers on the one hand and the UltraBus on the other. The protocol processor is responsible for handling the network traffic. The data path that couples the personality modules on either side of the protocol processor consists of two unidirectional 64-bit-wide buses with speed matching FIFO's at the interface boundaries. The buses operate independently and achieve peak transfers of 100 Mbyte/s.

The protocol processor consists of three components: the data acknowledgment and command block processor (DACP), the control processor (CP), and the transfer engine (TE). The DACP performs fast processing of protocol headers and request blocks. The CP is responsible for managing the network, such as setting up and deleting connections between network nodes. The TE rapidly moves data through the protocol processor.

Figure 16 depicts the protocol architecture supported by the UltraNet. The combination of the UltraNet firmware and software implements the industry standard TCP/IP protocols on top of the UltraNet, as well as UltraNet specific protocols. The lower levels of the network protocol, namely the transport, network, data link, and physical link, are implemented with the assistance of the UltraNet protocol processor and host or hub-resident adapter hardware.

### B. Digital Equipment Corporation's VAXCluster and HSC-70

1) *VaxCluster Concept:* Digital Equipment Corporation's VAXCluster concept represents one approach for providing networked storage service to client computers [33], [34]. The VAXCluster is a collection of hardware and software services that closely couple together VAX computers and hierarchical storage controllers (HSC's). A VAXCluster lies somewhere between a "long distance" peripheral bus and a communications network: a high-speed physical link cou-

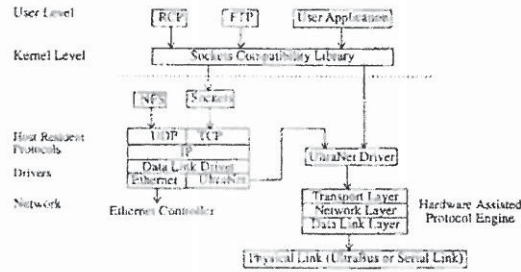


Fig. 16. UltraNet protocol architecture access to the UltraNet is through the standard UNIX socket interface. It is possible to use standard TCP/IP protocols on top of the UltraNet or UltraNet-specific protocols. The lower levels of the network are implemented with the assistance of the protocol engines and adapters throughout the UltraNet system.

ples together the processors, but message-oriented protocols are used to request and receive services. The VAXCluster concept is characterized by (1) a complete communications architecture, (2) a message-oriented computer interconnect, (3) hardware support for the connection to the interconnect, and (4) message-oriented storage controllers.

The hardware organization of a VAXCluster is shown in Fig. 17. Its elements include VAX processors, HSC storage controllers, and the computer interconnect (CI). The latter is a high-speed interconnect (dual path connections, 70 Mbit/s each), similar in operation to an Ethernet, although the detailed methods for media access are somewhat different. Physically, the CI is organized as a star network, but appears to processors as though it were a simple broadcast bus like the Ethernet. Up to 16 nodes can be interconnected by a single star coupler, with each link being no more than 45 m in length. A processor is connected to the CI via a CI port, a collection of hardware and software that provides the physical connection to the CI on one side and a high-level queue-based interface to client software on the other side.

The communications protocols layered onto the CI and CI ports support three methods of transmission: datagrams, messages, and blocks. Datagrams are short transmissions meant to be used for status and information requests, and are not guaranteed to be delivered. Messages are similar to datagrams except that delivery is guaranteed. Read/write requests and other device control transmissions to storage controllers are handled via messages. The hardware in the CI ports provides special support for block transfers: an ability to copy sequential large blocks of data from the virtual address space of a process on one processor to the virtual address space of another process on another CI node. Block transfers are exploited to move data back and forth between client nodes and the storage controllers.

An interesting aspect of the VAXCluster architecture is its support for a mass storage control protocol (MSCP), through which clients request storage services from storage controllers attached to the CI. A message-based approach has several advantages in the distributed environment embodied by the cluster concept. First, data sharing is simplified, since storage controllers can extract requests






Fig. 17. VAXcluster block diagram. A VAXcluster consists of client processors (VAX), server storage controllers (HSC), a high-speed interconnect (CI), adapters (CI port), and coupling hardware (Star Coupler). A message-oriented protocol is layered onto the interconnect hardware to implement client/server access to storage services.

from message queues and service them in any order they choose. Second, the protocols enforce a high degree of device independence, thus making it easier to incorporate new devices into the storage system without a substantial rewrite of existing software. Finally, the decoupling of a request from its servicing allows the storage controllers to apply sophisticated methods for optimizing I/O performance, including rearranging requests and breaking large requests into fragments that can be processed independently.

2) *HSC-70 Internal Organization* The internal organization of an HSC is shown in Fig. 18. The HSC was originally designed in the late 1970's, and has been in service for a decade. Its internal architecture was determined by the technology limits of that time. Nevertheless, there are a number of notable aspects about its organization. An HSC is actually a heterogeneous multiprocessor, with individual processors dedicated to specific functions. The three major subsystems are (1) the host interface, (2) the I/O control processor, and (3) the I/O device controllers. They communicate via shared control and data memories accessed via a control and data bus respectively.

The host interface, called a K.CI, is responsible for managing the transfer of messages over the CI bus. The hardware is based on an AMD bit-slice processor. The device controllers, called K.SDI's for disk interfaces and K.STI's for tape interfaces, use the same bit-slice processor. They implement device-specific read and write operations, as well as format, status, and seek operations (for disk) over Digital's proprietary device interfaces. Up to four devices can be controlled by a single K.X device controller, and up to eight K.X controllers can be attached to an HSC, for a total of 24 devices. All policy decisions are handled by the I/O control processor, which is based on a microprocessor implementation of the PDP-11 [35].




Fig. 18. HSC internal architecture. The host interface is managed by a dedicated bit slice processor called the K.CI. Devices are attached to K.SDI (disk) and K.STI (tape) device controllers. High-level control is performed by the P.I.O., a PDP-11 microprocessor programmed to coordinate the activities of the device controllers and the host interface.

The shared memory subsystem of the HSC plays a critical role in its ability to sustain I/O traffic. Private memories deliver instructions and data to the various processors, keeping them off the shared memory buses. Data structures used for interprocessor communications are located in the control memory. The control memory and bus support interlocked operation, making it possible to implement an atomic two-cycle read-modify-write. Data moving between I/O devices and the computer interconnect must be staged through the data memory. The sizes of both memories are rather modest by today's standards: 256K bytes in each.

The performance bottlenecks within the HSC come from two primary sources: bus contention and processor contention [36]. We examine bus contention first. Internal bus contention affects the maximum data rate that the controller can support. The controller's transfer bandwidth (Mbytes/s) is limited by its memory architecture and the implementation of the CI interface, both on the controller and on a processor with which it communicates. Because data must traverse the memory bus twice, the effective internal bandwidth to I/O devices is limited to 6.6 Mbyte/s. For example, on a device read, data must be staged from the device controller to the data memory over the memory bus, and then transferred once again over the bus to the CI interface. The HSC's software includes mechanisms for accounting for the amount of internal bandwidth that has been allocated to outstanding I/O requests. It will throttle I/O activity by delaying some requests if it detects saturation.

While this may appear to be a limitation, a more serious restriction is imposed by the HSC's CI interface itself. In general, it is designed to sustain on the order of 2 Mbyte/s. For some low-end members of the VAX family, even this may exceed the bandwidth of the host's CI interface. To avoid overrunning a host, and thus limit CI bandwidth wasted on retransmissions, the CI interface will only transmit a single buffer to a given client as long as there are buffers waiting to be sent to it.

Next we turn to processor contention. This is due to some extent to the design of the SDI disk interfaces. Each disk has a dedicated control bus, but a single data bus is shared among the devices attached to a controller. Thus,

high data bandwidth can be sustained by spreading disks among as many controllers as possible. For example, two disks on a single K.SDI will transfer fewer bytes per second than a configuration with one disk on each of two disk controllers.

3) *Typical I/O Operation Sequencing:* To understand the flow of data and processing through the HSC, we shall examine the processing steps of a typical disk read operation. The steps that we outline next are described at a high level. Considerably more detail, including the detailed data structures used, can be found in [35].

The first step is the arrival of the MSCP command over the CI bus. The message is placed in a K.CI reception buffer, where it is checked for well-formedness and validity. If it passes these checks, it is copied to a special data structure in the control memory, and pointers to this data structure are placed on a queue of work for the I/O policy processor.

The next step involves the execution of MSCP server software on the policy processor. The software is structured as a process that wakes up whenever there are pending requests in the work queue. The software examines the queue of commands, choosing the next one to execute based on the currently executing commands. It constructs a data structure that maps the MSCP command into physical disk operations, such as the disk seek command and a sequence of sector transfer requests. The original MSCP command message is modified to become a response message. The last phase is to place a pointer to the disk request data structure on a K.SDI work queue, to be found in the control memory.

The next thing that happens is the disk portion of the transfer. The K.SDI firmware reads the request on its work queue, extracts the seek command, and issues it to the appropriate drive. When the drive is ready to transfer, it indicates its status to the K.SDI. At this point the disk controller allocates buffers in the data memory, and stages the data as it comes in from disk to these buffers. When the list of sector transfers is complete, a completion message is placed on the work queue for the K.CI.

We are now ready to transfer the data from the controller's data memory to the host. The K.CI software wakes up when new work appears in its queue. It then generates the necessary CI message packets to transfer data from data memory out over the CI to the originally requesting host processor. As a data buffer is emptied, it is returned to a list of free buffers maintained in the control memory. When the last buffer of the transmission has been sent, the K.CI now transmits the MSCP completion message that was built by the I/O policy processor from the original request message.

The steps outlined above have assumed that processing continues without error. There are a number of error recovery routines that may be invoked at various points in the process described above. For example, if the transfer request within a K.SDI fails, the software is structured to route the request to error handling software to make the decision whether to retry or abort the request.

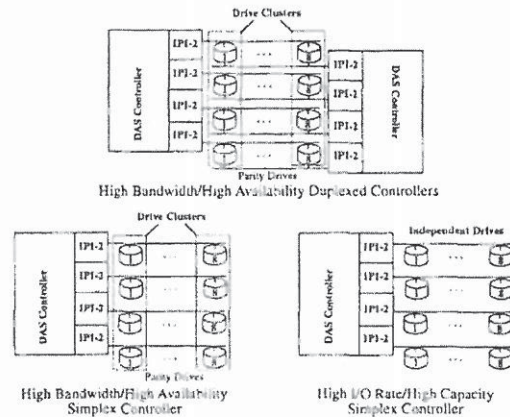


Fig. 19. Alternative disk organizations for the CDC DAS. The DAS can be configured in three alternative organizations: duplexed controller parity array for extremely high system availability, simplex parity array for high media availability, and a nonredundant organization for maximum I/O rate and capacity.

### C. Control Data Corporation Disk Array Subsystem (DAS)

1) *General Organization:* The CDC disk array subsystem is an example of an I/O controller design targeted for high-bandwidth environments. To this end, it supports multiple (4) IPI-2 interfaces to disks, which can burst at 10 Mbyte/s each, and multiple (2-4) IPI-1 interfaces to the host, each of which can burst transfer at 25 Mbyte/s. (The IPI interface is similar in concept to the SCSI protocols described in subsection II-C, but provide higher performance though they are more expensive to implement.) A single controller can handle up to 32 disk drives, organized into eight stripe units (called drive clusters by CDC) of four disks each.

The controller supports three alternative disk organizations: a high transfer rate/high availability mode with duplexed controllers, a simplex version of this organization, and a high transaction rate/high capacity organization. These are summarized in Fig. 19. The first organization, for high transfer rate and very high system availability, is distinguished by duplexed controllers, dual ported and spindle synchronized disk drives, and a 3 + 1 RAID level 3 parity scheme.

Enhanced system availability is achieved by the duplexed controllers and dual ported drives: if a single controller fails, then a path still exists from the host to a device through a functional controller. CDC claims 99.999% availability with a mean time to data loss (MTTDL) that exceeds 1 million hours with this configuration. This is probably a conservative estimate. A lost drive can be reconstructed within four minutes, assuming 1 Gbyte Seagate Sabre disk drives. The organization can sustain 36 Mbyte/s, assuming a sustained transfer bandwidth of 6 Mbyte/s per drive.

The second organization is characterized by single ported drives, organized into the RAID level 3 scheme, and a single controller. System availability is not as good as the previous organization: a failure within the controller renders the disk subsystem unavailable. However, media

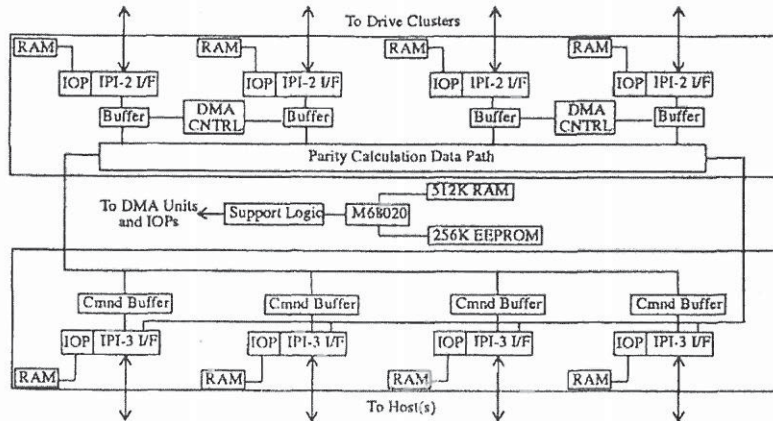


Fig. 20. Internal organization of CDC disk array controller. The DAC's internal structure consists of the disk interfaces, host interfaces, parity calculation logic, and a "traffic cop" microprocessor to determine the I/O strategy. I/O processors associated with each of the interfaces handle the low-level details of the interface protocols. Data movement is controlled by direct memory access engines associated with the disk interfaces.

availability is just as good because of the parity encoding scheme. This organization can sustain 18 Mbytes/s and also claims a 1 million hour MTDL.

The last organization represents a trade-off between performance, availability, and capacity. It gains capacity by dispensing with the parity drives, supporting a maximum of 32 Gbytes versus 24 Gbytes in the other two organizations, assuming 1 Gbyte drives. However, there is also no protection against data loss in the case of a disk crash. Data are no longer interleaved, thus sacrificing data bandwidth for a higher I/O rate. In the previous organizations, up to 8 I/O's can be in progress at the same time, one for each drive cluster. In this organization, 32 I/O's can simultaneously be in progress. The controller supports 500 random I/O's per second, approximately 16 I/O's per second per disk drive (this represents a disk utilization of 50%).

The controller's designers have placed considerable emphasis on providing support for very high data integrity within the controller and disk system. All internal data paths are protected by parity, data are written to disk with an enhanced ECC coding scheme (a 96 bit Reed-Solomon code that can correct up to 17 bit errors and even some 32 bit errors), and a large number of retries are attempted in the event of an I/O failure (three attempts at normal offset, all with ECC; three attempts at late and early data strobes with nominal carriage offset, all with ECC; three attempts at +/- carriage offset with nominal data strobes, all with ECC). All retries include possible correction from the RAID level 3 redundancy schemes.

2) *Controller Internal Organization:* Figure 20 shows the internal organization of the DAS controller. An I/O request can be traced as follows. The host issues the appropriate command to one of the IPI-3 interfaces. This is staged to a command buffer within the controller. The central control

microprocessor examines the command and determines how to implement it in detail.

Suppose that the command is a data write and that the array is organized into a RAID level 3 scheme. The control processor maps this logical write request into a stream of physical writes to the disks within a drive cluster. As the data stream across the host interface, they pass through the parity calculation data path, where the horizontal parity is computed. DMA controllers move data and parity from this data path to buffers associated with individual disk interfaces. I/O processors local to the IPI-2 disk interfaces manage the details of staging data from the buffers to particular disk drives. Read operations are performed in much the same manner, but in reverse.

Note that reconstruction operations can be performed without host intervention. Assume that the failed disk has been replaced by a new one. Under the control of the central microprocessor, data are read from the surviving members of the drive cluster. The data are streamed through the parity calculation data path, with the result being directed to the disk interface associated with the failed disk. The reconstituted data are then written to their replacement.

#### D. Maximum Strategies HIPPI-2 Array Controller

Maximum Strategies offers a family of storage products oriented toward scientific visualization and data storage applications for high-performance computing environments. The products offer a trade-off between performance and capacity, spanning from high Mbyte/s but low Mbytes (based on parallel transfer disks) to high-performance/high-capacity (based on arrays of disk arrays). In the following discussion, we concentrate on their HIPPI-based storage server.




Fig. 21. Strategy HiPPI controller block diagram. The Strategy controller couples multiple HiPPI interfaces to an 8 + 1 + 1 RAID level 3 disk organization.

Figure 21 shows the basic configuration of the Strategy-2 array controller. It supports one or two 100 Mbyte/s host HiPPI interfaces or a single 200 Mbyte/s interface. The controller supports a RAID level 3 organization calculated over eight data disks and one parity disk. Optionally, hot spares can be configured into the array. This allows reconstruction to take place immediately, without needing to wait for a replacement disk. It also helps the system achieve an even higher level of availability. Since reconstruction is fast, the system becomes unavailable only when two disks have crashed within a short period.

The controller can be configured in a number of different ways, representing alternative trade-offs between performance and capacity. The low capacity/high performance configuration stripes its data across four parallel transfer disks. This yields 3.2 Gbytes of capacity and can reach a 60 Mbyte/s data transfer rate. It provides no special support for high availability, such as RAID parity.

A second organization stripes across 8 + 1 parallel transfer disks implementing a RAID level 3 organization. This organization provides 6.4 Gbytes and achieves a 120 Mbyte/s transfer rate. Both configurations are called the Strategy HiPPI-SM storage server.

These organizations provide relatively little capacity for the level of performance provided. In addition, parallel transfer disks are quite expensive per Mbyte and have a poor reputation for reliability. An alternative configuration uses multiple ranks of 8 + 1 + 1 commodity disk drives. Maximum Strategies' HiPPI-S2 storage server is shown in Fig. 22. Back-end controllers (called S2's in Maximum Strategies' terminology) manage strings of eight disks each. Maximum Strategies makes use of older technology ESDI drives (5.25 in form factor, 1.2 Gbyte capacity each), which can share a common control path but require dedicated data paths. A maximum configuration can support ten of these: eight data strings, one parity string, and an optional hot spare string. The back ends are connected to the front end HiPPI interfaces through a 250 Mbyte/s data backplane and a conventional VME backplane used for control. Note the separation of control and data path. The high-bandwidth data transfer path is over HiPPI; the control path uses a lower latency (and lower bandwidth) VME interconnect. Parity calculations are handled in the front end. This organization can provide a 300 Gbyte capacity and a 144

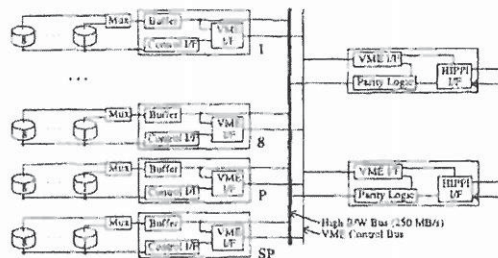


Fig. 22. High capacity strategy array. High capacity is achieved by using large numbers of commodity disk drives. These are coupled to the HiPPI front ends through a high-bandwidth data bus and a VME-based control bus.

Mbyte/s transfer rate.

Maximum Strategies also provides a storage server based on a VME-based host interface. The S2R storage server supports up to  $40 \times 5.25$  in ESDI drives, organized into an 8 + 1 + 1 RAID level 3 scheme that is four stripe units deep. This organization yields 38.4 Gbytes of capacity and an 18 Mbyte/s transfer rate.

The highest capacity/highest performance system combines the S2R-based arrays with the HiPPI-attached controller of Fig. 21. The result is an "array of disk arrays." The architecture calls for replacing the S2 controllers with S2R disk controllers. Each S2R array contains 37 disks, organized into four stripe units of eight data disks and one parity disk, plus one spare for the entire subarray. Up to ten subarrays can be controlled by a single HiPPI controller, yielding a system configured from a total of 370 disk drives, a 345 Gbyte data capacity, and a 144 Mbyte/s transfer rate.

#### E. AUSPEX NS5000 File Server

1) *General Overview:* AUSPEX has developed a special hardware and software architecture specifically for providing very high performance NFS file service. The system provides a file system function integrated with an ability to bridge multiple local area networks. They claim to have achieved a performance level of 1000 NFS 8 Kbyte read I/O operations per second, compared with approximately 100-400 I/O operations per second for more conventional server architectures [37].

They call their approach functional multiprocessing. Rather than building a server around a single processor that must simultaneously run the UNIX operating system and manage the network and disk interfaces, their architecture incorporates dedicated processors to separately manage these functions. By running specialized software within the network, file, and storage processors, much of the normal overhead associated with the operating system can be eliminated.

A functional block diagram of the NS5000 appears in Fig. 23. The system backbone is an enhanced VME bus that has been tweaked to achieve a high aggregate bandwidth (55 Mbyte/s). A conventional UNIX host processor (a SUN-3 or SUN Sparcstation board), the various special-purpose




Fig. 23. NS5000 block diagram. The server incorporates four different kinds of processors, dedicated to network, file, storage, and general-purpose processing. The server can integrate up to eight independent Ethernets through the incorporation of multiple network processors. The storage processor supports ten SCSI channels, making it possible to attach up to 20 disks to the server.

processors, and up to 96 Mbytes of semiconductor memory (the primary memory) can be installed in the backplane. We examine each of the special processors in the next subsection.

2) *Dedicated Processors:* A dedicated network processor board contains the hardware and software needed to manage two independent Ethernet interfaces. Up to four of these can be incorporated into the server to integrate a reasonably large number of independent networks. The board executes all of the necessary protocol processing to implement the NFS standard. Because the network boards implement their own packet routing functions, it is possible to pass packets from one network to another without intervention by the host. Some cached network packet headers are buffered in the primary memory.

The file processor board runs dedicated file system software factored out of the standard UNIX operating system. The board incorporates a large cache memory, partitioned between user data and file system metadata, such as directories and inodes. This makes it possible for the file system code to access critical file system information without going to disks.

The storage processor manages ten SCSI channels. Disks are organized into four racks of five 5.25 in disks each (20 disks per server). It is also possible to organize these into a RAID-style disk array, although the currently released software does not support the RAID organization at this time. Most of the primary memory is used as a very large disk cache. Because of the way the system is organized, most of the memory system and backplane bandwidth is dedicated to supporting data transfers between the network and disk interfaces.

The host processor is either a standard SUN-3 68020-based processor board or a Sparcstation host processor board. These run the standard Sun Microsystems' UNIX, as well as the utilities and diagnostics associated with the rest of the system.

3) *Software Organization:* A significant portion of Auspex's improved performance comes from the way in which the network and file processing software are layered onto the multiprocessor organization described above. The basic software architecture, its mapping onto the processors, and their interactions are shown in Fig. 24.




Fig. 24. Auspex NS5000 software architecture. The main data flow is represented by the heavy black line, with data being transmitted from the disks to the primary memory to the network interface. The primary control flow is shown by a heavy gray line. File system requests are passed between LFS (local file system) client software on the Ethernet processor to server software on the file processor. These are mapped onto detailed requests to the storage processor by the file system server. Limited control interactions involve the virtual file system interfaces on the host, and are denoted by dashed lines.

Consider an NFS read operation. Initially, it arrives at an Ethernet processors, where the network details are handled. The actual data read request is forwarded to a file processor, where it is transformed into physical read requests, assuming that the request cannot be satisfied by cached data. The read request is passed to the storage processor, which turns it into the detailed operations to be executed by the disk drives. Retrieved data are transferred from the storage processor to primary memory, from which the Ethernet processor can construct data packets to be sent to the client. Note the minimal intervention from the host processor and software.

#### F. Berkeley RAID-II Disk Array File Server

Our research group at the University of California, Berkeley, is implementing a high-performance I/O controller architecture that connects a disk array to an UltraNet network via a HIPPI channel. We call it RAID-II to distinguish it from our first prototype, RAID-I, which was constructed from off-the-shelf controllers [38]. Given the observations about the critical performance bottlenecks in file server architectures throughout this paper, our controller has been specifically designed to provide considerable bandwidth between the network, disk, and memory interfaces.

A block diagram for the controller is shown in Fig. 25. The controller makes use of a two-board set from Thinking Machines Corporation (TMC) to provide the HIPPI channel interface to the UltraNet interfaces. The disk interfaces are provided by a VME-based multiple SCSI string board from Array Technologies Corporation (ATC). The major new element of the controller, designed by our group, is the X-bus board, a crossbar that connects the HIPPI boards, multiple VME buses, and an interleaved, multiported semiconductor memory. The X-bus board provides the high-bandwidth data path between the network and the disks. The data path is controlled by an external file server through a memory-mapped control register interface.

The X-bus board is organized as follows. The board implements an 8 by 8 32-bit-wide crossbar bus. All crossbar

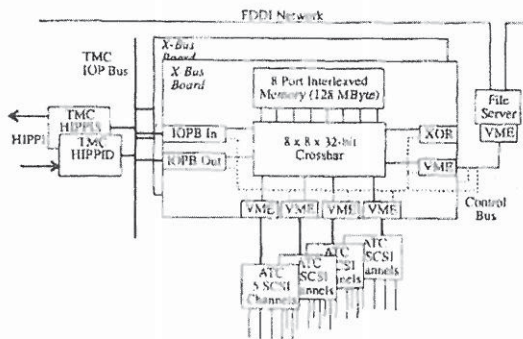


Fig. 25. RAID-II organization. A high-bandwidth crossbar interconnection ties the network interface (HIPPI) to the disk controllers (Array Tech) via a multiported memory system. Hardware to perform the parity calculation is associated with the memory system.

transfers involve the on-board memory as either the source or the destination of the transfer. The ports are designed to burst transfer at 50 Mbyte/s, and sustain transfers of 40 Mbyte/s. The crossbar is designed to provide an aggregate bandwidth of 320 Mbyte/s.

The controller memory is allocated eight of the crossbar ports. Data are interleaved across the eight banks in 32 word interleave units. Although the crossbar is designed to move large blocks from memory to or from the network and disk interfaces, it is still possible to access a single word when necessary. For example, the external file server can access the on-board memory through the X-bus board's VME control interface. Two of the remaining eight ports are dedicated as interfaces to the Thinking Machine I/O processor bus. The TMC HIPPI board set also interfaces to this bus. Since these X-bus ports are dedicated by their direction, the controller is limited to a sustained transfer rate to the network of 40 Mbyte/s.

Four more ports are used to couple to single board multistring disk controllers via the industry standard VME bus, one disk controller per VME bus. Because of the physical packaging of the array, 15 disks can be attached to each of these, in three stripe units of five disks each. Thus, 60 disk drives can be connected to each X-bus board, and a two X-bus board configuration consists of 120 disk drives.

Of the remaining two ports, one is dedicated for special hardware to compute the horizontal parity for the disk array. The last port links the X-bus board to the external file server. It provides access to the on-board memory—as well as the board's control registers (through the board's control bus). This makes it possible for file server software, running off of the controller, to access network headers and file metadata in the controller cache.

It may seem strange that there is no processor within the X-bus board. Actually, the configuration of Fig. 25 contains no less than seven microprocessors: one in each of the HIPPI interface boards, one in each of the ATC boards, and one in the file server (we are also investigating multiprocessor file server organizations). The processors

within the HIPPI boards are being used to handle some of the network processing normally performed within the server. The processors within the disk interfaces handle the low-level details of managing the SCSI interfaces. The file server CPU must do most of the conventional file system processing. Since it is executing file server code, the file server needs access only to the file system metadata, not user data. This makes it possible to locate the file server cache within the X-bus board, close to the network and disk interfaces.

Since a single X-bus board is limited to 40 Mbyte/s, we are examining system organizations that interleave data transfers across multiple X-bus boards (as well as multiple file servers, each with its own HIPPI interface). Multiple X-bus boards can share a common HIPPI interface through the IOP bus. Two X-bus boards should be able to sustain 80 Mbyte/s, more fully utilizing the available bandwidth of the HIPPI interface.

The controller architecture described in this subsection should perform well for large data transfers that require high bandwidth. But it will not do so well for small transfers where latency dominates performance more than transfer bandwidth. Thus we are investigating organizations in which the file server remains attached to a more conventional network, such as FDDI. Requests for small files will be serviced over the lowest latency network available to the server. Only very large files will be transferred through the X-bus board and the UltraNet.

## VI. SUMMARY AND RESEARCH DIRECTIONS

In this paper, we have made the case for generalizing the workstation-server storage architecture to the mainframe and high-performance computing environment. The concept of network-based storage is very compelling. It has been said that the difference between a workstation and a mainframe is the I/O system. The distinction will become blurred in the new system architectures made possible by high-bandwidth, low-latency networks coupled to the correct use of caching and buffering throughout the path from service requestor to service provider.

Nevertheless, many research challenges remain before this vision of ubiquitous network-based storage can be achieved. First, new methods are needed to effectively manage the complete and complex storage hierarchy as described in this paper. How should data be staged from tertiary to secondary storage? What are the effective prefetching strategies? How are data to be extracted from such large storage systems?

Second, it is time to apply a system-level perspective to storage system design. Throughout the I/O path, from host to embedded disk controller, we find buffer memories and processing capabilities. The current partitioning of functions may not be correct for future high-performance systems. For example, some searching and filtering capabilities could be migrated from applications into the devices. The memory in the I/O path could be better organized as caches rather than speed matching buffers, given enough

local intelligence about I/O patterns. A better approach for error handling is also possible given a system perspective. For example, in response to a device read error, a disk array controller could choose between retrying the read or exploiting horizontal parity techniques to reconstitute the data on the fly.

Third, new architectures are needed to break the bottlenecks, both hardware and software, between the network, memory, and I/O interfaces. The RAID-II controller tackles this at the hardware level, by providing a high-bandwidth interconnection among these components. At the software level, new methods need to be developed to reduce the amount of copying and memory remapping currently required for controlling these interfaces.

Fourth, today's high-bandwidth networks, such as FDDI and UltraNet, exhibit latencies that are somewhat worse than conventional Ethernets. Unfortunately, latency becomes a dominating factor as the overheads of data transfer scale down in higher bandwidth networks. New methods need to be developed to reduce this latency. One strategy is to increase the packet sizes, to better amortize the start-up latencies. A second strategy, demonstrated by the Autonet project at Digital Equipment Corporation's System Research Laboratory, is to construct a high-bandwidth network using point-to-point connections and an active switching network [39].


Finally, the whole issue of distributed and multiprocessor file/storage servers and their role in high-performance storage systems must be addressed. The technical issues include the methods for how to partition the file server software functions among the processors of a multiprocessor or a distributed collection of processors. The AUSPEX controller architecture is one approach to the former. The IEEE Mass Storage System Reference offers one model for the latter.

#### ACKNOWLEDGMENT

The author appreciates the careful reading of this manuscript and the detailed comments by P. Chen, A. Chervenak, E. Lee, E. Miller, S. Seshen, S. Strange, and the referees.

#### REFERENCES

- [1] *Network Operations Manual*, UltraNetwork Technologies, Part Number 06-0001-001, Revision A, 1990, chs. 2 and 3.
- [2] J. Hennessy and D. Patterson, *Computer Architecture: A Quantitative Approach*. San Mateo, CA: Morgan Kaufmann, 1990.
- [3] V. Cerf, "Networks," *Scientific American*, vol. 265, no. 3, pp. 72-85, Sept. 1991.
- [4] L. Tesla, "Networked computing in the 1990s," *Scientific American*, vol. 265, no. 3, pp. 86-93, Sept. 1991.
- [5] N. Negraponte, "Products and services for computer networks," *Scientific American*, vol. 265, no. 3, pp. 106-115, Sept. 1991.
- [6] S. P. Joshi, "High performance networks: Focus on the fiber distributed data interface standard," *IEEE Micro*, pp. 8-14, June 1986.
- [7] D. Clark, V. Jacobson, J. Romkey, and H. Salwen, "An analysis of TCP processing overhead," *IEEE Communications Magazine*, pp. 23-29, June 1989.
- [8] S. Heatly and D. Stokesberry, "Analysis of transport measurements over a local area network," *IEEE Communications Magazine*, pp. 16-22, June 1989.
- [9] H. Kanakia, "High performance host interfacing for packet-switched networks," Ph.D. dissertation, Department of E.E.C.S., Stanford University, 1990.
- [10] R. Watson and S. Mamrak, "Gaining efficiency in transport services by appropriate design and implementation choices," *ACM Trans. Comput. Syst.*, vol. 5, no. 2, pp. 97-120, May 1987.
- [11] E. Orenstein, "HPP1-based storage system," *Computer Technol. Rev.*, Apr. 1990.
- [12] Anon, *Strategy HPP1 Disk Array Subsystem Operation Manual*, Maximum Strategies, Part Number #HPP100, 1990.
- [13] D. E. Tolmie and M. G. Halvorsen, "HIPPI / serial-HIPPI," in *Proc. IEEE Spring Compton Conf.*, Feb. 1992, pp. 222-228.
- [14] W. McFarland et al., "HP's link interface chipset for serial-HIPPI," in *Proc. IEEE Spring Compton Conf.*, Feb. 1992, pp. 229-233.
- [15] Anon, "Fiber channel - Physical layer (FC-PH)," ANSI X3T9.3 Working Document, Revision 2.1, May 1991.
- [16] K. Hardwick, "HIPPI world - The switch is the network," in *Proc. IEEE Spring Comput. Conf.* (San Francisco, CA), Feb. 1992, pp. 234-238.
- [17] P. Borrill and J. Theus, "An advanced communications protocol for the proposed IEEE 896 FutureBus," *IEEE Micro*, pp. 42-56, Aug. 1984.
- [18] M. Teener, "A Bus on a diet - The serial bus alternative," in *Proc. IEEE Spring Comput. Conf.* (San Francisco, CA), Feb. 1992, pp. 316-321.
- [19] D. Gustavson and E. Kristiansen, "Scalable coherent interface: Links to the future," in *Proc. IEEE Spring Comput. Conf.* (San Francisco, CA), Feb. 1992, pp. 322-327.
- [20] M. Nelson, J. K. Ousterhout, and B. Welch, "Caching in the Sprite network file system," *ACM Trans. Comput. Syst.*, vol. 6, no. 1, pp. 134-154, Feb. 1988.
- [21] R. Katz, G. Gibson, and D. Patterson, "Disk system architectures for high performance computing," *Proc. IEEE* (Special Issue on Supercomputing), Dec. 1989.
- [22] S. Ranade and J. Ng, *Systems Integration for Write-Once Optical Storage*. Westport, CT: Meckler, 1990.
- [23] M. H. Kryder, "Data storage in 2000 - Trends in data storage technologies," *IEEE Trans. Magn.*, vol. 25, pp. 4358-4363, Nov. 1989.
- [24] Hewlett-Packard Corporation, "HP Series 6300 Model 20GB/A rewritable optical disk library system product brief," 1989.
- [25] Kodak Corporation, "Optical disk System 6800 product description," 1990.
- [26] Exabyte Corporation, "EXB-120 cartridge handling subsystem product specification," Part No. 510 300&-002, 1990.
- [27] E. Tan and B. Vermeulen, "Digital audio tape for data storage," *IEEE Spectrum*, vol. 26, pp. 34-38, Oct. 1989.
- [28] B. F. Feder, "The best of tapes and disks," *N. Y. Times* (Sunday Business Section), p. 9, Sept. 1, 1991.
- [29] K. Spencer, "The 60-second terabyte," *Canadian Res. Mag.*, June 1988.
- [30] R. Sandberg, D. Goldberg, S. Keiman, D. Walsh, and B. Lyon, "Design and implementation of the SUN Network Filesystem," in *Proc. USENIX Summer Conf.*, June 1985, pp. 119-130.
- [31] M. Rosenblum and J. Ousterhout, "The design and implementation of a log-structured file system," *ACM Trans. Comput. Syst.*, Feb. 1992.
- [32] S. W. Miller, "A reference model for mass storage systems," *Advances in Computers*, vol. 27, pp. 157-210, 1988.
- [33] N. P. Kronenberg, H. Levy, and W. D. Strecker, "VAXclusters: A closely coupled distributed system," *ACM Trans. Comput. Syst.*, vol. 4, no. 2, pp. 130-146, May 1986.
- [34] N. P. Kronenberg, H. M. Levy, W. D. Strecker, and R. J. Merewood, "The VAXcluster concept: An overview of a distributed system," *Digital Tech. J.*, vol. 5, pp. 7-21, Sept. 1987.
- [35] R. L. Lary and R. G. Bean, "The hierarchical storage controller: A tightly coupled multiprocessor as storage server," *Digital Tech. J.*, vol. 8, pp. 8-24, Feb. 1989.
- [36] K. H. Bates, "Performance aspects of the HSC controller," *Digital Tech. J.*, vol. 8, pp. 25-37, Feb. 1989.
- [37] B. Nelson, "An overview of functional multiprocessing for NFS network servers," AUSPEX Tech. Rep. 1, July 1990.
- [38] A. Chervenak and R. H. Katz, "Performance measurements of a disk array prototype," presented at ACM SIGMETRICS Conf., San Diego, CA, May 1990.
- [39] M. Schroeder et al., "Autonet: A high-speed self-configuring local area network using point-to-point links," DEC SRC Tech. Rep. 59, Apr. 1990.
- [40] A. Chervenak, "Performance measurements of the first RAID prototype," U.C. Berkeley Computer Science Division Report No. UCB/CSD 90/574, Jan. 1990.



**Randy H. Kutz** (Senior Member, IEEE) was born in Brooklyn, NY, on August 19, 1955. From 1973 through 1976 he attended Cornell University, receiving a A.B. degree in mathematics and computer science, and graduating Phi Beta Kappa with distinction in all subjects. He was a graduate student at the University of California, Berkeley, receiving the M.S. and Ph.D. degrees in computer science in 1978 and 1980 respectively. He held an IBM predoctoral fellowship.

In the period 1980–1981 he held research scientist positions at Bolt, Beranek, and Newman, Inc., and the Computer Corporation of America, Inc., in Cambridge, MA. From 1981 to 1983 he was an assistant professor in the Computer Sciences Department at the University of Wisconsin-Madison. In 1983 he joined the faculty of the University of California, Berkeley, as an assistant professor. He was promoted to associate professor in 1985 and to full professor in 1989.

Prof. Katz was awarded an IBM Faculty Development Award and a National Science Foundation Presidential Young Investigator Award in 1984. He serves on the advisory panel of the Microelectronics and Information Processing (M.I.P.S.) division of the Computer and Information Sciences (C.I.S.E.) directorate of the National Science Foundation. He has won three best paper awards and four best presentation awards at conferences in his technical field. He has published over 80 papers in the fields of database management, computer-aided design of electronics systems, parallel computer architectures, and high-performance mass storage systems. He is principal investigator of a NASA and DARPA sponsored contract to build a prototype high-performance disk subsystem for teraop computers. His current research interests include I/O controller design and high-performance striped disk and tape subsystems. He is a faculty investigator on the Sequoia 2000 Project, an effort to apply advanced storage technologies in support of Global Change Researchers throughout the state of California. In addition, he has served as a consultant to Intermetrics, the U.S. Air Force, Xerox, Texas Instruments, Hambrecht and Quist, and numerous small high technology firms. Prof. Katz is a member of the Association of Computing Machinery.