

IDA-based Disk Arrays

AZER BESTAVROS*
Harvard University

May 1991

Abstract

Parallel disks can improve I/O performance in a manner analogous to the use of parallel processors to improve computation times. However, due to their data storage function, reliability issues become very important for these parallel disk systems. Currently proposed schemes use shadowing or parity to achieve reliability in parallel disks. In this paper we introduce the idea of using the Information Dispersal Algorithm (IDA) of Michael O. Rabin [Rabin:87] to distribute data and redundancy information *uniformly* among multiple disks and compare the performance and reliability characteristics of shadowing, parity, and IDA. We discuss some ways to take advantage of the uniformity of data placement and argue that IDA is the algorithm of choice for achieving reliability and performance in parallel disk systems.

Index Terms: Reliability, Performance, Parallel I/O, RAID storage systems, Information Dispersal Algorithm.

Acknowledgement: I would like to thank Prof. Michael O. Rabin for his valuable comments. Also, I would like to thank Wing Wong and Danny Chen from Bell Labs for their help and suggestions.

*This work was done while the author was with AT&T Bell Laboratories, August 1989.

1 Introduction

Recently, parallel disk systems have emerged as a potential solution for achieving ultra-high capacity, performance and reliability at a reasonably low cost. This emerging technology is likely to have a great impact on the market of data-storage devices for mainframe computers, massively parallel computers and large database systems [Zachary:89]. The basic idea underlying the design of parallel disk systems is to use a large number of *inexpensive* small disks to attain very large capacities. The relatively low cost per megabyte of storage on these disks,¹ combined with their better volumetric and power efficiencies [Patterson:89], makes them very attractive when compared to larger disks. In [Bestavros:89c], we discussed a number of issues related to the design, performance and reliability of parallel disk systems (disk arrays).

A significant drawback of using a large number of parallel disks to improve I/O performance is the effect it has on reliability. More disks means more frequent failures and, thus, an increased probability of data loss. Without taking additional protective measures, it is well known that the probability of losing data increases almost *linearly* with the number of disks in the system. With large disk arrays,² the reliability of the system drops to unacceptable levels. For example, even with a small disk array of 32-64 disks, the MTTF (Mean Time To Fail) is expected to drop to a few days.

There are two approaches that can be adopted to improve the performance of parallel disk systems: fault-avoidance and fault-tolerance [Randell:79]. With fault-avoidance, the overall system's reliability is enhanced by using higher quality components, better assembly, extensive testing, regular maintenance, *etc.* This approach is not appropriate for parallel disk systems for several reasons. First, higher quality components translates to more expensive small disks. This defeats the underlying philosophy of parallel disk systems which relies on the fact that small disks are cheap! Second, with a large number of small disks, regular maintenance and testing becomes impractical. Finally, the fault-avoidance approach does not *scale* nicely. A fault-avoidance measure that works for a system of tens of disks is not likely to work for a system of hundreds or even thousands of disks! Fault-tolerance, on the other hand, seeks the same goal of improving the overall system reliability by using protective redundancy to recover from the effect of malfunctions. In a sense, malfunctions are expected to happen, but, if they do, their effects are automatically counteracted by the redundancy.

Interest in parallel disk systems has been spurred by the sharp decline in the price of small hard disk drives. The current ratio between the cost per Megabyte for large and small disks is somewhere between 2 and 5, and, with current market trends, this figure is only expected to rise. Previous research [Patterson:89] used minimal redundancy primarily to ensure acceptable levels of fault-tolerance. Performance gains from the added redundancy (if any) were viewed as a side effect. In our opinion, this should be reversed. With the aforementioned pricing trends, more redundancy should be used, primarily, to boost performance with fault-tolerance being a byproduct.

In this paper, we propose a novel technique for adding redundancy to disk arrays. Our approach is based on the Information Dispersal Algorithm (IDA) of Michael O. Rabin [Rabin:87]. We show that IDA is superior to current approaches, namely: shadowing [Bates:85] and parity [Gibson:88].

In shadowing, faults are tolerated by maintaining multiple identical copies of the data. In parity, redundant information is computed and stored with the data to enable the system to compensate for possible data losses. The IDA approach is similar to parity except that there is no distinction between parity information and data. Using the results we obtained in [Bestavros:89c], we analytically compare all three techniques and conclude that IDA is superior. It yields the best reliability and performance gains for every percent of added redundancy.

¹Roughly 3-10 Dollars per megabyte for 3.5-inch disks (*e.g.* Conners CP3100) compared to 20-45 Dollars per megabyte for the 10.5-inch and the 14-inch disks (*e.g.* DEC RA-81, IBM 3380).

²Some of these are estimated to have as many as 1000 disks! (see [Gibson:88]).

2 Shadowing

This is the simplest and most straightforward approach for increasing the reliability of a system. The idea is to keep $(r + 1)$ replicas of every block of data. Each one of these replicas is striped over a group of $m = n/(r + 1)$ disks. This scheme tolerates up to r failures out of the n disks by an r -fold increase in the required storage.

2.1 The Seek Time:

In this section, we analyze the seek time of a shadowed system and how it relates to the parameters n and r .³ We do not assume that any intelligent arm scheduling algorithm is used. Thus, we assume that the accesses are uniformly distributed over the disks diameter D . Notice that this assumption becomes invalid if a scheduling algorithm like the ones described in [Biton:89] is used. Also, we assume that the disks are asynchronous.⁴

Read Access:

To read a data block, the controller will have to wait for the first replica to become available. Since the number of disks involved in reading a replica is m , the expected seek time becomes the expected maximum of these m disk seeks.⁵ Using the results in [Bestavros:89c], we get:

$$E(\text{Seek for a group read}) = \frac{D}{2} \frac{\Gamma(m+1)}{\Gamma(m+1.5)} \sum_{i=1}^m \frac{\Gamma(i+0.5)}{\Gamma(i+1)}$$

Each one of these groups can be approximated with a *logical disk* having the same expected seek time. For a single disk, however, the expected seek time is one third its maximum seek distance. Thus, the logical disk would have a maximum seek distance of:

$$D' = 3 \frac{D}{2} \frac{\Gamma(m+1)}{\Gamma(m+1.5)} \sum_{i=1}^m \frac{\Gamma(i+0.5)}{\Gamma(i+1)}$$

The seek time for reading in a shadowed system is, thus, the minimum of the $r + 1$ logical disks. Using the analysis in [Bestavros:89c], we get:

$$E(\text{Shadowed read seek}) = \left(\frac{3}{2r+3} \right) \frac{D}{2} \frac{\Gamma(m+1)}{\Gamma(m+1.5)} \sum_{i=1}^m \frac{\Gamma(i+0.5)}{\Gamma(i+1)}$$

Table 7 in the Appendix shows the expected seek time for reading a shadowed organization with n disks and tolerating r failures with $D = 1$. Notice that for a constant level of fault-tolerance, the seek time tends to increase as the total number of disks per server increases. This is due to the fact that the number of disks per group increases, resulting in an increase in the expected seek time of the group (logical disk). Notice also that for a given number of disks per server, the expected seek time tends to decrease sharply as the level of fault-tolerance increases. This is due to two factors. First, the number of groups increases

³We assume that the seek time is dominant compared to other delays (*e.g.* latency and overhead).

⁴This assumption is pessimistic. In a number of situations, one might argue that other assumptions are more appropriate. For example, seek synchronization [Bestavros:89c] is a more appropriate assumption if the head movement of the different disks are correlated.

⁵For seek synchronous systems, the expected seek time should be *Dover3* (see [Bestavros:89c]).

making the expected minimum seek time decrease. Second, the number of disks per group decreases making the expected seek time per group decrease as well.

Write Access:

To write a data block, the controller will have to wait for all of the disks (all groups and all disks per group) to respond. Using the analysis from [Bestavros:89c], we get:

$$E(\text{Shadowed write seek}) = \frac{D}{2} \frac{\Gamma(n+1)}{\Gamma(n+1.5)} \sum_{i=1}^n \frac{\Gamma(i+0.5)}{\Gamma(i+1)}$$

The seek time for writing in a shadowed system is, thus, independent of the level of fault-tolerance. It increases as the total number of disks in the system increases.

When deriving the above estimate, we assumed that *all* the pieces (namely n) have to be written *before* committing the write request. That was necessary to maintain the same level of fault-tolerance throughout the whole operation. Another approach, however, would be to relax this restriction a little bit to improve performance. We illustrate this using an the following example:

Assume that a shadowing system tolerates up to two faults by maintaining three replicas of each stripe. Using the usual (conservative) approach, we will have to write all three stripes for every write request. As we have shown above, this will cause the seek and latency delays to increase. In a sense, this performance degradation is the price we have to pay for the increased protection against failures. The question, of course, is whether this price is justified! For instance, if we commit the write request after writing only one copy (the first one to finish), then the seek time would be changed from 0.543 to 0.143, a decrease of about seventy-five percent. One might argue that, by doing so, the level of fault-tolerance is jeopardized, since a failure in the disk with the only copy (before the other two copies get to be written) might result in an unrecoverable error. Although true, this is unrealistic. First, the probability of having a failure in the interval of time between the completion of the first write and the second is extremely small. Second, even if that failure occur, the second and third writes can still go through unless the whole system fails (say as a result of a power failure). As a matter of fact, we can still tolerate such unlikely situation by adopting a more conservative solution where we commit the write request after two of the copies get written (the first two to finish). This will make the seek time 0.314, a decrease of over forty percent from the original write-all approach. Assuming the occurrence of two failures before the third write is completed is unquestionably overly pessimistic.

Read/Write Access:

For a read/write mix of α , the expected seek time is given by:

$$E(\text{Seek time for an } \alpha \text{ read/write mix}) = \alpha E(\text{Read seek}) + (1 - \alpha) E(\text{Write seek})$$

In typical applications, read requests tend to be much more frequent than write requests. However, the use of buffer caches makes the actual (physical) read/write mix more balanced. Tables 8, 9, and 10 in the Appendix show the expected seek time in a shadowed system for a read/write mix of 0.25, 0.50 and 0.75, respectively.

2.2 The Transfer Time

Since the total number of disks per server is n , it follows that data can be striped into $m = n/(r+1)$ pieces. These pieces can be accessed in parallel, thus reducing the transfer time by a factor of $m = n/(r+1)$.

3 Parity

A major drawback of the shadowing approach is the excessive amount of redundancy needed to protect against failures. Instead of replicating data, redundancy can be added to the system so as to *correct* the erroneous information. In [Patterson:89], [Patterson:88] such an approach for tolerating a single failure in a group of N disks was suggested and termed *$N+1$ RAID*. The idea is to calculate and store parity information of a group of disks on a bit-per-disk basis. One parity block would be needed for every N blocks across the disks. Any single disk failure can be corrected simply by reading the rest of the disks in the group to determine what bit values of the failed disk would result in getting the proper parity.⁶

In this section, we generalize the parity approach to be able to tolerate more than one failure. We analyze the expected performance of such an organization and contrast it to the shadowing approach.

3.1 Level of redundancy and fault-tolerance

Let the number of data disks per server be m and assume that it is required to tolerate up to r simultaneous faults using the parity approach. The parity disks will partition the system into a number of parity groups. Each parity group consists of a number of data disks m and one parity disk.⁷ Obviously, each one of the parity groups can tolerate at most one fault. To be able to tolerate up to r faults, we should have every data disk in at least r parity groups. If no more than r faults occur, then the data on any disk is recoverable since that disk is guaranteed to be in at least one parity group with at most one fault.

One way of organizing the data and parity disks to guarantee the above condition is to imagine that the data disks are organized in an r -dimensional space.⁸ Obviously, every data disk will have r coordinates. Each one of these coordinates identifies a parity group that the data disk is a member of. For such an arrangement, one parity disk is needed per coordinate for each dimension. Let the number of parity groups (and thus the number of parity disks) needed to tolerate r faults be p . We have:

$$\begin{aligned} p &= r \lceil \sqrt[r]{m} \rceil^{r-1} \\ &\approx rm^{\frac{r-1}{r}} \end{aligned}$$

To tolerate one fault, we get $p = 1$ which corresponds to the *$N+1$ RAID* of [Patterson:88]. To tolerate more faults, the number of parity disks grows rapidly. For example, to tolerate 2 simultaneous faults amongst 64 data disks, we need 16 parity disks (25 percent added redundancy). In this case, the parity groups consist of 8 data disks and one parity disk. To tolerate 3 simultaneous faults, however, we need 48 parity disks (75 percent added redundancy). In this case, the parity groups consist of 4 data disks and one parity disk.

⁶The N information disks along with the parity disk form a parity group.

⁷Notice that an information disk can be in more than one parity group.

⁸The cases where $r = 2, 3$ were examined in [Gibson:88].

3.2 Seek Time

In this section, we analyze the seek time of a parity system. However, before doing that, we will discuss how reading and writing in such a system is accomplished. We assume that data is striped over the m data disks.

To write a data block, all $m + p$ disks are accessed in parallel to write all the data and parity information to the disks. In case of no failures, the access would take as long as the slowest of the $m + p$ disks. In case of failures, data should be written on different sectors or on standby units. Depending on the seriousness of the failure, this process might be more complicated. To read a data block, all m disks are accessed in parallel. In case of no failures, the access would take as long as the slowest of the m disks. If failure(s) are detected, the system will have to:

- [1] Lock all parity disks associated with the failed disk(s),
- [2] Read all the necessary parity groups,
- [3] Compute the data block that would give the correct parities,
- [4] Write back the corrected data on different sectors or on standby units,
- [5] Unlock the parity disks.

Since failures are not expected to be frequent, this performance overhead can be neglected on the average. In the above discussion, we have assumed that fine grain striping is used, making it necessary for every request to access all data disks. If coarse grain striping is used, it might be possible to fulfill the request by accessing *fewer* disks. For read requests, this might be an advantage, since the seek/latency delays are expected to be smaller. Write requests, however, become very expensive [Patterson:89]. In particular, to write data to a single disk in a parity system, the system will have to:

- [1] Lock all the disks in all the parity groups associated with the disk to be written,
- [2] Read the old data,
- [3] Read the old parity,
- [4] Perform the write request,
- [5] Compute and write back the new parity using the formula:
New parity = (old data \oplus new data) \oplus old parity, and
- [6] Unlock the parity disks.

For a system with even a small percentage of writes, the parity disks will quickly become a performance bottleneck. Thus, in the following analysis, we will assume that fine grain striping is used.

Read Access:

Using the formula for the expected seek time for the maximum of m seeks [Bestavros:89c], we get:

$$E(\text{Read seek in a parity system}) = \frac{D}{2} \frac{\Gamma(m+1)}{\Gamma(m+1.5)} \sum_{i=1}^m \frac{\Gamma(i+0.5)}{\Gamma(i+1)}$$

Write Access:

Using the formula for the expected seek time for the maximum of $n = m + p$ seeks [Bestavros:89c], we get:

$$E(\text{Write seek in a parity system}) = \frac{D}{2} \frac{\Gamma(n+1)}{\Gamma(n+1.5)} \sum_{i=1}^n \frac{\Gamma(i+0.5)}{\Gamma(i+1)}$$

3.3 The Transfer Time

The transfer to/from the m data disks can be done in parallel, resulting in decreasing the transfer time by a factor of $m = n - p$.

4 The Information Dispersal Approach

Recently, Michael O. Rabin [Rabin:87] devised a new algorithm for the secure and fault-tolerant communication and storage of information. The basic idea of this algorithm is to disperse the contents of a data block into n different pieces so that recombining *any* m of these pieces, $m \leq n$, is sufficient to reconstruct the original data block. Figure 1 illustrates the dispersal of m stripes into n stripes using IDA. Obviously, such an approach would tolerate up to $n - m$ faults. The algorithm uses *irreducible polynomial arithmetic* to disperse or reconstruct the data. If the size of the data to be dispersed is S then the size of each of the pieces is $(1/m)S$, making the total required storage $(n/m)S$. The major aspect of IDA is that the added redundant information is not identifiable (as is the case with the parity approach). Rather, it is distributed among the data blocks. This makes the different pieces of information in the system *uniform*. That is, there is no distinction between data and parity.

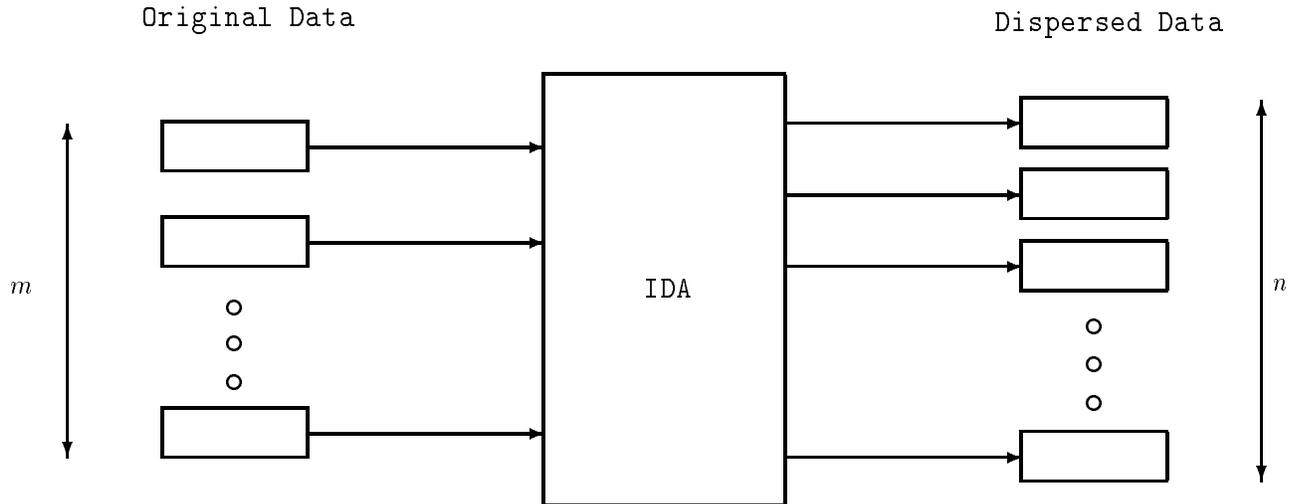


Figure 1: Data dispersal using IDA

4.1 Level of Redundancy and Fault-tolerance

To tolerate up to r simultaneous faults, IDA requires that the total number of dispersed stripes exceeds the minimum number of stripes needed for reconstruction by r . Thus, a total of $n = m + r$ stripes is needed for every m stripes of data, a redundancy of $100(n - m)/m$ percent. Obviously, IDA requires the least amount of redundancy compared to shadowing and parity.

4.2 The Seek Time

In this section, we analyze the seek time of IDA and how it relates to the parameters n and r . The estimates we get should be thought of as upper bounds. A number of optimizations (using intelligent layout of data stripes on the different cylinders and appropriate arm scheduling algorithms) are possible which would result in decreasing the seek times considerably.

Read Access:

To read a data block, all n disks are accessed in parallel and the first m of these to reply are used to reconstruct the data block. If no more than r failure(s) are detected, the system will have to recompute the missing stripes and rewrite them to different sectors or standby disks. Notice that this will not cause any delays for the read request. Using the formula for the expected seek time [Bestavros:89c] we get:

$$E(\text{Read seek in an IDA system}) = \frac{D}{2} \frac{\Gamma(n+1)}{\Gamma(n+1.5)} \sum_{i=n-m+1}^n \frac{\Gamma(i+0.5)}{\Gamma(i+1)}$$

Write Access:

To write a data block, all the n disks are accessed in parallel to write all the stripes. In case of no failures, the access would take as long as the slowest of the n disks. In case of a failure, data should be written on different sectors or on standby units. This, however, should not delay the commitment of the write request. Using the formula for the expected seek time for the maximum of the n disks [Bestavros:89c], we get:

$$E(\text{Write seek in an IDA system}) = \frac{D}{2} \frac{\Gamma(n+1)}{\Gamma(n+1.5)} \sum_{i=1}^n \Gamma(i+0.5)\Gamma(i+1)$$

Read/Write Access:

We tabulate the seek time for the IDA approach for different read/write mixes in the appendix (Tables 11, 12, 13, and 14.) Notice that these figures should be thought of as *upper bounds* since, as it will be explained later, much better write access performance can be achieved through intelligent arm scheduling, optimized layout of data blocks and adoption of a delayed write policy.

An important property of IDA can be inferred by looking at these tables. The larger the number pieces dispersed, n , the worse the seek time of both read and write accesses for a constant level of redundancy. Table 1 shows this trend for different values of n . The write seek time depends only on n , whereas the read seek time depends on the available redundancy (we show the read seek time for redundancy levels of 25, 50, and 75 percent). Despite this deterioration in the seek time, we might still favor a larger n to cut on the transfer time, especially with fairly large block sizes. Obviously, an optimum level of dispersion can be determined once the system parameters (block size, transfer rate, maximum seek time,... *etc.*) are given.

Expected seek time				
n	Read access			Write access
	$\frac{n-m}{n} = 0.25$	$\frac{n-m}{n} = 0.50$	$\frac{n-m}{n} = 0.75$	
4	0.3905	0.2381	0.1111	0.5937
8	0.4384	0.2629	0.1216	0.7005
16	0.4671	0.2771	0.1275	0.7835
32	0.4830	0.2848	0.1307	0.8451
64	0.4914	0.2888	0.1323	0.8899
128	0.4956	0.2903	0.1331	0.9219
256	0.4978	0.2919	0.1336	0.9447

Table 1: Trend of read and write seek times

4.3 The Transfer Time

The transfer to (or from) the n (or m) disks can be done in parallel, resulting in decreasing the transfer time by a factor of m .

4.4 The Overhead Time

To assemble/disassemble blocks of data using IDA, some computations have to be done. If fine grain striping is used, this overhead can be kept very small. In [Bestavros:89], we presented an architecture for a VLSI chip to implement IDA in real-time. The chip has been fabricated by MOSIS and tested in the VLSI lab, Harvard University. The performance of the chip was measured to be about 1 megabyte per second. This corresponds, to an overhead of 1 micro-second for a stripe of 8 bits. By using proper pipelining and more elaborate designs, we believe that the overhead delay can be reduced significantly. Even, if hardware support is not available to assemble/disassemble the data, the overhead time will still be minimal compared to seek and latency delays. As a matter of fact, the overhead time of IDA becomes an issue only when it exceeds the transfer time of that block.

4.5 Arm Scheduling

In [Bitton:88], arm scheduling algorithms were proposed for shadowed systems in order to decrease the seek delays for read requests. These algorithms are impossible to use with parity systems since there is no *choice* with respect to which disks to be accessed. With the IDA approach the system can select out of the n disks the best m disks to use for accessing the requested block. Notice that using such an algorithm for arm scheduling does not change our previous estimate of the seek time since the system is, in a way, *predicting* which disks will have the smallest seek time.

4.6 Delayed Write Optimization

The main advantage of the IDA approach is the symmetry of the different stripes. There is no distinction between data and redundancy. As we have shown earlier, this property can be exploited to enhance the performance of read requests. However, when writing the dispersed information to the different disks, we have assumed that the system has to wait until *all* the pieces are written before committing the write request. This restriction results in a degradation of the performance of the write requests compared to the

read requests. In systems where disk caching is used, the percentage of write requests is high. This means that the overall system performance is likely to deteriorate.

The requirement of writing all the pieces before committing the write request stems from the need to sustain the level of fault-tolerance. This might be necessary if that level is low. However, in a system where redundancy is introduced to boost performance as well as fault-tolerance, this restriction becomes questionable. To illustrate this point, consider a system where data blocks are dispersed into 16 pieces, of which any 8 are sufficient to reconstruct the original blocks. In such a system, the read seek time will be as low as $0.27D$, whereas the write seek time will be as high as $0.78D$. Such a system, tolerates up to 8 simultaneous faults. If we are willing to sacrifice this ultra-high level of fault-tolerance a little bit for a very short period of time, then a much better write seek time can be achieved. For instance, if we commit any write request after physically writing only 9 pieces, then the write seek time becomes $0.31D$, an improvement of more than half (actually over 60 percent). One might argue that, by doing so, the level of fault-tolerance is jeopardized since the system can only tolerate one failure right after the write request is committed. This concern is true but unrealistic. First, the probability of having two failures in the interval of time between the completion of the 9th and 10th writes is extremely small. Second, unless the failures are coordinated, the rest of the writes should proceed normally, thus restoring the level of fault-tolerance to its original level.

The flexibility of not writing all the pieces before committing a write request can also be exploited to increase the achievable concurrency. Again, we illustrate this by an example. Consider a system consisting of 64 disks where data blocks are dispersed into 16 pieces of which only 8 are sufficient to reconstruct the original data. If we insist on having all the physical writes go through before committing a write request then, obviously, the system can respond to *at most* 4 write requests at any time. On the other hand, if we adopt the policy of writing (say) only 9 of the pieces before committing a write request, then the system can service up to 7 write requests at any time. This would result in boosting the throughput of system by up to 75 percent and almost doubling the achievable concurrency. Moreover, the delayed writes can be done cheaply – for instance by waiting until the head of a disk is at a favorable position (using a scheduling algorithm similar to C-scan).

4.7 Controllable Redundancy and Dynamic Reconfiguration

One of the major advantages of using IDA is that the amount of redundancy (and thus the level of fault-tolerance and the performance gains) can be adjusted for the different needs of the system. For instance, critical files would be assigned a higher level of redundancy thus insuring ultra-high accessibility and better performance. On the other hand, less critical data can be assigned lower levels of redundancy (or even no redundancy at all). This flexibility is impossible to insure with parity systems. Another main advantage of using IDA is that, with a careful design, it is possible to add new disks to the system (and thus increase its capacity) dynamically. This is also possible with shadowed systems. For parity systems, however, the distinction between parity information and data makes such a reconfiguration more difficult.

5 Comparison

We will compare the three approaches we have presented (shadowing, parity and IDA) using typical examples. The examples and figures we will be presenting are only *suggestive*. As we have discussed in [Bestavros:89c], the actual performance depends heavily on other factors as well.⁹

⁹Organization, Operating system optimizations, Tuning of parameters, *etc.*

First, we will examine all three approaches for a fixed level of fault-tolerance. We will compare the required redundancy and the achievable performance. Next, we will fix the total capacity and the percentage of redundancy and compare the achievable performance and fault-tolerance.

5.1 Fixed Fault-tolerance

Consider an example system consisting of $n = 24$ disks. We will examine two situations. First, we will assume that it is required to tolerate one failure out of the 24 disks ($r = 1$). Next, we will assume that it is required to tolerate up to two simultaneous failures ($r = 2$). The conclusions we will draw out of this example can be easily extended for larger values of n or r .

Example-1: ($n = 24, r = 1$)

Table 2 shows the level of redundancy, the transfer speedup and the expected seek time for all three striping approaches. The expected seek time is calculated for different read/write mixes of 0, 20, 50, 80 and 100 percent reads.

Example-1 ($n = 24, r = 1$)			
	Striping approach		
	Shadowing	Parity	IDA
Redundancy (percent)	100.0	4.35	4.35
Transfer speedup	12.0	23.0	23.0
Seek Time ($\alpha = 0.0$)	0.82	0.82	0.82
Seek Time ($\alpha = 0.2$)	0.75	0.80	0.80
Seek Time ($\alpha = 0.5$)	0.64	0.78	0.78
Seek Time ($\alpha = 0.8$)	0.52	0.75	0.75
Seek Time ($\alpha = 1.0$)	0.45	0.73	0.73

Table 2: Comparison for a fixed level of fault-tolerance ($n = 24, r = 1$)

From this table, we notice that the amount of redundancy for a shadowed system is much higher than that of a parity system or a system using IDA. This means a much smaller capacity out of the available space. This also means that less parallelism is available for striping. This is reflected in the lower transfer speedup of shadowed systems when compared to either parity systems or IDA systems. For write accesses, the seek time is the same for all three approaches. The seek time for read accesses, however, is different; shadowing is the best followed by IDA and then parity.

Example-2: ($n = 24, r = 2$)

Table 3 shows the level of redundancy, the transfer speedup and the expected seek time for all three striping approaches. The expected seek time is calculated for different read/write mixes of 0, 20, 50, 80 and 100 percent reads.

5.2 Fixed Redundancy / Maximum Fault-tolerance

In the above examples, we have fixed the level of fault-tolerance. Our conclusion was that IDA requires a much lesser amount of redundancy compared to the other approaches. What if IDA is allowed to use as much redundancy as does shadowing or parity? What would be the achievable performance and

Example-2 ($n = 24, r = 2$)			
	Striping approach		
	Shadowing	Parity	IDA
Redundancy (percent)	200.0	50.0	9.1
Transfer speedup	8.0	16.0	22.0
Seek Time ($\alpha = 0.0$)	0.82	0.82	0.82
Seek Time ($\alpha = 0.2$)	0.71	0.81	0.79
Seek Time ($\alpha = 0.5$)	0.56	0.80	0.74
Seek Time ($\alpha = 0.8$)	0.40	0.79	0.70
Seek Time ($\alpha = 1.0$)	0.29	0.78	0.67

Table 3: Comparison for a fixed level of fault-tolerance ($n = 24, r = 2$)

fault-tolerance gains ?

Example-3: ($n = 8, m = 4$)

Consider an example system consisting of 8 disks where up to 4 can be used for redundant storage. When comparing the cost per megabyte for small disks (5 to 10 Dollars/MB) with that for large disks (20 to 45 Dollars/MB), this level of redundancy is acceptable. Assume that the data is to be striped on the other 4 disks so as to achieve a transfer speedup of 4. With shadowing, 4 disks will be used to hold the redundant copy of the striped blocks. Thus every stripe will be stored on two different disks, thus, tolerating one erasure. To read a data block, we will have to wait for 1 out of every 2 copies to become available. With parity, the four data disks will be arranged in two dimensions, with two parity disks for row-parity and two parity disks for column-parity, thus, tolerating up to 2 erasures. To read a data block, we will have to wait for *all* of the 4 data disks to respond. To write a data block, we will have to write *all* of the data and parity disks. With IDA, data blocks will be dispersed into 8 pieces such that *any* 4 will suffice to reconstruct the original data. To read a data block, we will have to wait for the first 4 pieces (*any* 4 pieces) to be available. To write a data block, we will have to write all of the 8 pieces. Table 4 shows a comparison of the relative performance and fault-tolerance of the three approaches.

Example-3 ($n = 8, m = 4$)			
	Striping approach		
	Shadowing	Parity	IDA
Tolerable Faults	1	2	4
Seek Time ($\alpha = 0.0$)	0.70	0.70	0.70
Seek Time ($\alpha = 0.2$)	0.63	0.68	0.61
Seek Time ($\alpha = 0.5$)	0.53	0.65	0.48
Seek Time ($\alpha = 0.8$)	0.43	0.61	0.35
Seek Time ($\alpha = 1.0$)	0.36	0.59	0.26

Table 4: Comparison for a fixed level of redundancy ($n = 8, m = 4$)

From this table, we notice that IDA is superior to both parity and shadowing in terms of level of fault-tolerance. Its performance is better (although close) to shadowing and a lot better than parity (especially when the read percentage is high).

5.3 Fixed Redundancy / Fixed Fault-tolerance / Maximum Performance

In the example above, we have fixed the level of redundancy and sought the highest possible level of fault-tolerance with any improvement in the performance seen as a *side effect*. As we have discussed earlier, one might add redundancy to the system to enhance its performance in the first place! What would be the achievable performance for a given level of redundancy and a given level of fault-tolerance ?

Example-4: ($n = 8, m = 4$)

Consider an example system consisting of 8 disks where up to 4 disks can be used for redundant storage. Also, assume that the system is required to tolerate at least one erasure at any point in time. With a pure shadowing approach, the best we can do is to use all of the available redundancy to guarantee the required level of fault-tolerance in an organization identical to that used in Example-3 (4 disks for data and 4 disks for the shadow). With parity, one can suggest an organization similar to the one in Example-3, except that a delayed write policy is used with one of the groups of parity disks (row-parity or column-parity). In this case, to write a data block, we will have to wait until 6 writes terminate (instead of 8). With IDA, we can use a delayed write policy with an arrangement identical to that of Example-3. To write a data block, we will have to write any 5 pieces. Table 5 shows a comparison of the relative performance and fault-tolerance of the three approaches. We tabulate two measures for fault-tolerance: transient and steady state. The transient level of fault-tolerance reflects the reliability of the system before the delayed writes terminate, whereas the steady-state level of fault-tolerance reflects the eventual reliability once the delayed writes are finished. From Table 5, we notice that IDA is superior to both parity and shadowing in terms of the achievable performance.

Example-4 ($n = 8, m = 4$)			
	Striping approach		
	Shadowing	Parity	IDA
Tolerable Faults (transient)	1	1	1
Tolerable Faults (eventual)	1	2	4
Seek Time ($\alpha = 0.0$)	0.70	0.66	0.34
Seek Time ($\alpha = 0.2$)	0.63	0.65	0.32
Seek Time ($\alpha = 0.5$)	0.53	0.63	0.30
Seek Time ($\alpha = 0.8$)	0.43	0.60	0.28
Seek Time ($\alpha = 1.0$)	0.36	0.59	0.26

Table 5: Comparison for fixed redundancy/fault-tolerance ($n = 8, m = 4$)

In all of the above examples, we have assumed that only one of the three approaches is used. Is it possible that a combination of (say) parity and shadowing might be superior (in terms of performance and/or reliability) when compared to IDA ? The answer to this question is not obvious. For instance, we have established that when tolerating up to 1 erasure, IDA and parity are identical in terms of the required redundancy and the achievable performance. What if we use parity to tolerate failures and shadowing to improve performance ? Our answer is *No*. To illustrate this, we consider the following example:

Example-5: ($n = 10, m = 4$)

Consider a system consisting of 10 disks where up to 6 disks can be used for redundant storage. Also, assume that the system is required to tolerate at least one erasure at any point in time. Using IDA, the solution is obvious: disperse data blocks into 10 pieces of which any 4 would be sufficient and use a delayed write strategy. One can imagine another arrangement where data is striped over 4 disks and shadowed

onto another 4 disks. To tolerate 1 erasure, a parity disk will be used with every group of 4 data disks. Table 6 shows the achievable performance and fault-tolerance for these two alternatives.

Example-5 ($n = 8, m = 4$)		
	Striping approach	
	Shadowing & Parity	IDA
Tolerable Faults (transient)	1	1
Tolerable Faults (eventual)	2	6
Seek Time ($\alpha = 0.0$)	0.350	0.270
Seek Time ($\alpha = 0.2$)	0.356	0.260
Seek Time ($\alpha = 0.5$)	0.365	0.240
Seek Time ($\alpha = 0.8$)	0.374	0.220
Seek Time ($\alpha = 1.0$)	0.380	0.210

Table 6: IDA compared to a combination of parity and shadowing

From this table, we notice that IDA is still superior to the combination of parity and shadowing in terms of the achievable performance and reliability.

6 Conclusion and Future Work

In this paper we have examined the relative performance of three approaches for adding redundancy to disk arrays; namely shadowing, parity [Patterson:89] and IDA (Information Dispersal Algorithm) [Rabin:87].

Shadowing is the simplest approach to add redundancy in a system of parallel disks. The idea is to replicate the data enough to tolerate any required number of faults. It has the highest (even prohibitive) cost for a given storage capacity and fault-coverage. It performs better than the other approaches especially with high read/write mixes and intelligent arm scheduling algorithms [Bitton:89].

Using the parity approach, the data disks are partitioned into parity groups. Each group is assigned a parity disk that can be used to recover from any failure of any disk in the group. The parity approach has the worst performance when compared to shadowing and IDA.

The IDA approach is provably optimal with respect to the amount of added redundancy and the achievable level of fault-tolerance. Its main advantage is its indifference in dealing with data and redundancy. This makes it possible to make more use of the added redundancy to enhance performance. Other advantages of using IDA include: intelligent arm scheduling, dynamic reconfiguration and controllable redundancy.

Previous RAID research [Patterson:89] used redundancy primarily to ensure acceptable levels of fault-tolerance. Performance gains (if any) were viewed as a side effect. In our opinion, this should be reversed. Redundancy should be used primarily to boost performance with fault-tolerance being a byproduct. In this respect, we have demonstrated that IDA is the right choice. With an affordable level of redundancy,¹⁰ IDA can be used to reduce disk access delays and improve concurrency while insuring ultra-high levels of reliability.

In this paper we tried to show that IDA is an optimum approach for achieving internal parallelism (striping) in RAID systems. Much more work remains in order to incorporate IDA with a RAID design.

¹⁰As a matter of fact, a level comparable to what shadowing or even parity will require to achieve a modest level of fault-tolerance.

This includes: designing and possibly implementing appropriate hardware prototypes, investigating the potential support available from operating systems, and analyzing the requirements of different applications in order to fine tune the system parameters for an optimized performance. The full power of IDA will prevail.

References

- [Avizienis:76] Avizienis A., *Fault tolerant systems*, IEEE transactions on computers December 1976.
- [Bates:85] Bates K.H. and TeGrotenhuis M., *Shadowing boosts system reliability*, Computer Design, April 1985.
- [Bertsekas:87] Dimitri Bertsekas and Robert Gallager, "Data Networks," *Prentice-Hall, Inc., Englewood Cliffs, NJ*, 1987.
- [Bestavros:89a] Azer Bestavros, *A VLSI chip for the real-time information dispersal and retrieval for security and fault-tolerance*, Technical Report, TR-06-89, Dept. of Computer Science, DAS, Harvard University, January 1989. Also, in Proceedings of the 19th International Conference on Parallel Computing, August 1990.
- [Bestavros:89c] Azer Bestavros, Danny Chen and Wing Wong, *The Reliability and Performance of parallel Disks* Technical Memorandum 45312-890707-01TM, Bell Laboratories, July 1989.
- [Bitton:89] Dina Bitton, *Arm Scheduling in Shadowed Disks*, COMPCON-89, the Thirty-fourth IEEE Computer Society International Conference, March 1989.
- [Gibson:88] Garth Gibson, Lisa Hellerstein, Richard Karp, Randy Katz and David Patterson, *Coding Techniques for Handling Failures in Large Disk Arrays*, Technical Report UCB/CSD 88/477, Computer Science Division, University of California, July 1988.
- [Kurzweil:88] Fred Kurzweil, Jr., *Small Disk Arrays - The emerging approach to high performance*, COMPCON-88, the Thirty-third IEEE Computer Society International Conference, March 1988.
- [McKusick:83] Marshall McKusick, William Joy, Samuel Leffler, and Robert Fabry, *A Fast File System for UNIX*, Technical Report UCB/CSD 83/147, Computer Science Division, University of California at Berkeley, July 1983.
- [Ng:89] Spencer Ng, *Some Design Issues of Disk Arrays*, COMPCON-88, the Thirty-fourth IEEE Computer Society International Conference, March 1989.
- [Patterson:89] David A. Patterson, Peter Chen, Garth Gibson, and Randy H. Katz, *Introduction to Redundant Arrays of Inexpensive Disks (RAID)*, COMPCON-89, the Thirty-fourth IEEE Computer Society International Conference, March 1989.
- [Rabin:87] Michael O. Rabin, *Efficient Dispersal of Information for Security, Load Balancing and Fault Tolerance* Technical Report, TR-02-87, Department of Computer Science, DAS, Harvard University - April 1987. Also appeared in the Journal of the Association for Computing Machinery, Vol. 36, No. 2, April 1989, pp. 335-348. April 1989.
- [Randell:79] Randell B., *System reliability and structuring*, Ch.1 Pp. 1-18 of "Computing systems reliability", Cambridge press, 1979.
- [Rennels:84] Rennels D. A., *Fault-tolerant computing - Concepts and examples*, IEEE transactions on computers December 1984.
- [Schulze:89] Martin Schulze, Garth Gibson, Randy Katz, and David Patterson, *How Reliable is a RAID?*, COMPCON-89, the Thirty-fourth IEEE Computer Society International Conference, March 1989.
- [Zachary:89] Pascal Zachary, *Disk Arrays may reshape storage of computer data*, The Wall Street-Journal (Technology), July 19, 1989.

Appendix

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.20000							
3	0.54286	0.20000	0.14286						
4	0.59365	0.28000	0.14286	0.11111					
5	0.63059	0.28000	0.14286	0.11111	0.09091				
6	0.65901	0.32571	0.20000	0.11111	0.09091	0.07692			
7	0.68174	0.32571	0.20000	0.11111	0.09091	0.07692	0.06667		
8	0.70046	0.35619	0.20000	0.15556	0.09091	0.07692	0.06667	0.0588	
9	0.71623	0.35619	0.23265	0.15556	0.09091	0.07692	0.06667	0.05882	0.05263
10	0.72974	0.37835	0.23265	0.15556	0.12727	0.07692	0.06667	0.05882	0.05263
15	0.77671	0.40904	0.27025	0.18095	0.14805	0.10769	0.09333	0.05882	0.05263
20	0.80545	0.43784	0.28243	0.21020	0.16190	0.12527	0.09333	0.08235	0.07368
25	0.82536	0.45110	0.30020	0.21967	0.17198	0.13700	0.10857	0.09580	0.07368
30	0.84019	0.46602	0.31275	0.22725	0.17973	0.14552	0.11873	0.09580	0.08571
35	0.85178	0.47380	0.31778	0.23349	0.18593	0.14552	0.12612	0.10476	0.08571
40	0.86117	0.48327	0.32615	0.24325	0.19104	0.15208	0.12612	0.11128	0.09373
45	0.86898	0.48852	0.33287	0.24716	0.19533	0.15732	0.13180	0.11128	0.09957
50	0.87560	0.49522	0.33577	0.25061	0.19902	0.16165	0.13635	0.11630	0.09957

Table 7: Expected seek time for shadowed systems ($\alpha= 1.00$)

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.40000							
3	0.54286	0.45714	0.44286						
4	0.59365	0.51524	0.48095	0.47302					
5	0.63059	0.54294	0.50866	0.50072	0.49567				
6	0.65901	0.57568	0.54426	0.52203	0.51698	0.51349			
7	0.68174	0.59273	0.56131	0.53908	0.53403	0.53054	0.52797		
8	0.70046	0.61439	0.57535	0.56424	0.54807	0.54458	0.54201	0.54005	
9	0.71623	0.62622	0.59533	0.57606	0.55990	0.55640	0.55384	0.55188	0.55033
10	0.72974	0.64189	0.60547	0.58615	0.57918	0.56658	0.56395	0.56201	0.56046
15	0.77671	0.68479	0.65009	0.62777	0.61954	0.60945	0.60586	0.59724	0.59569
20	0.80545	0.71355	0.67470	0.65664	0.64457	0.63541	0.62742	0.62468	0.62251
25	0.82536	0.73179	0.69407	0.67394	0.66201	0.65327	0.64616	0.64297	0.63744
30	0.84019	0.74665	0.70833	0.68695	0.67507	0.66652	0.65982	0.65409	0.65157
35	0.85178	0.75729	0.71828	0.69721	0.68532	0.67522	0.67037	0.66503	0.66026
40	0.86117	0.76670	0.72742	0.70669	0.69364	0.68390	0.67741	0.67370	0.66931
45	0.86898	0.77386	0.73495	0.71352	0.70057	0.69106	0.68468	0.67955	0.67662
50	0.87560	0.78050	0.74064	0.71935	0.70645	0.69711	0.69079	0.68577	0.68159

Table 8: Expected seek time for shadowed systems ($\alpha= 0.25$)

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.33333							
3	0.54286	0.37143	0.34286						
4	0.59365	0.43683	0.36825	0.35238					
5	0.63059	0.45530	0.38672	0.37085	0.36075				
6	0.65901	0.49236	0.42950	0.38506	0.37496	0.36797			
7	0.68174	0.50373	0.44087	0.39643	0.38632	0.37933	0.37420		
8	0.70046	0.52833	0.45023	0.42801	0.39569	0.38869	0.38356	0.37964	
9	0.71623	0.53621	0.47444	0.43589	0.40357	0.39657	0.39145	0.38753	0.38443
10	0.72974	0.55405	0.48120	0.44265	0.42851	0.40333	0.39820	0.39428	0.39119
15	0.77671	0.59288	0.52348	0.47883	0.46238	0.44220	0.43502	0.41776	0.41467
20	0.80545	0.62165	0.54394	0.50783	0.48368	0.46536	0.44939	0.44390	0.43957
25	0.82536	0.63823	0.56278	0.52251	0.49867	0.48118	0.46697	0.46058	0.44952
30	0.84019	0.65310	0.57647	0.53372	0.50996	0.49285	0.47946	0.46799	0.46295
35	0.85178	0.66279	0.58478	0.54263	0.51886	0.49865	0.48895	0.47827	0.46875
40	0.86117	0.67222	0.59366	0.55221	0.52610	0.50663	0.49365	0.48623	0.47745
45	0.86898	0.67875	0.60093	0.55807	0.53216	0.51315	0.50039	0.49013	0.48427
50	0.87560	0.68541	0.60569	0.56310	0.53731	0.51862	0.50597	0.49595	0.48758

Table 9: Expected seek time for shadowed systems ($\alpha= 0.50$)

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.26667							
3	0.54286	0.28571	0.24286						
4	0.59365	0.35841	0.25556	0.23175					
5	0.63059	0.36765	0.26479	0.24098	0.22583				
6	0.65901	0.40904	0.31475	0.24809	0.23293	0.22244			
7	0.68174	0.41472	0.32044	0.25377	0.23862	0.22813	0.22044		
8	0.70046	0.44226	0.32512	0.29178	0.24330	0.23281	0.22512	0.21923	
9	0.71623	0.44620	0.35355	0.29572	0.24724	0.23675	0.22906	0.22317	0.21853
10	0.72974	0.46620	0.35692	0.29910	0.27789	0.24013	0.23243	0.22655	0.22191
15	0.77671	0.50096	0.39687	0.32989	0.30522	0.27495	0.26418	0.23829	0.23365
20	0.80545	0.52975	0.41319	0.35901	0.32279	0.29532	0.27136	0.26313	0.25663
25	0.82536	0.54466	0.43149	0.37109	0.33532	0.30909	0.28777	0.27819	0.26160
30	0.84019	0.55956	0.44461	0.38048	0.34484	0.31919	0.29909	0.28190	0.27433
35	0.85178	0.56829	0.45128	0.38806	0.35239	0.32209	0.30753	0.29152	0.27723
40	0.86117	0.57775	0.45991	0.39773	0.35857	0.32935	0.30988	0.29875	0.28559
45	0.86898	0.58364	0.46690	0.40262	0.36375	0.33524	0.31610	0.30070	0.29192
50	0.87560	0.59031	0.47073	0.40686	0.36816	0.34013	0.32116	0.30612	0.29358

Table 10: Expected seek time for shadowed systems ($\alpha= 0.75$)

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.20000							
3	0.54286	0.31429	0.14286						
4	0.59365	0.39048	0.23810	0.11111					
5	0.63059	0.44589	0.30736	0.19192	0.09091				
6	0.65901	0.48851	0.36064	0.25408	0.16084	0.07692			
7	0.68174	0.52261	0.40326	0.30381	0.21678	0.13846	0.06667		
8	0.70046	0.55069	0.43837	0.34476	0.26285	0.18914	0.12157	0.05882	
9	0.71623	0.57434	0.46793	0.37925	0.30165	0.23182	0.16780	0.10836	0.05263
10	0.72974	0.59461	0.49326	0.40881	0.33491	0.26840	0.20743	0.15082	0.09774
15	0.77671	0.66506	0.58132	0.51154	0.45049	0.39554	0.34516	0.29839	0.25454
20	0.80545	0.70818	0.63523	0.57443	0.52124	0.47336	0.42947	0.38872	0.35052
25	0.82536	0.73804	0.67255	0.61797	0.57022	0.52724	0.48784	0.45126	0.41697
30	0.84019	0.76028	0.70035	0.65041	0.60671	0.56738	0.53133	0.49785	0.46647
35	0.85178	0.77767	0.72209	0.67577	0.63524	0.59877	0.56533	0.53429	0.50518
40	0.86117	0.79176	0.73970	0.69631	0.65835	0.62419	0.59287	0.56379	0.53653
45	0.86898	0.80347	0.75433	0.71339	0.67756	0.64532	0.61576	0.58832	0.56259
50	0.87560	0.81340	0.76675	0.72787	0.69386	0.66324	0.63518	0.60912	0.58469

Table 11: Expected seek time for IDA systems ($\alpha= 1.00$)

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.40000							
3	0.54286	0.48571	0.44286						
4	0.59365	0.54286	0.50476	0.47302					
5	0.63059	0.58442	0.54978	0.52092	0.49567				
6	0.65901	0.61638	0.58442	0.55778	0.53447	0.51349			
7	0.68174	0.64196	0.61212	0.58726	0.56550	0.54592	0.52797		
8	0.70046	0.66302	0.63494	0.61154	0.59106	0.57263	0.55574	0.54005	
9	0.71623	0.68076	0.65415	0.63198	0.61258	0.59512	0.57912	0.56426	0.55033
10	0.72974	0.69596	0.67062	0.64951	0.63103	0.61440	0.59916	0.58501	0.57174
15	0.77671	0.74879	0.72786	0.71042	0.69515	0.68141	0.66882	0.65713	0.64616
20	0.80545	0.78114	0.76290	0.74770	0.73440	0.72243	0.71146	0.70127	0.69172
25	0.82536	0.80353	0.78716	0.77351	0.76157	0.75083	0.74098	0.73183	0.72326
30	0.84019	0.82021	0.80523	0.79274	0.78182	0.77198	0.76297	0.75460	0.74676
35	0.85178	0.83325	0.81936	0.80778	0.79765	0.78853	0.78017	0.77241	0.76513
40	0.86117	0.84382	0.83080	0.81996	0.81047	0.80193	0.79410	0.78683	0.78001
45	0.86898	0.85260	0.84032	0.83008	0.82112	0.81306	0.80567	0.79881	0.79238
50	0.87560	0.86005	0.84839	0.83867	0.83016	0.82251	0.81549	0.80898	0.80287

Table 12: Expected seek time for IDA systems ($\alpha= 0.25$)

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.33333							
3	0.54286	0.42857	0.34286						
4	0.59365	0.49206	0.41587	0.35238					
5	0.63059	0.53824	0.46898	0.41126	0.36075				
6	0.65901	0.57376	0.50982	0.45654	0.40992	0.36797			
7	0.68174	0.60218	0.54250	0.49277	0.44926	0.41010	0.37420		
8	0.70046	0.62558	0.56941	0.52261	0.48166	0.44480	0.41102	0.37964	
9	0.71623	0.64528	0.59208	0.54774	0.50894	0.47402	0.44201	0.41229	0.38443
10	0.72974	0.66217	0.61150	0.56927	0.53232	0.49907	0.46858	0.44028	0.41374
15	0.77671	0.72088	0.67901	0.64412	0.61360	0.58612	0.56093	0.53755	0.51562
20	0.80545	0.75682	0.72034	0.68994	0.66335	0.63941	0.61746	0.59709	0.57799
25	0.82536	0.78170	0.74895	0.72167	0.69779	0.67630	0.65660	0.63831	0.62116
30	0.84019	0.80023	0.77027	0.74530	0.72345	0.70378	0.68576	0.66902	0.65333
35	0.85178	0.81473	0.78694	0.76378	0.74351	0.72528	0.70856	0.69303	0.67848
40	0.86117	0.82646	0.80043	0.77874	0.75976	0.74268	0.72702	0.71248	0.69885
45	0.86898	0.83622	0.81166	0.79118	0.77327	0.75715	0.74237	0.72865	0.71578
50	0.87560	0.84450	0.82117	0.80174	0.78473	0.76942	0.75539	0.74236	0.73014

Table 13: Expected seek time for IDA systems ($\alpha= 0.50$)

Total Disks	Number of tolerable faults								
	0	1	2	3	4	5	6	7	8
1	0.33333								
2	0.46667	0.26667							
3	0.54286	0.37143	0.24286						
4	0.59365	0.44127	0.32698	0.23175					
5	0.63059	0.49206	0.38817	0.30159	0.22583				
6	0.65901	0.53114	0.43523	0.35531	0.28538	0.22244			
7	0.68174	0.56239	0.47288	0.39829	0.33302	0.27428	0.22044		
8	0.70046	0.58813	0.50389	0.43369	0.37226	0.31697	0.26629	0.21923	
9	0.71623	0.60981	0.53000	0.46349	0.40530	0.35292	0.30491	0.26033	0.21853
10	0.72974	0.62839	0.55238	0.48904	0.43361	0.38373	0.33801	0.29555	0.25574
15	0.77671	0.69297	0.63017	0.57783	0.53204	0.49083	0.45305	0.41797	0.38508
20	0.80545	0.73250	0.67778	0.63219	0.59229	0.55638	0.52347	0.49291	0.46425
25	0.82536	0.75987	0.71075	0.66982	0.63400	0.60177	0.57222	0.54479	0.51906
30	0.84019	0.78026	0.73531	0.69785	0.66508	0.63558	0.60854	0.58343	0.55990
35	0.85178	0.79620	0.75451	0.71978	0.68938	0.66202	0.63695	0.61366	0.59183
40	0.86117	0.80911	0.77007	0.73753	0.70906	0.68343	0.65995	0.63814	0.61769
45	0.86898	0.81984	0.78299	0.75229	0.72542	0.70123	0.67907	0.65848	0.63918
50	0.87560	0.82895	0.79396	0.76480	0.73929	0.71633	0.69528	0.67574	0.65742

Table 14: Expected seek time for IDA systems ($\alpha= 0.75$)