Considerations in the Design of a RAID Prototype

Martin E. Schulze

Computer Science Division
Department of Electrical Engineering and Computer Sciences
University of California at Berkeley
Berkeley, CA 94720

ABSTRACT

Disk drive arrays are receiving attention from computer researchers looking for higher performance mass storage. RAID (Redundant Array of Inexpensive Disks) appears to offer benefits over storage systems based on large format disks. This paper investigates disk drive and disk support hardware failures as factors in the data integrity of RAID. Through the use of redundancy in the disk array, support hardware failures can be eliminated as a major factor in the data integrity of RAID. This paper also presents some practical considerations in the design of a RAID prototype and discusses how such a prototype could be constructed at Berkeley in the near future. The prototype will most likely be based on 5.25 inch disk drives and the Small Computer System Interface (SCSI) bus.

August 25, 1988

Table of Contents

Considerations in the Design of a RAID Prototype

Martin E. Schulze

Computer Science Division
Department of Electrical Engineering and Computer Sciences
University of California at Berkeley
Berkeley, CA 94720

1. Introduction

To keep pace with ever increasing computer speeds, computer designers are searching for ways to increase the performance of disk drive based mass storage systems. For example, in supercomputing it is common practice to stripe data files across multiple read/write heads to increase bandwidth during reads or writes. In transaction processing, it is desirable to spread a database across multiple actuators to allow many concurrent accesses even if overall bandwidth is not very high. Disk drive arrays can accommodate both these needs and are currently receiving attention from computer researchers.

Arrays of small drives, in particular, are receiving lots of attention. There are several reasons for this [GGI 88], [Vasudeva 88]. Data storage cost per megabyte is now less for small drives (5.25 inch and 3.5 inch) than for larger format drives (14 inch and 10.5 inch). Small drives offer better efficiency in terms of volume (MB/cu. ft.) than larger format drives. This volumetric efficiency can be translated into a smaller footprint (MB/sq. ft.), which is a measure of the floorspace required per unit of storage. Small drives also offer better efficiency in terms of power (MB/Watt) than larger format drives. These metrics are important to customers faced with floorspace and air conditioning constraints in their machine rooms.

One hindrance to using large numbers of disks to improve I/O performance is the impact on reliability and data integrity. More disks mean more disk failures and an increase in the possibility of data loss. This can be addressed by including redundancy as a feature of the disk array. The acronyms **RAID** (Redundant Array of Inexpensive Disks) and **SLED** (Single Large Expensive Disk) were introduced by [Patterson 88]. I will call this work the *RAID paper*. The RAID paper presented a taxonomy of five levels of RAID and showed RAID's potential for large improvements in performance and data integrity when compared to a SLED based mass storage system. I will briefly review the various RAID levels.

RAID level 1 is mirroring, where identical information is stored on a group of two disks. This scheme has 100 percent redundancy overhead. RAID level 2 has more data disks per group with multiple check disks arranged in Hamming Code fashion to identify and recover a disk failure. This allows redundancy overhead to be reduced to 20 - 40 percent. RAID level 3 recognizes that disk failures identify themselves, either by the intelligence in the disk controller or by the ECC (Error Correction Code) written at the end of each sector to guard against data errors. Only one check disk per group is needed and the data on this disk is just parity information for the data disks in each group. The redundancy overhead is further reduced to 5 - 20 percent. RAID level 4 retains the data integrity characteristics of RAID level 3 but allows independent reads and writes to individual drives within a group. Performance is improved while keeping the same redundancy overhead. RAID level 5 retains the data integrity characteristics of RAID level 3 and RAID level 4 but improves performance by distributing (or rotating) parity information over all drives within a group.

Commercial interest in the area of redundant disk arrays is growing. Many companies use mirroring (RAID level 1). Tandem Computers, for example, implements all disks this way. Thinking Machines has a RAID level 2 for its Connection Machine. Micropolis' 1804 is a RAID level 3 based on 5.25 inch drives. Pacstor's Integra III appears to be a RAID level 4 or 5 and is based on 3.5 inch drives. RAID level 5 appears to be the most interesting option for further study.

This report contains two parts: a more complete analysis of RAID data integrity than the coverage given in the RAID paper and [Vasudeva 88]; and a section covering some practical considerations in the design and implementation of a SCSI (Small Computer System Interface) based RAID prototype. It does not address or evaluate RAID performance in any detail.

2. RAID Data Integrity

This section of the report begins with some definitions, moves on to discuss disk drive reliability and the process of how it is evaluated, and finishes by examining RAID data integrity.

2.1. Definitions

At this point, some definitions are necessary in order to clearly define the terminology that will be used in this report. Much of this terminology is adopted from [Maxion 88], [Quantum 87], and [Siewiorek 82].

2.1.1. Failures, faults, and errors

A *failure* is a detectable physical change to hardware. Failures may be repaired by the replacement of a physical component. A *fault* is an event which interferes with normal operation and can be either *soft (transient)* - not readily repeatable, or *hard* - repeatable with high probability. Hard faults may be caused by failures, while soft faults are more likely caused by environmental factors or insufficient design margins. An *error* is a manifestation of a fault by an incorrect value. Errors, therefore, can be either soft or hard. For example, when discussing the performance of magnetic mass storage devices, *soft error rate* and *hard error rate* are terms commonly used to describe the frequency at which data errors occur.

2.1.2. Reliability and availability

System reliability $R(t)$ is defined as the conditional probability that the system will not fail during the interval $[0, t]$, given that it is operational at time $t = 0$. Often, constant failure rate is assumed and the reliability is expressed in terms of Mean Time To Failure (MTTF).

System availability $A(t)$ is the probability that the system is available at the instant of time t. Averaged over time, this function expresses the expected fraction of time that the system is available to do useful work. Availability is usually expressed as a percentage (hopefully close to 100%).

There are two techniques employed to achieve higher system reliability and availability:

fault avoidance and *fault tolerance*. Fault avoidance aims to minimize faults by the use of high reliability components and by the use of conservative design practices such as careful signal routing, low system operating temperatures, and low component load factor designs. The goal of fault avoidance is to reduce the possibility of a fault. In contrast to this, fault tolerance aims to negate the harmful effects of faults when they occur. This is done with redundancy, often in the form of extra hardware. Without fault tolerance, the failure of a component on a module may cause the failure of the module, which may cause the failure of the entire system. With fault tolerance (in the form of redundant modules, for example), the failure of a component on a module may cause the failure of the module but (hopefully) will not cause the failure of the entire system.

When considering the reliability of a mass storage system, two terms are necessary: *data integrity* and *hardware reliability*. Data integrity is the overall goal; hardware reliability is one factor in how this goal is achieved. Data integrity can be expressed in terms of Mean Time Between Data Loss (MTBDL) and hardware reliability can be expressed in terms of Mean Time To Failure (MTTF). The two terms should be considered separately because data integrity considers certain types of errors (e.g. magnetic media defects) that are not considered under the heading of hardware reliability.

2.2. Disk drive reliability and data integrity

Data integrity and therefore disk reliability are of prime importance to disk drive customers. Many small disk drives offer reliabilities specified as MTTF = 30,000 to 40,000 hours of normal usage. To help define normal usage, disk manufacturers specify operating lifetime (usually 5 years), after which the product should be removed from service. Manufacturers also specify maximum allowable spindle start/stop cycles (usually 10,000) for their products. These MTTF figures cover only hard faults caused by component failures; other types of faults are excluded. Reliability Research Inc., which monitors the reliability of products in the IBM compatible mainframe market, reports disk drive reliabilities as Millions of Start I/Os Between Failures. This reliability metric expresses the amount of useful work (each read or write command is considered a

Start I/O) a disk drive performs between failures, and appears to be a worthwhile metric in addition to MTTF; but has yet to be adopted by disk manufacturers.

Disk drives have error rates associated with the servo system (positioning the heads) and the read/write system (reading and writing data). These error rates are normally specified by the manufacturer as shown in Table I. The error rates in Table I are from [CDC 88] and [Quantum 87] and are typical of disk products available today. These error rates apply over all specified operating conditions except shock and vibration and assume that the magnetic media is properly flaw mapped during initial formatting.

Disk Drive Error Rates			
Type of Error	Error Rate	Recovery	Long Term Consequences
Recoverable Seek Error	< 1 error in 10^6 seeks	retry	none
Random Recoverable Data Error (Soft Error: ECC recoverable)	< 1 error in 10^{10} bits read	retry or ECC	none
Repeatable Recoverable Data Error (Hard Error: ECC recoverable)	< 1 error in 10^{12} bits read	retry or ECC	Sector may be reallocated. Data is rewritten to new sector.
Unrecoverable Data Error (Hard Error: unrecoverable)	< 1 error in 10^{14} bits read	none	Data in one sector is lost. Sector may be reallocated.
Miscorrected Data Error	< 1 error in 10^{21} bits read	none	One sector of incorrect data is returned to host.

Table I: Disk Drive Error Rates

A recoverable seek error is a seek in which the drive does not locate the desired cylinder on the first try but is successful during retry operations. The rest of the error types are all data errors associated with reading. A data error is defined as one sector read incorrectly, as determined by the Error Correcting Code (ECC). Since disk drives use the same head for both writing and reading, it is not possible to verify a write operation as it happens. For this reason, there are no data errors during writes. Random recoverable data errors are those which do not exhibit a repeating error pattern. The data is recovered by rereading or applying ECC correction. These are soft errors and are generally related to the signal-to-noise ratio of the system. Repeatable recoverable errors are those which exhibit a repeating error pattern on retry reads. The data is recovered by rereading or applying ECC correction. These are hard errors and are most likely due to media defects. As such, the sector in question may be reallocated and the data rewritten to the new location. Unrecoverable data errors are hard errors that are not recoverable by either retrying or by ECC. These errors result in the loss of one sector of data. In this case, the sector should be reallocated. Miscorrected data errors are those for which ECC correction has been performed but has resulted in incorrect data. This type of error is particularly harmful because the error is believed to be corrected, but one sector of incorrect data is returned to the host.

Now let's consider the frequency of errors listed in Table I. Take the case of a disk drive under a heavy workload: 50 seeks/sec and reading 512 KB/sec sustained. Using these numbers one can calculate a lower bound on the mean time to next error for: Recoverable seek error - 5.6 hours, Random recoverable data error - 40 minutes, Repeatable recoverable data error - 2.8 days, Unrecoverable data error - 276 days, and Miscorrected data error - 7.6 million years. None of the recoverable errors happen frequently enough to have a measurable impact on performance (because retry or recovery is a short process), so they appear to offer no threat to RAID. Losing one sector of data due to an unrecoverable data error every 276 days is terrible, but RAID can protect against this the same way it can protect against catastrophic drive failure. A miscorrected data error is the worst thing a mass storage device can do. This type of error should never happen.

It should be noted that all error rates in Table I are specified as upper limits. Many disk manufacturers design their products to internal error rate goals an order of magnitude better than the published specifications. A mature disk product based on proven technology will most likely exhibit soft and hard errors at much lower rates than the specified limits.

2.3. How disk drive reliability is evaluated

To insure data integrity, disk drive manufacturers must be concerned with the reliability of their products. At Digital Equipment Corporation's Storage Systems Division in Colorado, evaluating reliability of disk products is a multi-faceted process [Anderson 88]. This process is described below.

2.3.1. Reliability estimation

During product development, hardware reliability is estimated using an in-house software package based on the data and techniques of MIL-HDBK-217. The MIL-HDBK-217 model is based on electronic component failure data collected by the Department of Defense. The data are used to establish a mathematical model that estimates the frequency of hard faults caused by component failures; other types of faults are not addressed. In the MIL-HDBK-217 model, the reliability function is assumed to satisfy the equation $R(t) = e^{-\lambda t}$. This function assumes that the time to failure is distributed as an exponential random variable. This assumption is common in reliability analysis and has been verified for some electronic equipment in a study of failure data from the Cm* system at Carnegie Mellon University [Maxion 88].

Component failures are assumed to be random, independent events. The failure rate λ for individual integrated circuits is assumed to be constant with respect to time and is given by a formula that includes the following variables: a *learning factor* based on the maturity of the fabrication process, a *quality factor* based on incoming screening of components, a *temperature factor* based on ambient operating temperature, an *environmental factor* based on the operating environment, and several *complexity factors* based on the number of active devices and the number of pins. With the assumption of exponential $R(t)$ and constant failure rate λ, reliability can be expressed as MTTF $= \frac{1}{\lambda}$. Calculating system MTTF with the MIL-HDBK-217 model then reduces to summing failure rates for all components to get an overall failure rate; and taking the reciprocal of this number to find the MTTF.

2.3.2. Bringing a product to market

At Digital Colorado, a new product undergoes a number of testing phases before it is brought to market. These include Design Verification Test, Design Maturity Test, and Field Test.

The first phase is Design Verification Test (DVT), where several dozen prototype units are tested in various system configurations. The purpose of DVT is twofold: to verify that the design meets or exceeds the product specifications for functionality and performance; and to gather failure and error rate data to improve the design.

Once DVT is completed, the next phase is Design Maturity Test (DMT) where production units are subjected to extremes of voltage, temperature, and humidity. DMT tests the maturity of the design, the maturity of the manufacturing process, and verifies that the product operates correctly under all conditions. Again, a population of several dozen units is tested; failure and error rate data are collected and analyzed to point out design weaknesses. As many system configurations as possible are tested to insure that the new product is backward compatible with existing hardware and software. An estimation of product reliability is made from the failure data collected during DVT and DMT.

The next phase is internal and external Field Test, where a population of about 100 units is distributed to customers both inside and outside Digital. Digital's Customer Service Support Engineering (CSSE) organization monitors field test by weekly phone calls. The customers are asked to provide information on problems, failures, power-on hours, and on the type of workload that their field test units are experiencing. Based on this data, a report summarizing the behavior of field test units is issued every two weeks. Similar to DVT and DMT, reliability of field test units is estimated from the field failure data. Field Test is completed when previously agreed upon milestones are reached. These milestones are intended to insure that the product is ready for market and might be something like bug fix frequency less than one per week or outstanding failures less than 5 percent.

2.3.3. Insuring reliability during production

Once a Digital disk product is brought to market, several methods are used to track its reliability over time. These include Ongoing Reliability Test (ORT) and analysis of field failures.

ORT is essentially an ongoing small scale Design Maturity Test. Periodically, several newly manufactured units are tested under the same conditions that the initial DMT units were tested. Data on failures and errors are collected and analyzed. Although ORT is an expensive program to implement throughout the production lifetime of a product, it has demonstrated its worth by consistently predicting failure behavior of units in the field. Digital has found that field failure behavior lags ORT failure behavior by several months. This lag is nothing but the difference in time from when drives are manufactured to when they are first placed into service.

Along with ORT, analysis of field failures is performed based on data from written reports submitted by customer service engineers. To help improve availability, software tools are used to predict disk failures, so that corrective action can be taken prior to a failure, thus minimizing downtime. Maxion and Siewiorek state that symptoms of disk failure can be seen in system event logs up to two weeks prior to catastrophic failure.

2.4. How accurate are estimated reliability numbers?

MIL-HDBK-217 failure rate predictions tend to be pessimistic for new technologies [Anderson 88], [Maxion 88]. This is due to the time required to obtain sufficient data to calibrate failure models for new technologies and the frequency at which the models are updated. A new version of MIL-HDBK-217 is usually published every four or five years. In one study, it was found that MIL-HDBK-217B was a factor of 16 to 64 pessimistic in estimating failure rates of MOS RAMs and ROMs [Siewiorek 82]. The MIL-HDBK-217B model was published in 1974 and was probably developed with 1972 data. MOS was not a mature technology at the time the model was developed and this may account for the inaccuracy. The same study found that MIL-HDBK-217B was accurate to within a factor of 2 in estimating failure rates of TTL SSI and MSI parts. (TTL was a more mature technology than MOS at that time.) It appears that estimated reliabilities can be relatively accurate for mature technologies but may not be for new technologies. Concerns about absolute accuracy aside, the MIL-HDBK-217 models are certainly useful for making design decisions when comparing implementations of the same technology. MIL-HDBK-217E, dated October 1986, is the latest publication in this series [MIL 86].

2.5. RAID data integrity estimates

The redundancy of RAID is an application of fault tolerance to address the problem of data loss due to disk drive failures. Small disk drives such as the ones we are considering for RAID are not standalone units but require support hardware: power supplies, SCSI Host Bus Adapters (HBAs), cooling equipment, and cabling. To get a good picture of RAID data integrity all parts of the mass storage system should be considered. RAID is based on the concept of the *parity group*, a group of disks sharing a common parity check disk. When RAID is implemented with SCSI, a second type of grouping that emerges is the *SCSI group*, a group of disks sharing a common SCSI cable and HBA. There is also the *power group*, a group of disks sharing a common power supply and the *cooling group*, a group of disks sharing a common fan. The interaction of these groups is a major influence on RAID data integrity.

My analysis of RAID data integrity considers only hardware failures from components that are replicated many times within RAID. It ignores the RAID I/O processor that connects the RAID to the host. This may be a single point of failure. It also ignores AC line power failures and any occurrences of data loss due to software or operator problems.

Table II shows the reliability of various components of a SCSI based RAID. The MTTF figures for the disk drive and the Host Bus Adapter are estimated reliability numbers from manufacturers. The MTTF figures for power supplies are from MIL-HDBK-217D [Bardos 86]. The MTTF for fans is estimated from MIL-HDBK-217E part code 801 (electric motor, < 1 horsepower) with 4 solder connections used in a ground benign (machine room) environment at 40 degrees Celsius. The MTTF figure for SCSI cables is estimated from MIL-HDBK-217E part code 1105 (printed wiring board connector) with 50 active pins and 50 milliamps per pin used in a ground benign environment at 40 degrees Celsius with 0.04 mate/unmate cycles per 1000 hours.

The MTTF figure for power cables is estimated from MIL-HDBK-217E part code 1103 (power connector) with 4 active pins and 2.5 amps per pin used in a ground benign environment at 40 degrees Celsius with 0.04 mate/unmate cycles per 1000 hours.

Reliability of RAID Components		
Component	MTTF	Source of MTTF Number
Small Disk Drive	40,000 hrs	[CDC 88]
100W Switching Power Supply	174,000 hrs	[Bardos 86]
300W Switching Power Supply	123,000 hrs	
500W Switching Power Supply	85,000 hrs	
DC Brushless Fan	195,000 hrs	[MIL 86] - Part Code 801
SCSI Cable (7 disk drives)	21,000,000 hrs	[MIL 86] - Part Code 1105
Power Cable (7 disk drives)	10,000,000 hrs	[MIL 86] - Part Code 1103
SCSI Host Bus Adapter	120,000 hrs	[Moren 88]

Table II: Reliability of RAID Components

A formula for determining the data integrity of RAID was introduced in the RAID paper. I will paraphrase it here as

$$MTBDL_{RAID} = \frac{(MTTF_{disk})^2}{n_G * (G+1) * G * MTTR} \qquad [1]$$

$MTBDL_{RAID}$ = Mean Time Between Data Loss
$MTTF_{disk}$ = Mean Time To Failure (of an individual disk)
n_G = number of parity groups
G = number of data disks per parity group
$MTTR$ = Mean Time To Repair after failure

This formula is valid for any system with redundancy groups where each group can tolerate one fault at a time (two or more concurrent faults per group means failure). This type of fault tolerance is sometimes called "N+1 redundancy." MTTR (Mean Time To Repair) is the time to replace a failed disk and reconstruct data onto the new disk. Because the above formula does not account for support hardware failures, any failure rate contribution of support hardware will limit estimated MTBDL to a lower value. Some calculations on a sample RAID will illustrate this.

Consider a RAID consisting of $n_G = 7$ parity groups. Each parity group has $G = 7$ data disks and one parity disk. There are 56 disks total. Using $MTTF_{disk} = 40{,}000$ hours, and setting $MTTR = 12$ hours, Formula 1 from the RAID paper yields MTBDL = 340,000 hours. This is a period of about 39 years and can be considered an upper bound on the estimated data integrity of this RAID (given our assumptions). This 56 disk RAID would require as support hardware eight SCSI HBAs with cables, eight 300 Watt power supplies (each with a power cable), and eight fans for cooling. Remembering the assumption that MTTF = $\frac{1}{\lambda}$, we can add the failure rates of the support hardware components to the failure rate of the 56 disk RAID to get an overall failure rate. (This assumes that any support failure may cause data loss.) From this we estimate MTBDL to be 5650 hours for this RAID. This is a time period of about 235 days and represents a factor of 60 decrease in MTBDL from the simple estimate that considers only disk drive failures. By judicious placement of parity, SCSI, power, and cooling groups we can do much better.

2.5.1. Using the inherent redundancy in RAID

If parity groups are mapped onto the disk array orthogonal to SCSI, power, and cooling groups then no single hardware failure will cause data loss. This is illustrated for our sample 56 disk RAID in Figure 1.

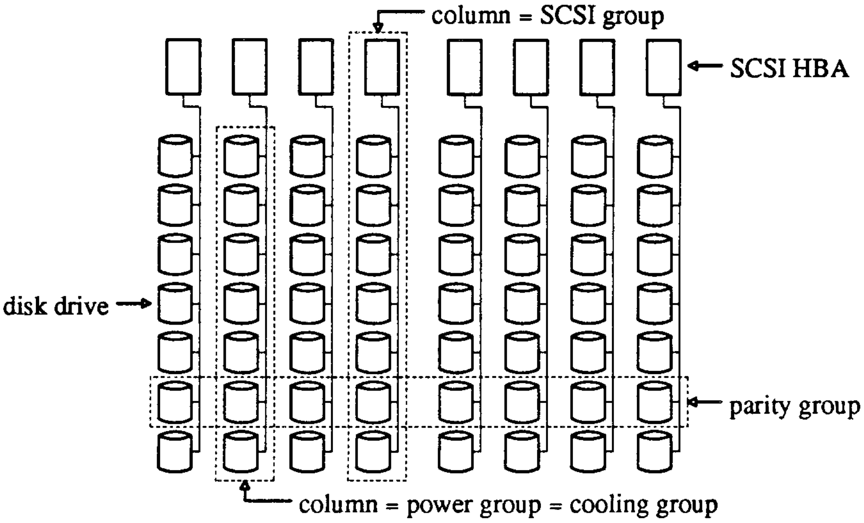

Figure 1: Groupings in Disk Array

This scheme does not have any explicit fault tolerance of the support hardware, but uses the redundancy of parity groups to protect against support hardware failures as well as disk failures. A good way to implement this is to make SCSI groups the same as power groups the same as cooling groups, and then map parity groups onto the array orthogonal to the other groups. This is convenient since SCSI, power, and cooling groups all require physical proximity of the drives but parity groups have no such requirement.

In Figure 1, there are 56 disks divided into 8 columns. Each column is a SCSI group, power group, and cooling group. The parity groups are mapped onto the disk array so that any column failure affects only one disk per group. The data integrity of this configuration can be estimated by the formula

$$MTBDL_{RAID} = \frac{(MTTF_{disk})^2}{\left[\frac{(MTTF_{disk})^2}{(MTTF_{column})^2} + \frac{2n_G * MTTF_{disk}}{MTTF_{column}} + n_G \right] * (G+1) * G * MTTR} \qquad [2]$$

$MTBDL_{RAID}$ = Mean Time Between Data Loss
$MTTF_{disk}$ = Mean Time To Failure (disk)
$MTTF_{column}$ = Mean Time To Failure (column)
n_G = number of parity groups
G = number of data disks per parity group
$G+1$ = number of columns (support hardware)
$MTTR$ = Mean Time To Repair after failure

This formula is derived in the Appendix. It considers the interaction of two types of failures: disk drive failures and column (support hardware) failures, and it assumes that the number of columns is equal to the number of disks (data and check) per parity group. It reduces to Formula 1 from the RAID paper as $MTTF_{column}$ tends to infinity.

As an example of the use of this formula, again consider the 56 disk RAID discussed earlier and shown in Figure 1. For this RAID, $n_G = 7$, $G = 7$, $MTTF_{disk} = 40{,}000$, and $MTTR = 12$ hours. Column reliability is obtained from the MTTF figures given in Table II. I will assume each column has one SCSI HBA, one 300 Watt power supply, one fan, one SCSI cable, and one power cable. Remembering that for reliability estimation, MTTF = $\frac{1}{\lambda}$, so the failure rates for the items in a column can be summed to get the column failure rate and column MTTF. Performing this computation for our example gives $MTTF_{column} = 46{,}000$ hours. Plugging values into Formula 2 yields $MTBDL_{RAID}$ = 119,500 hours or about 13.6 years. Comparing this value against the 5650 hours computed earlier shows that clever organization has recovered a factor of 20 in MTBDL without any extra cost or hardware. MTBDL is still a factor of three smaller than the value estimated by Formula 1 from the RAID paper, however.

Now consider a different organization where there is one smaller power supply per disk drive instead of one power supply per column. Power supply failures are now counted as disk failures rather than column failures. I will use 174,000 hours as the MTTF of the smaller power supply. $MTTF_{disk}$ becomes 32,500 hours and $MTTF_{column}$ becomes 74,000 hours, while the other values all remain unchanged. The formula yields $MTBDL_{RAID}$ = 117,800 hours or about 13.4 years. This is virtually the same as the last result and says that as long as parity groups and power groups are orthogonal, there is no data integrity benefit to a power supply per disk configuration.

2.5.2. Using explicit redundancy

To approach the data integrity estimate given by Formula 1 from the RAID paper we must consider explicit redundancy in the disk array. The impact of support hardware failures on MTBDL can be decreased by implementing redundancy in each column. This is illustrated in Figure 2.

Employing redundancy in each column has the effect of increasing $MTTF_{column}$ in Formula 2, thus increasing $MTBDL_{RAID}$. Power and cooling groups can be made redundant, for instance; and SCSI HBAs can be duplicated on each SCSI cable. Further investigation of our sample RAID will illustrate this.

For this RAID, $n_G = 7$, $G = 7$, $MTTF_{disk} = 40{,}000$, and $MTTR = 12$ hours. Originally, each column had one SCSI HBA, one 300 Watt power supply, one fan, one SCSI cable, and one power cable; and $MTTF_{column}$ was calculated to be 46,000 hours. Formula 2 estimated $MTBDL_{RAID}$ to be

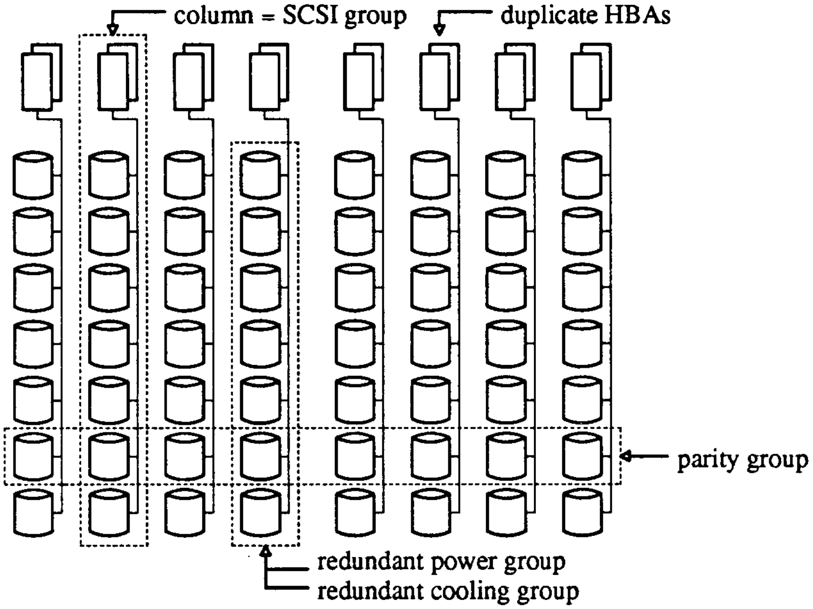

Figure 2: Redundant Groupings in Disk Array

119,500 hours for this configuration. As redundancy is applied to power, SCSI, and cooling groups, $MTTF_{column}$ increases. Using the higher $MTTF_{column}$ values in Formula 2 then gives an increase in $MTBDL_{RAID}$. This progression is shown in Table III for our sample 56 disk RAID.

Table III shows the increase in MTBDL as fault tolerance is applied to power groups, SCSI groups, and cooling groups in that order. At each stage, $MTTF_{column}$ is limited by non-redundant hardware and this in turn limits $MTBDL_{RAID}$ (estimated by Formula 2) to a smaller value than the 340,000 hours estimated by Formula 1. Each row of Table III compares the estimated $MTBDL_{RAID}$ from Formula 2 against the theoretical maximum estimated by Formula 1. With redundant power supplies, single failures of HBAs and fans limit $MTTF_{column}$ to 73,500 hours and $MTBDL_{RAID}$ to 159,600 hours. With redundant power supplies and duplicate HBAs, single cooling failures limit $MTTF_{column}$ to 189,400 hours and $MTBDL_{RAID}$ to 238,000 hours. (I ignore for the moment the SCSI limit of 8 devices per cable to allow the consideration of duplicate HBAs.) The bottom row of Table III shows that with redundancy applied to all support hardware in a column except cabling, $MTTF_{column}$ is limited by the cabling to 6,600,000 hours. Even so, this allows $MTBDL_{RAID}$ to be 336,000 hours. This is very nearly equal to the maximum estimate given by

Formula 1 and shows how it is possible to eliminate disk array support hardware failures as an important factor in the data integrity of RAID.

RAID Data Integrity with Redundant Support Hardware			
	$MTTF_{column}$	$MTBDL_{RAID}$	Percent of Maximum
RAID as is	46,000 hrs	119,500 hrs	35%
add redundant P.S.	73,500 hrs	159,600 hrs	47%
add duplicate HBAs	189,400 hrs	238,000 hrs	70%
add redundant fans	6,600,000 hrs	336,000 hrs	99%

Table III: RAID Data Integrity with Redundant Support Hardware

It is not possible to make individual SCSI and power cables fault tolerant in our RAID. Throughout this analysis, I have assumed that parity groups are mapped onto the disk array such that any single cabling failure affects only one disk per parity group. This may not be the case in general, and cabling failures can become a major factor limiting $MTBDL_{RAID}$ to lower values than those displayed in Table III, even if all other support hardware is fault tolerant.

3. Some Practical Considerations in RAID

This section of the report covers some of the practical considerations that must be addressed in assembling an array of disks and discusses how a SCSI based RAID prototype could be constructed at UC Berkeley in the near future.

3.1. Physical packaging

The packaging for a RAID should provide features to satisfy requirements imposed by the architecture of a large array of disks. Besides power and cooling, there are many other requirements. The footprint (MB/sq. ft.) should be as good or better than what is available in SLED technology. This can be an important selling point for RAID since it is positioned as a replacement for SLED. Cable management must also be addressed, as the number of cables will grow in proportion to the size of the array. Easy access should be provided to allow replacement of failed hardware (disks, power supplies, fans, cabling). This is particularly important if online maintenance is a goal. If easy access to all components is not possible, access to the cabling should be sacrificed first. Cabling should be the most reliable component in a RAID.

As system complexity increases, insuring that service personnel correctly identify failures and then replace only what has failed becomes a bigger challenge. The host or RAID I/O processor must provide unambiguous status and error messages to operators and service personnel. The packaging should provide an easy method for identifying and replacing failed components once a failure has been detected by the system. Many 5.25 inch disk drives have a Ready or Fault light; this can be useful in locating a particular unit and verifying the failure before replacement.

Our concept of RAID includes a requirement for scalability. We would like to provide a mass storage system that can scale an order of magnitude or more. To accommodate this, a packaging scheme based on a unit of modularity is desirable. Providing for small minimum expansion size is also desirable for applications where storage needs grow gradually. Finally, the packaging should provide for containment of radiated and conducted emissions from the electronics. This is not a serious issue for a prototype, but must be considered for any commercial design.

3.2. Power

Small disk drives in the 5.25 inch and 3.5 inch form factors do not have built in power supplies, but require +12 Volt DC and +5 Volt DC from an external supply. Disk drives have dynamic power requirements, particularly the +12 Volt used in the spindle motor and actuator. Spinup power requirements (lasting 15 - 30 seconds as the disks begin to spin) may be twice as much as typical operating power. Luckily, spinup of SCSI disk drives can be software controlled by the START/STOP UNIT command. This feature is important in RAID because it allows the host to stagger spinups to reduce peak power consumption during system powerup. This allows the RAID to operate with smaller, more reliable power supplies.

Various power supply configurations are possible for a RAID. One power supply per disk drive is feasible. In this configuration, the power group is the individual disk and the RAID redundancy guarding against disk failures also guards against power supply failures. Each supply must be sized to accommodate the full disk drive spinup power. This can be as much as 60 Watts for a 5.25 inch drive.

One power supply per SCSI cable is another possibility. In this configuration, the power supply group coincides with the SCSI group. The impact on data integrity from a power supply failure would be similar to the impact from an HBA failure. A power supply for seven 5.25 inch disk drives would need to be about 420 Watts if sized so all drives could spinup simultaneously. If drive spinups were sequenced in software and only two drives were allowed to spinup simultaneously, a 300 Watt supply would be sufficient for seven drives.

A larger power supply for several SCSI groups is also a possibility. The limit to this is packaging. Recall that to accommodate RAID scalability, a packaging scheme based on an independent unit of modularity is desirable. Data loss from power supply failure becomes a possibility if more than one disk per parity group is supported by the failed power supply. Fault tolerance can be applied to this problem by using N+1 redundancy. Commercial power supplies offering N+1 redundancy and online replacement of power modules are available in ratings from 300 Watts on up.

For applications requiring data integrity and very high availability, some form of emergency backup power or uninterruptible power supply would be appropriate. With MTTF of North American urban power typically equal to 2 months, power outages are a major source of failures among computer users who do not have emergency backup power [Gray 85]. Several types of uninterruptible power supplies are commercially available and can be easily added to any design [McGowan 87]. Power supplies will represent only 2 - 5 percent of the cost of a RAID, so cost probably won't be an issue in deciding what type of power configuration is best for a particular design.

3.3. Cooling

Electronic equipment generates heat, and disk drives are no exception. While many small disk drives are specified to perform satisfactorily with convection air cooling, the dense packaging desirable for RAID will most likely dictate some sort of forced air cooling. Although testing is the best way to determine cooling requirements, the following formula approximates the amount of airflow required for a given application.

$$Q = \frac{3.16W}{\Delta T} \qquad [3]$$

Q = air flow required (Cubic Feet per Minute or CFM)
W = power dissipated by the equipment (Watts)
ΔT = temperature rise of the air above incoming ambient (deg. F)

This formula is commonly used by cooling equipment manufacturers as a first approximation for determining cooling requirements. It is based on the heat capacity and standard density of air at sea level. For altitudes above sea level, more air flow is required to maintain the same ΔT temperature rise due to the decrease in the density of air. A good "rule of thumb" is to use 15 degrees Fahrenheit for ΔT. This will result in effective cooling without oversizing the fan.

Using cooling to reduce component operating temperatures is an application of fault avoidance to increase system reliability. Naturally the cooling equipment also has a failure rate and this must be taken into account. Fortunately, air moving devices such as fans provide a good opportunity to apply N+1 redundancy. Tandem uses this technique to insure that the failure of any single fan does not cause overheating of any part of the system.

3.4. A RAID prototype

In order to better understand RAID, we are constructing a prototype in the Computer Science Division at the University of California at Berkeley. I describe here how an array of 56 5.25 inch disks can be assembled along with a minicomputer as the RAID I/O processor to form a RAID prototype. As a research project, the first RAID prototype may make sacrifices in the areas of capacity, reliability, and convenience features for the sake of economy and simplicity. The goals of the first RAID prototype are (in rough order of importance):

1) To provide an array of embedded SCSI interface disks that can be logically configured into any one of the various RAID levels. This array of disks will serve as a platform for RAID performance studies and RAID I/O processor architecture development.

2) To demonstrate the basic concepts behind RAID, such as parity data redundancy and data reconstruction. (So Professor Patterson can pull a drive out of an operating RAID without causing data loss.)

3) To gather information for an accurate comparison between RAID mass storage systems and SLED mass storage systems in the areas of performance, cost, capacity, footprint, power, and data integrity.

4) To gather data on the failure behavior of the components of RAID: disk drives, power supplies, cooling equipment, SCSI HBAs, and cabling.

5) To provide familiarity with Uninterruptible Power Supplies and their potential for use with RAID. (So Professor Katz can pull the plug out of the wall without interrupting operation.)

6) To provide insights into scalability limitations of the RAID architecture and of the particular implementation (e.g. SCSI) we have chosen.

7) The process of assembling RAID hardware should be a valuable learning experience.

For this prototype, we have chosen the SCSI bus as the interface to the disk array. SCSI is an industry standard, high level, device independent interface [ADSI 85]. It offers several benefits that make it desirable for RAID. SCSI is a daisy chained bus shared by up to 8 devices. It allows disk drives to disconnect from the bus during the relatively long time periods associated with seeks and rotational delays. Many SCSI drives now offer track buffers, which accept read data from the platters at the drive's internal data rate and then transfer the data across the SCSI bus at a higher rate. Together, these two features minimize bus usage per I/O request and allow multiple drives to share the same cable and HBA without serious performance degradation. The SCSI bus has several limitations which may affect the scalability of a SCSI based RAID but should not affect the operation of a prototype. There is the limit of 8 devices per SCSI bus. The Host Bus Adapter (HBA) counts as one of those devices, leaving a maximum of 7 embedded SCSI disk drives per cable. The single-ended SCSI cable is limited to 6 meters in length. SCSI offers a differential option that increases maximum cable length to 25 meters, but this is not commonly supported by small disk drives. SCSI burst data transfer bandwidth is currently limited to a maximum of 5 MB/sec. Two possibilities for future disk interfaces in RAID are SCSI-2 and IPI-3.

For this prototype proposal we have chosen the CDC Wren IV as our storage unit. The Wren IV is a full height 5.25 inch embedded SCSI disk drive with a formatted capacity of 344 MBytes. It implements SCSI disconnect/reconnect and has a 32 KByte data buffer. The Wren IV appears to be a mature product that offers high capacity and good performance without pushing the state of the art in magnetic areal density. MTTF is specified as 40,000 hours [CDC 88].

A logical block diagram of the Berkeley RAID-I prototype is shown in Figure 3. A Sun 4 minicomputer is the RAID I/O processor. A number of SCSI/VME Host Bus Adapters (HBAs) reside in the backplane of the Sun 4 and act as the interface to the disk array. In Figure 3, eight HBAs are shown, each with a SCSI cable to connect to a number of disks. With a limit of eight devices per SCSI cable, this RAID can accommodate up to 56 disk drives.

The small disk drives are not standalone units, but require some sort of physical packaging. Certainly there are many forms that this packaging could take. One idea is a hinged panel offering access to two surfaces, each populated with disk drives. This seems more appropriate for the smaller and lower weight 3.5 inch disks than for 5.25 inch disks. A second idea is the "wall of disks" where each disk would be mounted in an individual slot. This scheme would work equally well for 3.5 inch disks and 5.25 inch disks. A third idea is to base the disk drive packaging on the standard rack. This alternative is explored in detail below.

3.4.1. Disk array rack

The standard rack is a good choice for packaging RAID for a number of reasons. RAID is positioned as a replacement for large disks (SLEDs) and the standard package for large disks is the rack or similar sized box. Each rack can operate independently of its peers and provides a good unit of modularity for expanding RAID. The rack provides an enclosure to contain radiated emissions from the electronics. For a prototype, the rack offers the benefits of standard parts and hardware.

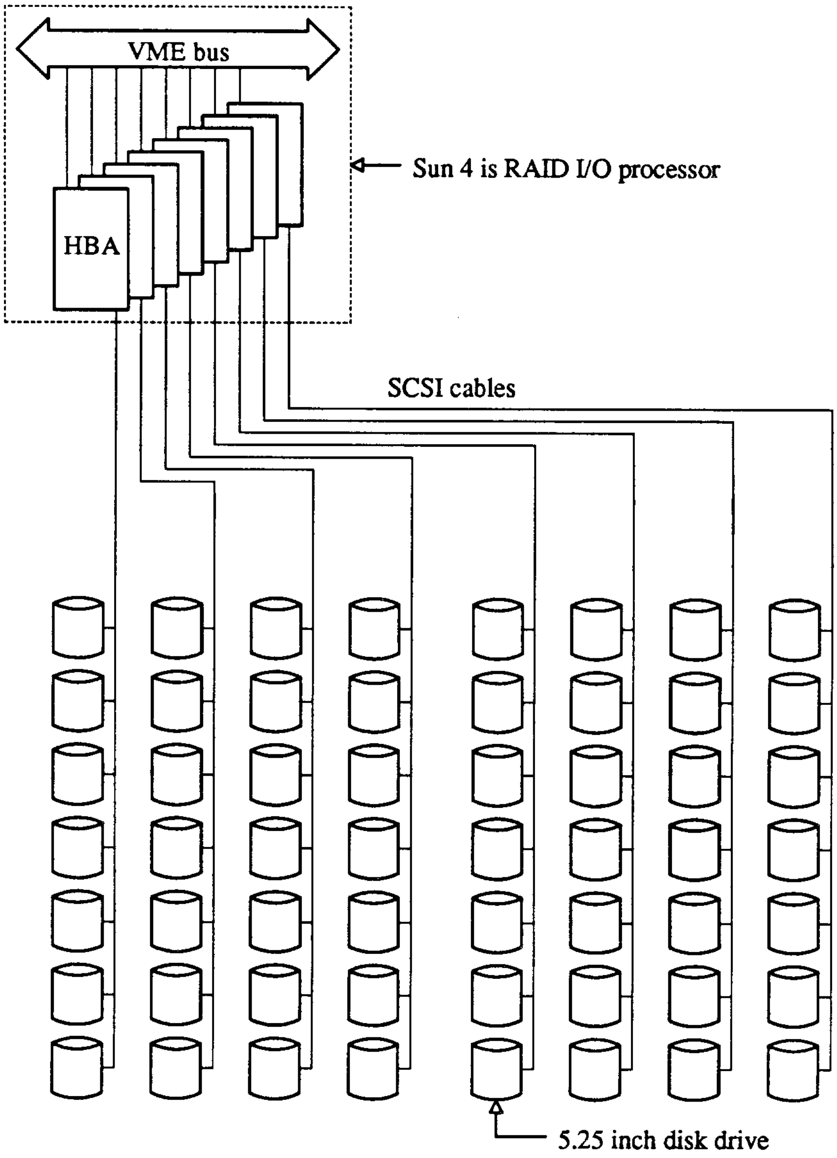

Figure 3: Berkeley RAID-I Prototype

The disk array of Figure 3 could be packaged in two racks as Figure 4 shows. Each disk array rack contains 28 full height 5.25 inch disk drives (16 in the front and 12 in the rear) and 4 SCSI cables along with power supplies and fans. Each rack conforms to the standard footprint of 22 inches wide by 30 or 36 inches deep and can be substituted directly for a rack of large format disks in the machine room.

Using 34 watts as the typical power of a Wren IV and assuming a power supply efficiency of 75 percent, the typical power consumption per rack can be calculated to be 1300 Watts. The

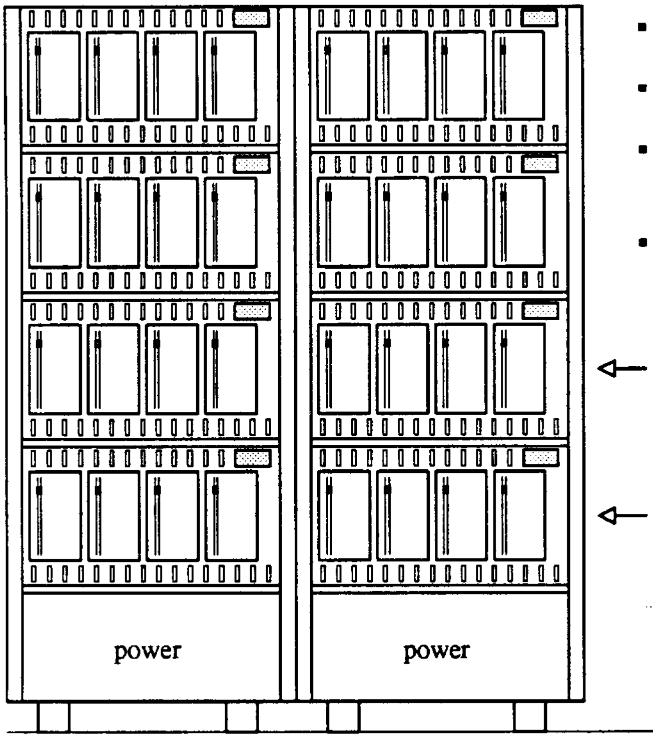

- Two standard racks
- Each rack 48" high x 22" wide x 30" deep
- 28 full height 5.25 inch disk drives and 4 SCSI cables per rack
- About 1.3 Kilowatts per rack

4 pullout drive packages per rack

Each drive package contains one SCSI cable and 7 embedded SCSI disk drives

Figure 4: RAID-I Prototype Disk Array

28 drives in a rack are divided into four SCSI groups; each group occupying a separate disk drive package. Before further discussing the rack, I will discuss the disk drive package in detail.

3.4.2. Disk drive package

The disk array rack described above contains four disk drive packages, each containing seven 5.25 inch disk drives sharing a common SCSI cable. Four disk drives are accessible from the front and three from the rear. Each disk package is mounted on chassis slides allowing it to be pulled forward to gain access to the interior. For a RAID constructed of half height 5.25 inch drives, 14 drives and two SCSI cables would fit into the same package; and for a RAID constructed of 3.5 inch drives, 21 drives and three SCSI cables would fit into the same package. A front view of the disk drive package is shown in Figure 5 and a rear view is shown in Figure 6.

Vents along the top and bottom allow for cooling. Connections for power and SCSI are at the bottom rear of the package. In this design, identification is provided to aid service personnel

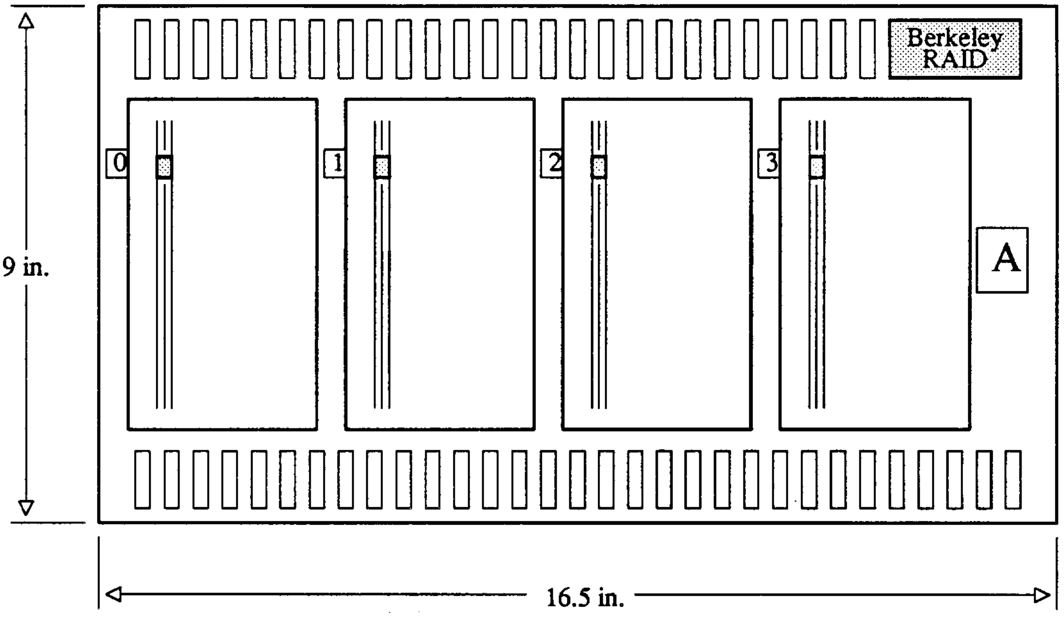

Scale: □ = 1 inch x 1 inch

- Four 5.25 inch disk drives in front
- Vents top and bottom for cooling
- Drive package ID letter uniquely identifies package and SCSI cable
- Drive ID numbers correspond to SCSI bus addresses

Figure 5: Disk Drive Package - Front View

in locating a particular drive within the RAID for service or maintenance. To do this, two pieces of information are necessary: which SCSI group and which drive in that group. Disk drive package is equivalent to SCSI group, so each package is identified by an ID letter corresponding to the SCSI cable and HBA. Each disk drive is identified by an ID number which corresponds to its SCSI address and is located near the front panel LED of the Wren IV. While the LED is not controllable from outside the Wren IV, observed behavior of the LED can be helpful in verifying operation (or non-operation) of a particular drive in question. The drive ID number also tells what address should be set on a drive being installed in a particular slot. SCSI bus addresses are set by placing jumpers on the disk drive before installation. This is an important step since two devices with duplicate addresses on the same SCSI cable will cause problems.

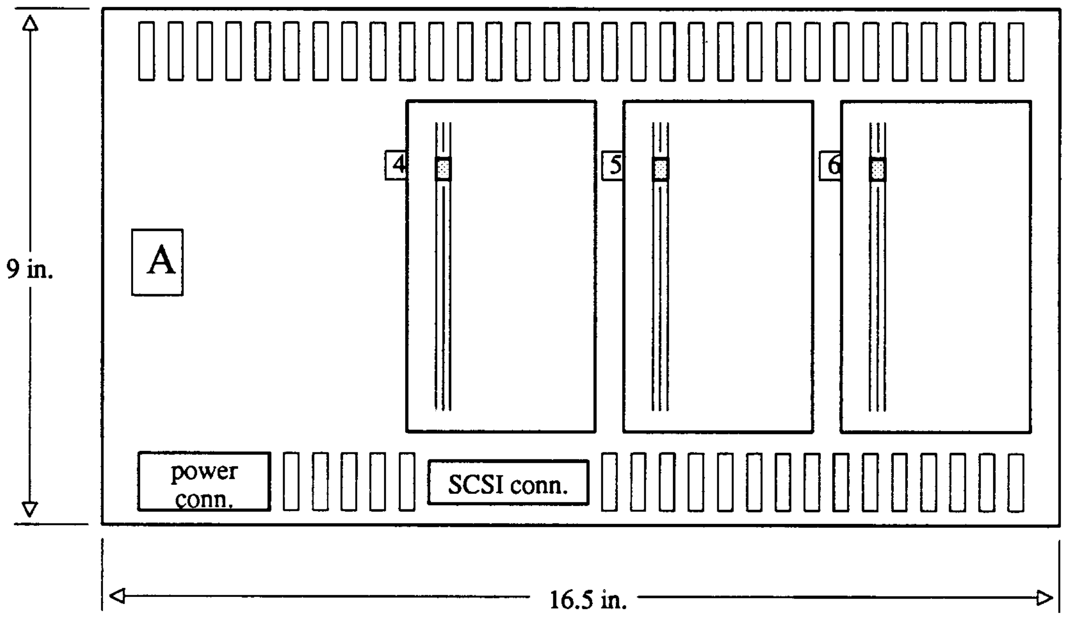

Scale: □ = 1 inch x 1 inch

- Three 5.25 inch disk drives in rear
- Vents top and bottom for cooling
- Drive package ID letter uniquely identifies package and SCSI cable
- Drive ID numbers correspond to SCSI bus addresses

Figure 6: Disk Drive Package - Rear View

A top view of the disk drive package is shown in Figure 7. Pulling the package forward out of the rack would give this view and allow access to the inside of the package. This figure shows the seven drives daisy-chained together by the SCSI cable. Resistive termination is required at the end of the cable and this is shown as a separate block. The Wren IV offers internal SCSI termination, but for RAID it is wise not to use this option. If the internal termination were utilized in the last drive on the cable, removal and installation of that drive would differ from the rest. On-line maintenance of the last drive would interrupt SCSI bus activity.

There is space to include power supplies for the drives, either one per drive or a larger supply for all seven. Alternatively, one large supply could be located in the bottom of the rack and

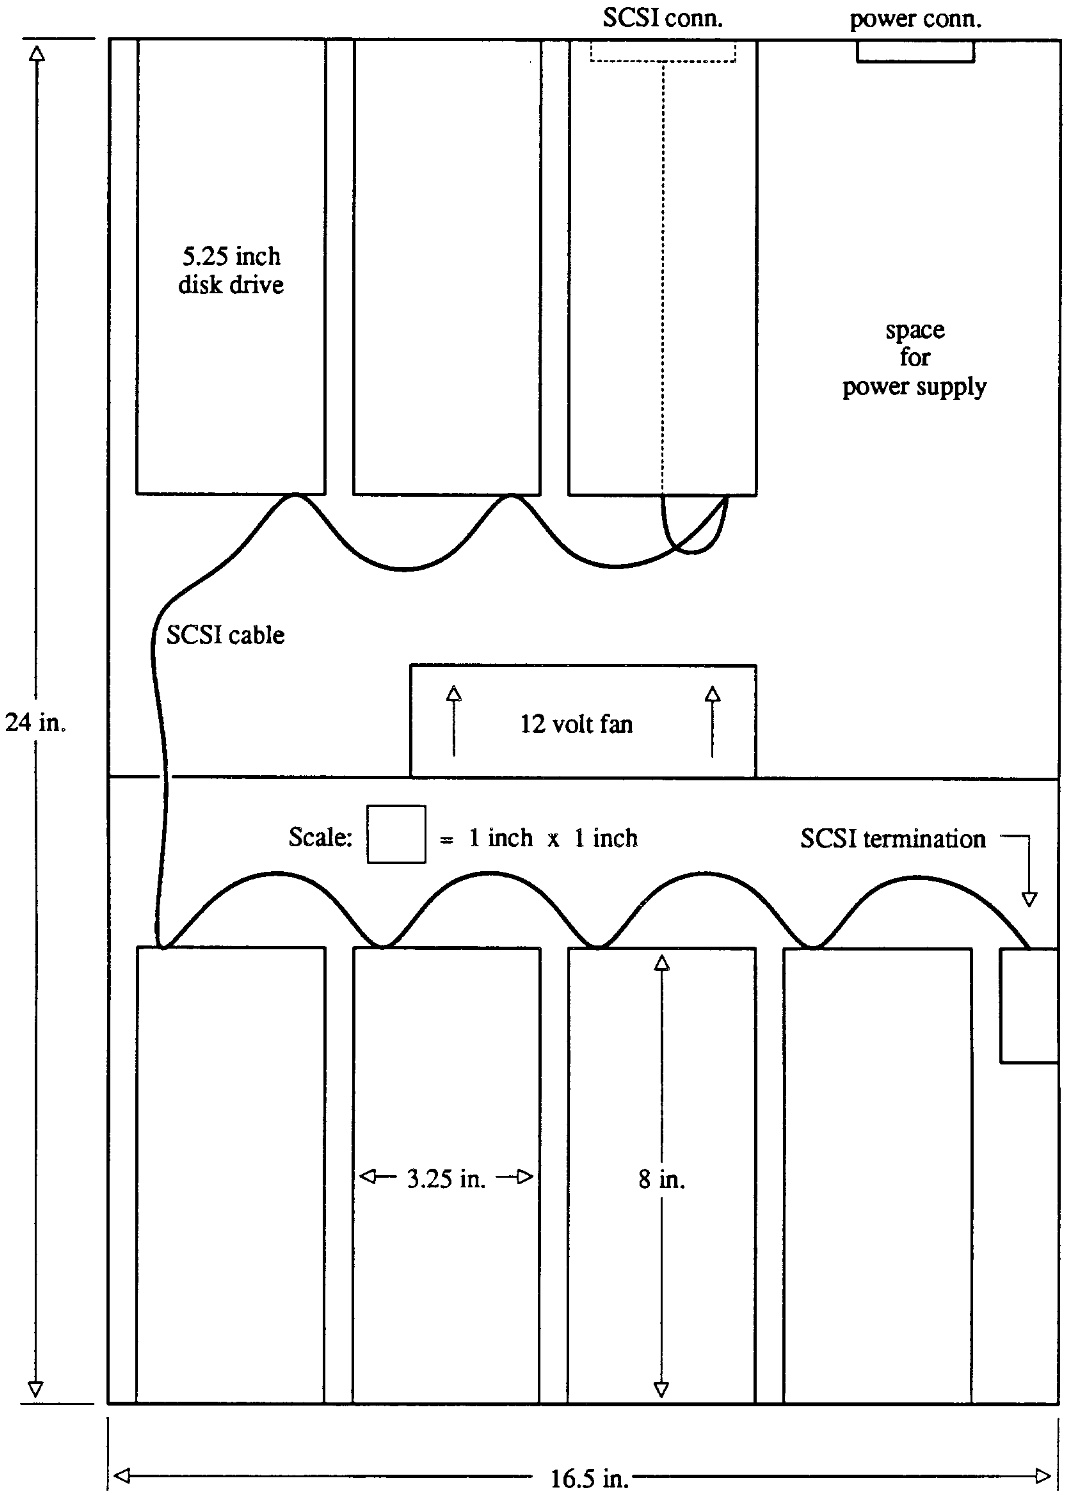

Figure 7: Disk Drive Package - Top View

power distributed to each disk package in the form of +12 Volts DC and +5 Volts DC. A fan will most likely be necessary to cool this package and could be sized by the formula given earlier. The time-averaged worst case power consumption (excluding spinup) of the Wren IV is 36.6 Watts. If only drives were installed in this package, the fan would need to cool 260 Watts. If a 75 percent efficient power supply were also included within this space, the fan would need to cool 350 Watts.

This package offers access to both the front and the rear of each disk drive. One drawback to this design is the extra SCSI cable length required to allow the package to be pulled forward out of the rack. About 1.2 meters of SCSI cable are needed inside the package and about 2.5 meters of cable are needed to reach the bulkhead at the bottom of the rack from the rear of the drive package (including the slack to allow the package to be pulled forward). This leaves 2.3 meters of cable to reach the HBA in the RAID I/O processor. While this could limit the scalability of this design for a single-ended SCSI based RAID, it is not a problem for a first prototype.

The FRUs (Field Replaceable Units) for this package are: disk drive, power supply, fan, and cabling. Many small disk drives have several FRUs: the HDA (Head Disk Assembly containing the stored data) and one or two electronics modules. For RAID it makes sense to treat the entire drive as an FRU. After a failure, it is easier and faster to reconstruct data onto a new drive than to attempt to repair a failed drive and reinstall it. A static control wrist strap should be included inside the package to encourage proper ESD (Electro-Static Discharge) procedures by service personnel. This is inexpensive insurance protection for the valuable disk drives.

3.4.3. An alternative rack organization

If online maintenance is a goal, then an alternative rack organization might be preferable for this RAID. This is shown in Figure 8. This rack is the same physical size and houses the same number of disk drives as the design presented earlier, but differs in internal construction.

For online maintenance, manual human contact with cables and connectors should be minimized. This requires extra packaging engineering effort but allows the cabling to be hidden

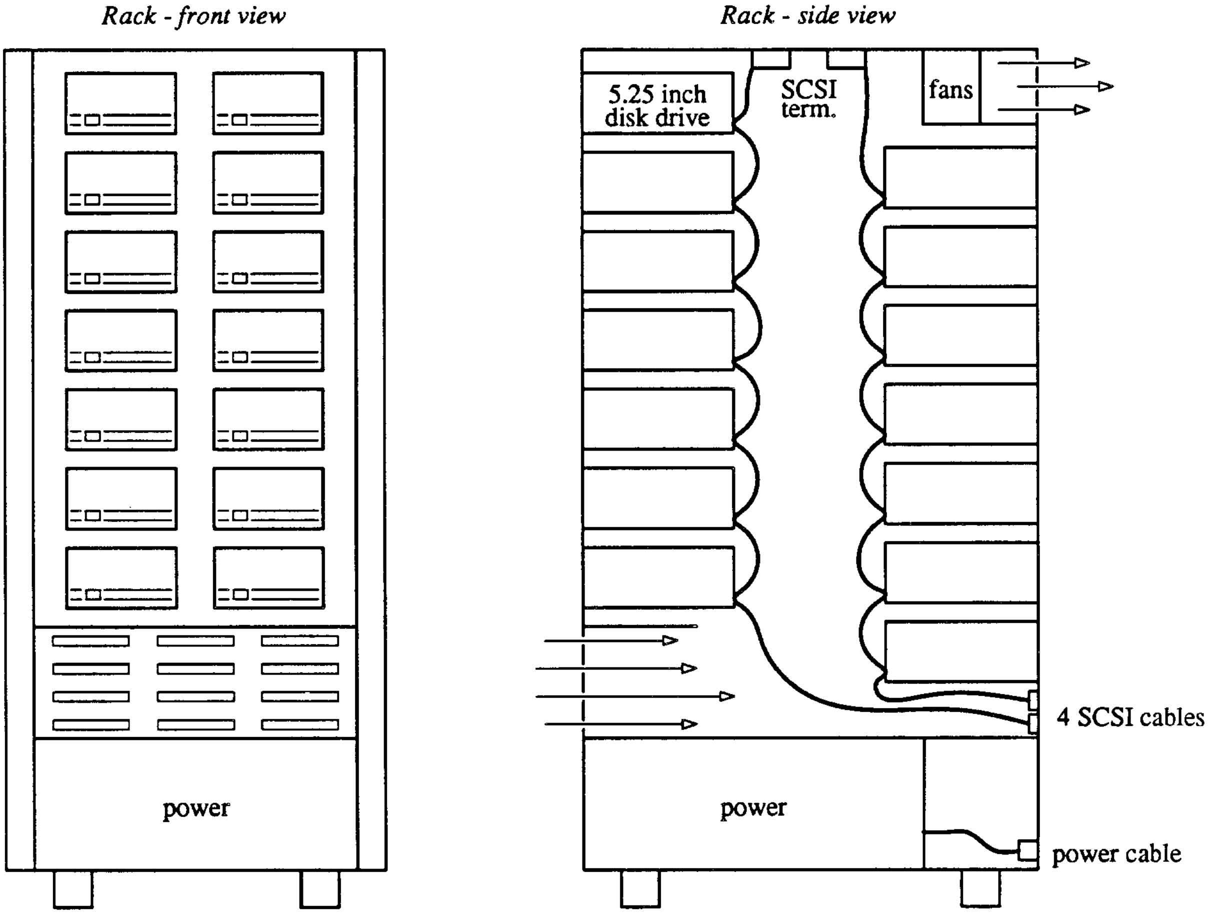

Figure 8: Alternative Rack Organization

inside the rack. There are two other advantages to this alternative design. Since a number of fans work together to cool the entire rack, N+1 redundancy is easily implemented. This design needs less SCSI cable within the rack (2 meters vs. 3.7 meters for the other design), leaving more cable length available to reach the RAID I/O processor.

The power supply shown at the bottom of the rack could be four separate supplies (one for each SCSI group) or it could be one large supply with redundancy. Data capacity, footprint, and power consumption of this rack would be the same as the design presented earlier.

3.5. The Berkeley RAID-I

To complete the hardware of the Berkeley RAID-I, each SCSI cable is connected to a

VME/SCSI Host Bus Adapter (HBA) residing in the backplane of the RAID I/O processor. This was shown in Figure 3 with a Sun 4 as the I/O processor. Software running in the Sun 4 implements the algorithms necessary for the functionality of the RAID I/O processor. For a first prototype, implementing the algorithms in software is a wise choice, as software is easy to change and manipulate. Scaling this RAID prototype to larger numbers of disks can be accomplished by adding more disk drives and HBAs. This can be done by adding complete disk array racks or by incrementally adding drives to a partially filled rack.

Five examples of magnetic disk storage racks are displayed in Table IV. Two examples of RAID level 5 are shown; one based on the 344 MByte Wren IV and the other based on the 574 MByte Wren V. These RAIDs are compared with racks of Fujitsu M2361A "Super Eagle" disks and NEC D2362 disks as examples of minicomputer storage; and with the IBM 3380 model AK4 disk unit as an example of mainframe storage.

Storage Rack Comparison					
Characteristics	IBM 3380-AK4	4 Fujitsu M2361A	8 NEC D2362	RAID level 5 (Wren IV)	RAID level 5 (Wren V)
Year introduced	1987	1986	1987	1988	1988
Formatted Capacity (MB)	7,560	2,400	5600	8,260	13,780
Number of actuators	4	4	8	28	28
MB/actuator	1,910	600	700	344	574
Magnetic areal density (Mb/sq. in.)	35.0	16.4	22.9	24.6	24.6
Platter diameter (inches)	14	10.5	9	5.25	5.25
Number of transfer paths	4	4	8	4	4
Burst transfer rate (MB/sec)	12	10	19.7	16	16
I/O's/sec/actuator	41	38	43	30	33
I/O's/sec/rack	164	152	344	840	924
Footprint (MB/sq. ft.)	805	440	1020	1,500	2,500
Power/rack (watts)	1,870	1,930	2,000	1,300	1,300

Comparison of SCSI based RAID level 5 with racks of Fujitsu M2361A "Super Eagle" disks and NEC D2362 disks as examples of minicomputer storage; and with IBM 3380 model AK4 disk unit as an example of mainframe storage. The IBM unit has its own cabinet while the rest each occupy one standard rack 22" wide x 36" deep. The RAIDs assume 24 data disks and 4 parity disks. The rotated parity of RAID level 5 allows all 28 actuators to be counted for I/O comparisons. "I/O's/sec/actuator" means the number of average accesses possible per second per actuator. Average access time = average seek time + average rotational delay.

Table IV: Storage Rack Comparison

The IBM unit has its own cabinet while the rest each occupy one standard rack 22 inches wide by 36 inches deep. This RAID level 5 assumes 24 data and 4 parity disks, but due to its rotated parity, all 28 actuators can be counted for I/O comparisons.

As Table IV shows, the RAID storage racks offer improvements over the other disks in terms of capacity, footprint, power, and I/O's per second. The number of I/O's per second for RAID is much greater than the other racks because the RAID has many more actuators. (For this table, "I/O's/sec/actuator" means the number of average accesses possible per second per actuator. Average access time = average seek time + average rotational delay.) This table clearly shows RAID's potential for SLED-beating performance.

4. Conclusions

4.1. RAID data integrity

RAID (Redundant Array of Inexpensive Disks) appears to offer many benefits over mass storage systems based on large format disks. Through fault tolerance, the data integrity of RAID can be much better than that of a standard mass storage system. When analyzing the effects of hardware failures on the data integrity of RAID, it is important to consider the failure rate of disk array support hardware. A Small Computer System Interface (SCSI) based system such as we are considering requires power supplies, SCSI Host Bus Adapters (HBAs), cooling equipment, and cabling as support hardware for the disks. Using the inherent redundancy in RAID is a good first step in minimizing the effects of support hardware failures on data integrity and availability. This is accomplished by arranging the RAID parity groups so that any *single* support hardware failure affects only one disk per parity group. To eliminate disk support hardware failures as a major factor in the data integrity of RAID requires the use of explicit redundancy in the disk array.

Despite redundancy in RAID, there may still be single points of failure. The RAID I/O processor and AC line power are two examples. These were not addressed in this report, only the disk array itself was examined. This report did not consider the impact of disk drive unrecoverable data errors (one sector read incorrectly) on RAID data integrity. These types of errors are specified separately from MTTF and could be a major factor in the data integrity of RAID. Also, this report did not consider software or operator errors as factors in the data integrity of RAID, although these will most likely play a major role.

4.2. Practical considerations in RAID

SCSI is a high level, standard interface implemented by many products. These qualities make it a good choice for a RAID prototype. The SCSI features of disconnect/reconnect, data buffering in the disk drive, and good data transfer rate will boost RAID performance. SCSI does have drawbacks that may limit the performance and scalability of this type of RAID. SCSI bus overhead may become a performance limitation if seven disk drives are placed on each cable (as shown in this report). The single ended implementation of SCSI limits cable length to 6 meters. Many SCSI HBAs offer differential (which stretches to 25 meters) as an option, but this is not widely available on disk drives. For future implementations, we look toward a longer, higher bandwidth bus that supports more devices. These features may be available in SCSI-2 or IPI-3.

The packaging for a RAID prototype can take many forms, but the standard rack is a good choice for several reasons. RAID is positioned as a replacement for large format disks and the standard package for large disks is the rack or similar sized box. The rack is also a convenient sized package for providing power and cooling; and is a good unit of modularity for scaling RAID to larger systems. A RAID level 5 of 5.25 inch disk drives housed in a standard rack can potentially offer better I/O performance, footprint, and data integrity for lower power and cost than a mass storage system based on large format disks.

Acknowledgements

I must first thank Digital Equipment Corporation for sponsoring my education at Berkeley. The people in the GEEP (Graduate Engineering Education Program) office deserve thanks for all their help. Janis Ackerman, in particular, always offered kind words and encouraging comments in response to my monthly reports. Also, everyone at Digital Colorado who was involved in my GEEP experience deserves my gratitude for indulging my wish to attend school for a year.

Next I must thank Professor R.H. Katz and Professor D.A. Patterson for accepting me into their group and giving me the opportunity to perform this work. Their guidance and suggestions helped shape my research.

Those who reviewed earlier drafts of this report and offered comments deserve thanks for substantially improving the presentation. This group of people includes Professor Katz, Professor Patterson, Garth Gibson, Peter Chen, Ken Lutz, and Steven Sprouse.

Lastly I must thank the University of California at Berkeley and the city of Berkeley for providing an interesting and stimulating environment for my studies. I learned as much about myself as I did about Computer Science.

Appendix: Data Integrity Calculation

The formula for Mean Time Between Data Loss (MTBDL) is derived from probability theory. There are two types of failures: disk failures and column (support hardware) failures. Since it takes two concurrent failures to cause data loss in this RAID, there are four possible scenarios for this to happen.

1) Any disk failure followed by a second disk failure within the same parity group before the first failure is repaired. The probability of this occurring can be approximated as

$$\frac{n_G * (G+1)}{MTTF_{disk}} * \frac{G * MTTR}{MTTF_{disk}}$$

2) Any disk failure followed by the failure of any column not containing the failed disk before the first failure is repaired. The probability of this occurring can be approximated as

$$\frac{n_G * (G+1)}{MTTF_{disk}} * \frac{G * MTTR}{MTTF_{column}}$$

3) Any column failure followed by any disk failure (not within that column) before the first failure is repaired. The probability of this occurring can be approximated as

$$\frac{G+1}{MTTF_{column}} * \frac{n_G * G * MTTR}{MTTF_{disk}}$$

4) Any column failure followed by a second column failure before the first failure is repaired. The probability of this occurring can be approximated as

$$\frac{G+1}{MTTF_{column}} * \frac{G * MTTR}{MTTF_{column}}$$

$MTBDL_{RAID}$ = Mean Time Between Data Loss
$MTTF_{disk}$ = Mean Time To Failure (disk)
$MTTF_{column}$ = Mean Time To Failure (column)
n_G = number of parity groups
G = number of data disks per parity group
$G+1$ = number of columns (support hardware)
$MTTR$ = Mean Time To Repair after failure

Summing these four probabilities yields an overall failure rate λ, which approximates the probability of data loss by this RAID due to disk or support hardware failures. For our reliability model we have assumed MTBDL = $\frac{1}{\lambda}$, so we can invert the expression to find MTBDL. This formula can be checked by letting $MTTF_{column}$ tend to infinity. If this is done, the expression reduces to Formula 1 from the RAID paper.

$$MTBDL_{RAID} = \frac{(MTTF_{disk})^2}{\left[\frac{(MTTF_{disk})^2}{(MTTF_{column})^2} + \frac{2n_G * MTTF_{disk}}{MTTF_{column}} + n_G \right] * (G+1) * G * MTTR} \qquad [2]$$

References

[ADSI 85] Adaptive Data Systems Inc., *SCSI GUIDEBOOK*, Second Edition, Adaptive Data Systems, Pomona CA, June 1985.

[Anderson 88] K.L. Anderson, Digital Equipment Corporation, private communication, June 1988.

[Bardos 86] P. Bardos, "The reliability of switch mode power supplies - 15 years on," *Electronic Engineering*, vol. 58, no. 715, July 1986, pp. 37-44.

[CDC 88] Control Data Corporation, *Product Specification for Wren IV SCSI Model 94171-344*, Control Data OEM Product Sales, Minneapolis MN, January 1988.

[GGI 88] Gartner Group Inc., Conference proceedings of the *Gartner Group Fourth Annual Computer Storage Conference*, March 1988, Phoenix AZ.

[Gray 85] J. Gray, "Why Do Computers Stop and What Can Be Done About It?" Tandem Technical Report 85.7, Tandem Computers, Cupertino CA, November 1985.

[Maxion 88] R.A. Maxion and D.P. Siewiorek, "Symptom-Directed Diagnosis of Distributed Computing Systems," *1986/1987 Research Review*, Computer Science Department, Carnegie Mellon University, Pittsburgh PA, 1988.

[McGowan 87] K.J. McGowan, "Three Types of UPS Give Users A Choice For Clean, Reliable Power," *Computer Technology Review*, vol. VII, no. 13, October 1987, pp. 123-127.

[MIL 86] U.S. Department of Defense, *Military Handbook: Reliability Prediction of Electronic Equipment, MIL-HDBK-217E*, October 1986.

[Moren 88] W.D. Moren, Ciprico Inc., private communication, July 1988.

[Patterson 88] D.A. Patterson, G. Gibson, and R.H. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)," *ACM SIGMOD Conference*, June 1988, Chicago IL.

[Quantum 87] Quantum Corporation, *OEM Manual* and *Programmers Manual* for Q200 Series disk drives, Quantum Corporation, Milpitas CA, May 1987.

[Siewiorek 82] D.P. Siewiorek and R.S. Swarz, *The Theory and Practice of Reliable System Design*, Digital Press, Bedford MA, 1982.

[Vasudeva 88] A. Vasudeva, "A Case For Disk Array Storage System," *Systems Design and Networks Conference Proc., Mass Storage Trends and Systems Integration*, ed. Kenneth Majithia, unpublished, April 1988.