

PARAID: A Gear-Shifting Power-Aware RAID

CHARLES WEDDLE, MATHEW OLDHAM, JIN QIAN, and AN-I ANDY WANG

Florida State University

PETER REIHER

University of California, Los Angeles

and

GEOFF KUENNING

Harvey Mudd College

Reducing power consumption for server-class computers is important, since increased energy usage causes more heat dissipation, greater cooling requirements, reduced computational density, and higher operating costs. For a typical data center, storage accounts for 27% of energy consumption. Conventional server-class RAID configurations cannot easily reduce power because loads are balanced to use all disks, even for light loads.

We have built the power-aware RAID (PARAID), which reduces energy use of commodity server-class disks without specialized hardware. PARAID uses a skewed striping pattern to adapt to the system load by varying the number of powered disks. By spinning disks down during light loads, PARAID can reduce power consumption, while still meeting performance demands, by matching the number of powered disks to the system load. Reliability is achieved by limiting disk power cycles and using different RAID encoding schemes. Based on our five-disk prototype, PARAID uses up to 34% less power than conventional RAID configurations while achieving similar performance and reliability.

Categories and Subject Descriptors: D.4.2 [Operating Systems]: Storage Management—*Secondary storage*; D.4.8 [Operating Systems]: Performance—*Measurements*

General Terms: Design, Experimentation, Measurement, Performance

Additional Key Words and Phrases: Power savings, energy efficiency, RAID

This research was supported by NSF Grant CNS-0410896. Opinions, findings, and conclusions or recommendations expressed in this document do not necessarily reflect the views of the NSF, FSU, UCLA, Harvey Mudd College, or the U.S. Government.

Authors' addresses: C. Weddle, M. Oldham, J. Qian, and A.-I.A. Wang, Department of Computer Science, Florida State University, Tallahassee, FL 32306; email: {weddle, oldham, qian, awang}@cs.fsu.edu; P. Reiher, Department of Computer Science, University of California, Los Angeles, 405 Hilgard Ave., Los Angeles, CA 90095; email: reiher@cs.ucla.edu; G. Kuenning, Department of Computer Science, Harvey Mudd College, 301 Platt Blvd., Claremont, CA 91711; email: geoff@cs.hmc.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permission@acm.org. © 2007 ACM 1553-3077/2007/10-ART13 \$5.00 DOI = 10.1145/1289720.1289721 <http://doi.acm.org/10.1145/1289720.1289721>

13:2 • C. Weddle et al.

ACM Reference Format:

Weddle, C., Oldham, M., Qian, J., Wang, A.-I. A., Reiher, P., and Kuenning, G. 2007. PARaid: A gear-shifting power-aware RAID. *ACM Trans. Storage* 3, 3, Article 13 (October 2007), 33 pages. DOI = 10.1145/1289720.1289721 <http://doi.acm.org/10.1145/1289720.1289721>.

1. INTRODUCTION

The disk remains a significant source of power usage in modern systems. In web servers, disks typically account for 24% of the power usage; in proxy servers, 77% [Carrera et al. 2003; Huang et al. 2003]. Storage devices can account for as much as 27% of the electricity cost in a typical data center [Zhu et al. 2004a]. The energy spent to operate disks also has a cascading effect on other operating costs. Greater energy consumption leads to more heat dissipation, which in turn leads to greater cooling requirements [Moore et al. 2005]. The combined effect also limits the density of computer racks, which leads to more space requirements and thus higher operating costs.

Data centers that use large amounts of energy tend to rely on RAID to store much of their data, so improving the energy efficiency of RAID devices is a promising energy reduction approach for such installations. Achieving power savings on commodity server-class disks is challenging for many reasons, as follows.

- (1) RAID performance and reliability must be retained for a solution to be an acceptable alternative.
- (2) To reduce power, a server cannot rely on caching and powering-off disks during idle times because such opportunities are not as frequent on servers [Gurumurthi et al. 2003; Carrera et al. 2003; Zhu et al. 2004a].
- (3) Conventional RAID balances the load across all disks in the array for maximized disk parallelism and performance [Patterson et al. 1988], which means that all disks are spinning, even under a light load. To reduce power consumption, we must create opportunities to power-off individual disks.
- (4) Many legacy reliability encoding schemes rely on data and error-recovery blocks distributed among disks in constrained ways to avoid correlated failures. A solution needs to retro-fit legacy reliability encoding schemes transparently.
- (5) Server-class disks are not designed for frequent power cycles, which reduce life expectancy significantly. Therefore, a solution needs to use a limited number of power cycles to achieve significant energy savings.

Some existing approaches use powered-down RAIDs for archives [Colarelli and Grunwald 2002] or trade performance for energy savings [Pinheiro et al. 2001]. Some studies have exploited special hardware such as multispeed disks [Carrera et al. 2003; Li et al. 2004; Zhu et al. 2005]. Although simulation studies show promising energy savings, multispeed disks are still far from ubiquitous in large-scale deployments [Li et al. 2004; Yao and Wang 2006]. With the aid of nonvolatile RAM, approaches that use existing server-class drives have been recently made available [Li et al. 2004; Yao and Wang 2006; Pinheiro et al. 2006],




Fig. 1. UCLA Computer Science Department web server activity from August 11 through August 14, 2006.

but RAID reliability encoding constraints limit the number of spun-down drives (e.g., one for RAID-5).

We have designed, implemented, and measured the power-aware RAID (PARAID), which is deployable with commodity server-class disk drives without special hardware. PARAID introduces a skewed striping pattern that allows RAID devices to use just enough disks to meet the system load. PARAID can vary the number of powered-on disks by “gear-shifting,” or switching among sets of disks to reduce power consumption. Compared to a conventional 5-disk RAID, PARAID can reduce power consumption by up to 34% while maintaining comparable performance and reliability. Moreover, PARAID reuses different RAID levels so that the underlying RAID technology can evolve independently.

Beyond the power savings obtained by PARAID, the process of creating a real energy measurement framework produced some useful insights into the general problem of measuring energy consumption and savings. These insights are also discussed in this article.

2. OBSERVATIONS

Overprovisioned resources under RAID. Load balancing allows a conventional RAID device to maximize disk parallelism and performance, and ensures that no disk becomes a bottleneck. This uniformity simplifies data management and allows all disks to be accessed in the same way. However, uniform striping is not favorable for energy savings. Load balancing significantly reduces opportunities to power-off disks because all disks in the array need to be powered to serve a file, even if a RAID receives relatively light loads, when fewer powered disks would be sufficient.

Cyclic fluctuating load. Many system loads display cyclic fluctuations [Chase et al. 2001]. Figure 1 shows the web traffic gathered at the UCLA Computer Science Department across one week. The load fluctuations roughly follow daily cycles. Depending on the types of traffic, different systems may exhibit various fluctuation patterns, with varying ranges of light to heavy loads [Iyengar et al. 2000].

We can exploit these patterns by varying the number of powered disks, while still meeting performance needs and minimizing the number of power

13:4 • C. Weddle et al.

switches. A few strategically timed power cycles can achieve significant power savings.

Unused storage space. Storage capacity is outgrowing demand for many computing environments, and various large-scale installations report only 30% to 60% storage allocation [Asaro 2005; Gray 2005; Levin 2006]. Researchers have been looking for creative ways to use the free storage (e.g., trading-off capacity for performance [Yu et al. 2000] or storing every version of file updates [Santry et al. 1999]).

Additionally, many companies purchase storage with performance as the top criterion. Therefore, they may need many disks for parallelism to aggregate bandwidth, while the associated space is left largely unused. Further, administrators tend to purchase more space in advance to avoid frequent upgrades. Unused storage can then be used opportunistically for data-block replication to help reduce power consumption.

Performance versus energy optimizations. Performance benefits are realized only when a system is under a heavy load, and may not result in an immediate monetary return. Energy savings, however, are available at once and could, for example, be invested in more computers. Also, unlike performance, which is purchased in chunks as new machines are acquired, energy savings can be invested immediately and compounded over computer lifetimes. Therefore, if a server usually operates below its peak load, optimizing energy efficiency is attractive.

3. POWER-AWARE RAID

The main design issues for PARAID are how to skew disk striping to allow opportunities for energy savings, and how to preserve performance and reliability.

3.1 Skewed Striping for Energy Savings

PARAID exploits unused storage to replicate and stripe data blocks in a skewed fashion so that disks can be organized into hierarchical overlapping sets of RAIDs. Each set contains a different number of disks, and can serve all requests via either its data blocks or replicated blocks. Each set is analogous to a gear in automobiles, since different numbers of disks offer different levels of parallelism and aggregate disk bandwidth.

The replicated blocks are soft states in the sense that they can be easily reproduced. Thus, as storage demands rise, replicated blocks can be reclaimed by reducing the number of gears. Unlike memory caches, these soft states persist across reboots.

Figure 2 shows an example of replicated data blocks persisting in soft states in the unused disk regions. By organizing disks into gears, PARAID can operate in different modes. When operating in gear 1 with disks 1 and 2 powered, disks 3 and 4 can be powered off. As the load increases, PARAID upshifts into second gear by powering up the third disk.

By adjusting the number of gears and of disks in each gear, PARAID provides disk parallelism and bandwidth that can follow the fluctuating performance demand curve closely throughout the day. By creating opportunities to spin-down disk drives, PARAID conserves power.

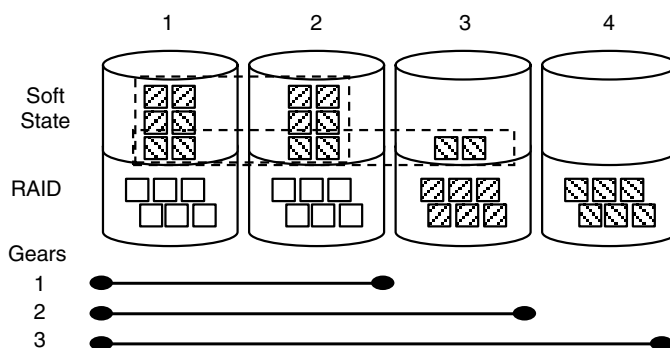


Fig. 2. Skewed striping of replicated blocks in soft state, creating three RAID gears over four disks.

While more gears can match the performance demand curve more closely, the number of gears is constrained by the unused storage available and the need for update propagation when switching gears. To minimize overhead, the gear configuration also needs to consider the number of gears and gear switches.

3.2 Preserving Peak Performance

PARAID matches the peak performance of conventional RAIDs by preserving the original disk layouts when operating at the highest gear. This constraint also allows PARAID to introduce minimal disturbances to the data path when the highest gear is in use.

In low gears, since PARAID offers less parallelism, the bandwidth offered is less than that of a conventional RAID. Fortunately, the number of requests affected by this performance degradation is significantly smaller compared to those affected during peak hours. Also, as bandwidth demand increases, PARAID will upshift the gear to increase disk parallelism.

PARAID also can potentially improve performance in low-gear settings. As a gear downshifts, the transfer of data to the soft state from disks about to be spun down warms up the cache, thus reducing the effect of seeking between blocks stored in different gears.

3.3 Retaining Reliability

To retain conventional RAID reliability, PARAID must be able to tolerate disk failures. To accomplish this goal, PARAID needs to supply the data redundancy of conventional RAIDs and address the reduced life expectancy of server-class disks due to power cycles.

PARAID is designed to be a device layer sitting between an arbitrary RAID device and its physical devices. Thus, PARAID inherits the level of data redundancy, striping granularity, and disk layout for the highest gear provided by that RAID. For example, a PARAID device composed with a RAID-5 device would still be able to rebuild a lost disk in the event of disk failure (the details of failure recovery will be discussed in Section 4.4).

Because PARAID power-cycles disks to save energy, it must also address a new reliability concern. Power-cycling reduces the MTTF of a disk, which is

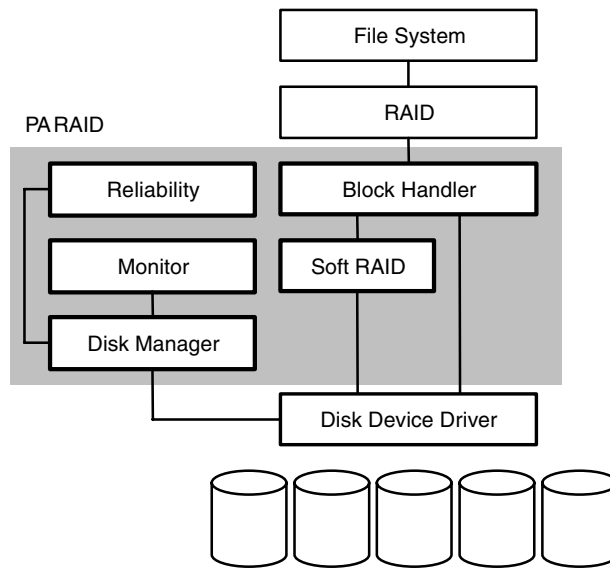


Fig. 3. PARAID system components.

designed for an expected number of cycles during its lifetime. For example, the disks used in this work have a 20,000-power-cycle rating [Fujitsu 2007]. Every time a disk is power-cycled, it comes closer to eventual failure.

PARAID limits the power-cycling of disks by inducing a bimodal distribution of busy and idle disks. Busier disks stay powered on, and more idle ones often stay off, leaving a set of middle-range disks that are power-cycled more frequently. PARAID can then prolong the MTTF of a PARAID device as a whole by rotating the gear-membership role of the disks and balancing their current number of power cycles.

Further, PARAID limits the number of disk power cycles used for gear shifting. By rationing power cycles, PARAID can operate with an eye to targeted life expectancy. For example, if the disks have a five-year life expectancy due to the system upgrade policy and the disks are expected to tolerate 20,000 cycles, then each disk in the array cannot be power-cycled more than 10 times a day. Once any of the disks has reached the rationed number of power cycles for a given period, PARAID can operate at the highest gear without energy savings. The power-saving mode resumes at the next rationing period.

4. PARAID COMPONENTS

PARAID has four major components (i.e., a block handler, monitor, reliability manager, and disk manager; see Figure 3) responsible for handling block I/Os, replication, gear shifting, update propagation, and reliability.

4.1 Disk Layout and Data Flow

PARAID is a new device layer in the conventional software RAID multidevice driver. The block handler under PARAID transparently remaps requests from a

Table I. PARAID Disk Layout with One 4-Disk Gear and One 5-Disk Gear, Each Running RAID-5.

	Disk 1	Disk 2	Disk 3	Disk 4	Disk 5
Gear 1 RAID-5	(1-4)	8	12	((1-4), 8, 12)	
	16	20	(16, 20, -)	-	
Gear 2 RAID-5	1	2	3	4	(1-4)
	5	6	7	(5-8)	8
	9	10	(9-12)	11	12
	13	(13-16)	14	15	16
	(17-20)	17	18	19	20

Each table entry contains either a single block number or a parenthesized list, denoting a parity block for the blocks in the list. “-” means an empty block.

conventional RAID device and forwards them to other soft-state RAID devices or individual disk devices.

PARAID currently delegates RAID regions to store replicated soft states for individual gears. The highest gear reuses the original RAID level and disk layout to preserve peak performance. When the highest gear is active, PARAID forwards requests and replies with minimal disturbance to the data path.

However, the data and parity blocks of D disks cannot be striped across fewer disks to achieve the same level of redundancy. If we simply assigned the D th block to one of the still-powered disks, it would be possible for a single drive to lose both a data block and a parity block from the same stripe, while the block stored on the powered-off disk might be out-of-date.

To provide reliability, the soft-state replicated blocks stored in each gear use the same RAID level. For example, consider a 5-disk RAID-5 (see Table I). Gear 2 uses all 5 disks; gear 1 uses 4. When disk 5 is spun down, its blocks must be stored on the remaining 4 disks. This is done by creating a 4-disk soft-state RAID-5 partition; the data and parity blocks from disk 5 are stored in this partition as if they were normal data blocks arriving directly from the application. If necessary, the soft-state partition can be removed to recover space whenever disk 5 is spinning.

The synchronization between disk 5 of gear 2 and the blocks in gear 1 resembles the data flow of RAID1+0. Disk 5 is “mirrored” using RAID-5 on gear 1, with synchronization performed during gear shifts. By using the underlying RAID-5 code for disk layout and parity handling, the PARAID code is drastically simplified as compared to trying to deal with those details internally.

For all gears (including the case where all disks are powered), if either a read or write request is sent to a powered disk, the disk simply serves the request. If a request is sent to a powered-off disk, then PARAID will remap the request to a replicated block stored on a powered disk. For writes, a remapped update is later propagated to neighboring gears during gear shifts.

The required unused storage depends on the RAID level, number of gears, and number of disks in each gear. For RAID-5, $D > 3$ disks, M gears with G_i disks within the i th gear ($1 \leq i \leq M \leq (D - 3 + 1)$, $3 \leq G_i < G_{i+1} < G_M = D$), the percentage storage consumption S_i of the total RAID for the i th gear can

13:8 • C. Weddle et al.

be solved with M equations.

$$\begin{cases} \sum_{i=1}^M S_i = 1 \\ \left(\sum_{j=i}^M S_j \right) (G_i - G_{i-1}) = (G_{i-1} - 1) S_{i-1}, \quad i = 2..M \end{cases} \quad (1)$$

For a disk in the lowest gear (Figure 2, disk 1), the sum of the percentage usage of disk space by each gear must be one. Also, for a gear (Figure 2, gear 2) to be able to shift to a lower gear (Figure 2, gear 1), the lower gear must store all the content of the disk(s) (Figure 2, disk 3) that are about to be spun down, with their parity information created for the lower gear.

PARAID uses around $(D - G_1)/(D - 1)$ of the total RAID-5 storage to store soft states. This estimate is largely based on the number of disks in the lowest gear, not the number of gears or the number of disks in intermediate gears, so the overhead of gear switching and the time spent in each gear will determine optimal gear configurations.

The target percentage of energy savings for an active system (not specific to RAID-5) is described by formula (2), where P_{standby} is the power consumption for a spun-down disk (which is typically not zero, for reasons described in the performance section), and $P_{\text{active/idle}}$ is the average power consumption for either a busy or idle disk to compute disk power-savings for busy or idle loads.

Power savings increase with more disks, fewer disks in the lowest gear, and a higher $P_{\text{active}}/P_{\text{standby}}$ ratio. Since spun-down disks still consume power, it is better to install PARAID with large disks with unused space rather than buying more disks later.

$$\text{Percentage_energy_savings} = 1 - \frac{(D - G_1)P_{\text{standby}} + G_1P_{\text{active/idle}}}{DP_{\text{active/idle}}} \quad (2)$$

For this article, an upshift means switching from a gear with G_i disks to G_{i+1} disks; a downshift, switching from a gear with G_i disks to G_{i-1} disks. A gear switch can be either an upshift or downshift.

4.2 Update Propagation

When a powered-off disk misses a write request, it must synchronize to replace the stale data either when powered on or just before the stale information is accessed. If there is a lot of stale data, fully synchronizing a disk can be slow. The on-demand approach updates stale data only when it is accessed, allowing the gear shift to take place much more swiftly, but the full synchronization approach is simpler to build. The on-demand approach is not applicable for downshifts, since PARAID needs to finish the propagation before spinning-down drives.

The disk manager captures outstanding writes to powered-off disks. For full synchronization, the disk manager reissues outstanding writes to the disk when it is powered on, possibly rereading some data from replicated soft states stored in the current gear.

For on-demand synchronization, the PARAID block I/O handler uses a dirty-block list. If a referenced dirty block is not cached, PARAID will retrieve the block from the original gear and return it to the requestor. PARAID will then write that block to the target-gear disks, effectively piggybacking the synchronization step at access time.

The disk manager must track stale block locations for synchronization. This list of dirty blocks is stored on disk in case of system failure and is replicated in memory for fast access.

A failed disk can stop the gear-shifting process. Disks can also fail during synchronization. However, the list of outstanding writes is maintained and replicated with the same RAID mechanism throughout the disk failure and recovery process. Once the failed disk recovers, synchronization can resume.

The choice of on-demand or full synchronization for upshifting is configurable. On-demand allows PARAID to be more responsive to sudden request bursts, at the cost of tracking additional writes for unsynchronized disks. The full synchronization approach may be preferable for few gear shifts and a read-dominated workload, since the number of blocks to be synchronized is small. The full synchronization method is also available for manual maintenance, such as when an administrator would need to have a consistent system state before pulling out a disk.

4.3 Asymmetric Gear-Shifting Policies

The disk manager performs shifts between gears. The PARAID monitor decides when a shift is needed, and the disk manager then performs the actual power cycles.

Switching to a higher gear is aggressive so that the PARAID device can respond quickly to a sharp and sustained increase in workload. However, the algorithm should be resilient to short bursts, or it will lead to little energy savings. Downshifting needs to be done conservatively so that wild swings in system activity will not: (1) mislead the PARAID device into a gear that cannot handle the requests; or (2) cause rapid oscillations between gears and significantly shorten the life expectancy of disks (or, considering disk power-cycle rationing, quickly turn off the gear shifting behavior for the day, limiting possibilities for power savings).

Upshifts. To decide when to upshift, the monitor must know whether the current gear has reached a predetermined utilization threshold in terms of busy RAID milliseconds within a time window. Interestingly, we could not check the disk-busy status directly, since this probe would spin up a powered-down disk. Instead, an active RAID device is marked busy from the point when a request enters the RAID queue to when the completion callback function is invoked. Since multiple RAID requests can overlap, should a request be completed with an elapsed time of t milliseconds, we mark the prior t milliseconds as busy.

The threshold and time window are static but configurable, and are set to 80% (based on prior studies [Carrera et al. 2003]) and 32 seconds (based on empirical experience), respectively. The intent is that within the time it takes

13:10 • C. Weddle et al.

to spin up the disk and propagate updates, the utilization threshold will not reach 100%. The use of an online algorithm to set thresholds automatically will be future work.

To track the system load, the monitor keeps a moving average of utilization $0 \leq U \leq 1$ for each gear. The purpose of averaging is to filter out short bursts of requests that are frequently seen in real-world workloads. The monitor also keeps a moving standard deviation S . If the utilization plus its standard deviation exceeds the threshold $0 \leq T \leq 1$, an upshift is performed.

$$U + S > T \quad (\text{upshift condition}) \quad (3)$$

The addition of standard deviation makes upshift more aggressive; however, since both the moving average and standard deviation lag behind the actual load, the policy is more responsive to changes that lead to sustained activities.

Downshifts. To decide when to downshift, the utilization of the lower gear $0 \leq U' \leq 1$ needs to be computed as given next, with associated moving standard deviation S' . If their sum is below the threshold T , the lower gear can now handle the resulting load, with associated fluctuations.

$$U' + S' < T \quad (\text{downshift condition}) \quad (4)$$

A complication arises when each gear is stored in RAID with parity blocks. Take Table I as an example. Suppose gear 2 contains a 5-disk RAID-5, and the 5th disk is replicated in gear 1 with a 4-disk RAID-5. After a downshift (i.e., spinning down the 5th disk), a write to blocks 1, 2, 3, 4, 8, 12, 16, and 20 will result in a parity update for gear 2, and another parity update for gear 1. Therefore, to compute the downshift threshold, the monitor must track recent write activity and inflate the percentage of write accesses A_{write} to those blocks by a weight W of 1.5x (specific to RAID-5, where writes to 1 data block and 1 parity block can be increased to 1 data block and 2 parity blocks). Otherwise, the lower gear will be unable to handle the resulting load, and will shift back up. Therefore, assuming a sufficiently large number of data blocks where the discrete modulo effects of the parity block locations can be overlooked, U' is computed with the following formula.

$$U' = U \frac{G_i}{G_{i-1}} \left[A_{\text{read}} + A_{\text{write}} \left(\frac{G_{i-1}(G_{i-1} - 1)}{G_i(G_i - 1)} (1 - W) + W \right) \right] \quad (5)$$

4.4 Reliability

The reliability manager rations power cycles and exchanges the roles of gear memberships to prolong the life expectancy of the entire PAROID. The reliability manager is also responsible for recovering a PAROID device upon disk failure. When PAROID fails at the highest gear, the recovery is performed by the RAID of the highest gear. When PAROID fails in other gears, the recovery is first performed by the lowest gear containing the failed disk, since the parity computed for the disks in that gear is sufficient to recover the soft states stored on the failed disk. The recovered soft-state data then is propagated to the next higher gear before recovering that gear, and so on. In the worst case,

the number of bytes needing to be recovered for a single-drive failure is the size of a single disk.

Although PARAID may take much longer to recover in the worst case due to cascaded recoveries, the average recovery time can be potentially reduced by recovering only modified content in the intermediate gears and frequent switching to the highest gears. To illustrate, should a PARAID host read-only content, recovery only involves switching to the highest gear and performing the recovery with the underlying RAID once, since no cascaded update propagations are needed. With modified content, PARAID can selectively recover only the modified stripes as well as stripes used to recover modified stripes at intermediate gears, and propagate them to the highest gear where a full recovery is performed. Assuming that 2% of disk content is modified per day [Kuenning and Popek 1997] and that PARAID switches to the highest gear 10 times a day, lightweight cascaded recovery is theoretically possible.

One might argue that PARAID can lengthen the recovery time, and thus reduce the availability of PARAID. On the other hand, PARAID reduces power consumption, and the associated heat reduction can extend drive life by about 1 percent per degree Celsius [Herbst 2007]. Therefore, the tradeoff requires further studies which are beyond the scope of this article.

5. IMPLEMENTATION

PARAID was prototyped on Linux (2.6.5), which was chosen for its open source and software RAID module. The block I/O handler, monitor, disk manager, and reliability manager are built as kernel modules. A PARAID user administration tool runs in user space to help manage the PARAID devices. For reliability, data blocks for all gears are protected by the same RAID level. Although we have not implemented drive rotations, our gear-shifting policies and the characteristics of daily work cycles have limited the number of disk power-cycles quite well. We have not implemented the mechanisms to recover only modified stripes in intermediate gears to speedup cascaded recovery.

Linux uses the `md` (multiple device) device driver module to build software RAIDs from multiple disks. For the PARAID block handler implementation, we changed `md` to make it PARAID-aware. The driver's data path is intercepted by the PARAID device layer so that requests from conventional RAID are redirected to the block queues of PARAID, which remaps and forwards requests to other conventional RAID-device queues.

During initialization, the PARAID-aware `md` module starts a daemon that provides heartbeats to the PARAID device and calls the monitor periodically to decide when to gear-shift. The disk manager controls the power status of disks through the disk device I/O control interface.

As an optimization, to limit the synchronization of content of a powered-off disk only to updated content, the disk manager keeps a per-disk red-black tree of references to outstanding blocks to be updated. This tree is changed whenever an update is made to a clean block on a powered-off disk. The upkeep of this data structure is not CPU-intensive. Currently, the disk manager synchronizes all modified blocks after bringing powered-off disks back online, by iterating

13:12 • C. Weddle et al.

through the tree for each disk and reissuing all outstanding writes. For each block to be synchronized, the disk manager reads the block from the original gear, and then writes it to the disks being brought back online. Once synchronization is complete, the gear-shifting manager switches to the new gear. Note that the red-black tree is only an optimization. In the case of losing this tree, gear content will still be fully propagated. A new tree can be constructed once PARaid gear switches to the highest gear.

Currently, PARaid serves requests from the current gear until the target gear completes synchronization, a conservative method chosen for implementation ease and to assure that no block dependency is violated through update ordering. In the future, we will explore using back pointers [Adb-El-Malek et al. 2005] to allow the new gear to be used during update propagation.

For the PARaid monitor, we currently use 32-second time windows to compute moving averages of disk utilization. The choice of this time window is somewhat arbitrary, but works well for our workloads and can tolerate traffic bursts and dampen the rate of power cycles. Further investigation of the gear-shifting triggering conditions will be future work.

The `mkraid` tool, commonly used by Linux to configure RAIDs, had to be changed to handle making PARaid devices and the insertion of entries in `/etc/raidtab`. Additional `raidtab` parameters had to be defined to be able to specify the gears. PARaid contains 3,392 lines of modifications to the Linux and `Raidtools` sourcecode. Since the PARaid logic is contained mostly in the Linux Software RAID implementation, it should be portable to future Linux kernel versions and software RAID implementations in other operating systems. We inserted four lines into `raid0.c` and `raid5.c` to set a flag to forward the resulting I/O requests to PARaid.

6. PERFORMANCE EVALUATION

Since the study of energy-efficient approaches to RAIDs is relatively recent, most prior work has been done analytically or via simulations. Analytical methods provide a fundamental understanding of systems. Simulation studies allow for the exploration of a vast parameter space to understand system behaviors under a wide range of scenarios. Both analysis and simulation, however, are limited in their validity by the degree to which their models match real-world behavior. We chose implementation and empirical measurements to see if we could overcome unforeseen physical obstacles and conceptual blind spots to bring us one step closer to a deployable prototype. When we designed, implemented, and evaluated PARaid, we discovered why an empirical study is difficult for systems designed to save energy.

- Prototyping PARaid was the first barrier, since the system had to be stable enough to withstand heavy benchmarking workloads.
- Commercial machines are not designed for energy measurements, so we had to rewire drives, power supplies, and probes for power measurements.
- The conceptual behaviors of a system are far from close to its physical behaviors; therefore, we had to adjust our design along the way.

- Most benchmarks and workload generators measure the peak performance of a system at steady state, which is not applicable for measuring energy savings, for which we need to capture daily workload fluctuations.
- For trace replays, since our physical system configuration was largely fixed, we had to try to match different trace environments with our physical environments in terms of the memory size, traffic volume, disk space consumption, and so on.
- Although many research trace replay tools are available, more sophisticated ones tend to involve kernel hooks and specific environments. Incompatibility of kernel versions prevented us from leveraging many research tools.
- Finally, since the process cannot be easily automated and inexpensively parallelized, measuring energy savings on a server was very time consuming.

Considering these challenges, we document the experimental settings we used to obtain our results. We demonstrate the power savings and performance characteristics of PARAID by replaying a web trace (Section 6.1) and the Cello99 trace [HP Labs 2007] (Section 7). The web workload contains 98% reads and is representative of a very large class of useful workloads. The Cello99 workload is I/O-intensive, and contains 58% reads. We used the PostMark benchmark [Katcher 1997] (Section 8) to demonstrate PARAID's performance under peak load. To demonstrate that PARAID can reuse different RAID levels, it was configured with RAID-0 for the web workload, and RAID-5 for the Cello99 workload. The PostMark benchmark stresses the gear-shifting overhead. All experiments were conducted five times. Error curves were removed from graphs for clarity. Generally, the standard deviations are within 5% of the measured values, with the exception of latency and bandwidth, which tend to be highly variable.

6.1 Web Trace Replay Framework

The measurement framework consisted of a Windows XP client and a Linux 2.6.5 server. The client performed trace playback and lightweight gathering of measurement results, and the server hosted a web server running on a RAID storage device [Fujitsu 2007] (Table II and Figure 4). On the server, one disk was used for bootstrapping and five were used to experiment with different RAID levels. The client and server were connected directly by a CAT-6 crossover cable to avoid interference from extraneous network traffic.

To measure power consumption, we used an Agilent 34970A digital multimeter. Each disk probe was connected to the multimeter on a unique channel, and the multimeter sent averaged data to the client once per second per channel via a universal serial bus.

To measure the power consumed by a given disk, we inserted a $0.1\text{-}\Omega$ resistor in series in the power-supply line (see Figure 5). The multimeter measured the voltage drop across the resistor, V_r . The current I through the resistor (which is also the current used by the disk) can be calculated as V_r/R . Given the voltage drop across the disk V_d , its power consumption is then V_d times I (this assumes

13:14 • C. Weddle et al.

Table II. Hardware Specifications

	Server	Client
Processor	Intel Xeon 2.8 Ghz	Intel Pentium 4 2.8 Ghz
Memory	512 Mbytes	1 Gbytes
Network	Gigabit Ethernet	Gigabit Ethernet
Disks [Fujitsu 2007]	Fujitsu MAP3367 36.7Gbytes 15K RPM SCSI Ultra 320 8MB on-disk cache 1 disk for booting 5 disks for RAID experiments Power consumption: 9.6 W (active) 6.5 W idle (spinning) 2.9 W standby (spun-down, empirically measured)	Seagate Barracuda ST3160023AS 160 Gbytes 7200 RPM SATA




Fig. 4. The measurement framework.

a relatively slow variation in a disk’s power requirements, which fortunately turned out to be the case).

In the measurement system, we removed each disk from the server and introduced a resistor into its +12V and +5V power lines. The +12V line supplied power to the spindle motor; the +5V line provided power to the disk electronics. The SCSI cable was connected directly to the motherboard, allowing the cable to maintain the same performance as if the disks were connected to the SCSI hot swappable backplane in the server.

On the client, the Agilent multimeter software logged the data using Microsoft Excel XP. The multithreaded trace driver, implemented in Java 1.5, was designed to replay web access log traces and collect performance numbers. Associated requests generated from the same IP address are each handled by a separate thread to emulate users clicking through webpages. The trace driver also collected server-side and end-to-end performance numbers.

The server hosted an Apache 2.0.52 web server on top of an ext2 file system operating over a RAID storage device that is described in Table II.




Fig. 5. The resistor inserted in series between the power supply and disk adapter.




Fig. 6. Energy consumption characteristics of PARAID versus conventional RAID.

6.2 Web Server Workload

Workload characteristics affect PARAID's ability to save energy. Under a constant high load, PARAID will not have opportunities to downshift and save energy. Under a constant light load, trivial techniques like turning everything on and off can be used to save energy (see Figure 6). In practice, though, workloads tend to show cyclic fluctuations. The chosen workload needs to capture these fluctuations to demonstrate PARAID's energy savings.

We chose a web server workload from the UCLA Computer Science Department. Since the web content is stored in a decentralized fashion via NFS mounts, we only report the hardware configuration of the top-level web server, which is a Sun Ultra-2 with 256Mbytes of RAM, 200 Mhz UltraSPARC I CPU, one 2Gbyte system disk and one 18Gbyte external SCSI disk, running Apache 1.3.27. Activity was captured from August 10 to August 16, 2006. Various NFS file systems contained approximately 32Gbytes of data and $\sim 500K$ files. We recreated the file system based on the referenced files in the trace. For each full path referenced, every directory in the full path and the file was created according to the order of replay. The file blocks stored on the web server were refilled with random bits. Also, the replay did not include dynamic file content, which accounts for the relatively few references in this trace.

We chose a 30-hour segment of the whole trace starting from 6 PM, August 12, 2006. The duration included around 95,000 requests, with 4.2Gbytes of data, of which 255Mbytes are unique. Although the workload is light, it captures

13:16 • C. Weddle et al.

the essence of read-mostly cyclic loads and sheds light on PARAID system behaviors, gear-shifting overhead, and the practical implementation limits on power savings.

6.3 Web Trace Replay Experimental Settings

PARAID was compared with a RAID-0 device. The PARAID device used 5 disks, with 2 disks in gear 1 and 5 in gear 2. Both client and server were rebooted before each experiment and PARAID was configured to start with the lowest gear, with gear content prepopulated. The client replayed trace log entries to the server. Due to the hardware mismatch and light trace workload, the collected trace was accelerated at different speeds to illustrate the range of possible savings with different levels of workloads. Experiments included a 256x speedup, which is close to a zero-think-time model, translating to 241 requests/second. With this reference point, we lowered the speedup factor to 128x and 64x, which correspond to 121 and 60 requests/second. All three loads offer few opportunities for the entire 5-disk RAID to be power-switched as a whole. Timing dependent on human interactions, such as the time between user mouse clicks on links (i.e., reference intervals by the same IP), was not accelerated.

6.4 Power Savings

Figure 7 compares the power consumption of PARAID and RAID-0. Due to the effects of averaging, power spikes are not visible.

PARAID demonstrates a 34% overall savings (ratios of areas under the curves) at 64x. The results approximately match the 33 to 42% range based on Eq. (2), indicating that further load reduction will yield limited energy benefits. However, turning off 3 out of 5 drives achieves nowhere near 60% energy savings, neither for PARAID nor other RAID systems that save power by spinning-down disks. Powering off a disk only stops it from spinning its platter and therefore only the 12V line is shut off. Power is still needed for the 5V line that powered the electronics, so that it can listen for a power-up command and pass commands along the daisy-chained SCSI cable.


Based on our measurements, spinning up a disk can consume 20 to 24W. Also, a spun-down disk still consumes 2.9W, noticeably higher than the 1.0W to 2.5W extracted from various datasheets and used in many simulations [Gurumurthi et al. 2003; Huang et al. 2003; Pinheiro and Bianchini 2004; Zhu et al. 2004, 2004b, 2005]. The results show that variations in physical characteristics can change the expected energy savings significantly. In our case, if we replace our Fujitsu [2007] with the commonly cited IBM Ultrastar 36Z15 [IBM 2007], we anticipate an additional 5% of energy savings.

Table III summarizes the overall PARAID energy savings.


The second observation is that the traffic pattern observed in the web log does not correlate well with the disk power consumption. Although this finding reveals more about the nature of caching than the energy benefits of PARAID, it does suggest the value of further investigations into the relationship between server-level and after-cache device-level activities.

Table III. Percent Energy Saved for Web Replay


Speed-up	Power savings
256x (241 req/sec)	10%
128x (121 req/sec)	28%
64x (60 req/sec)	34%



(a) 256x speedup




(b) 128x speedup



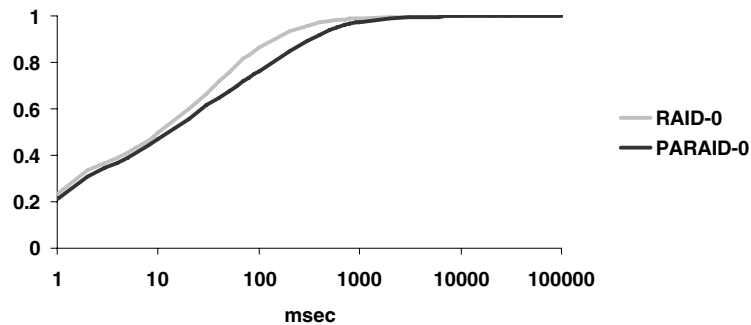
(c) 64x speedup

Fig. 7. Power consumption for web replay.


13:18 • C. Weddle et al.



(a) 256x speedup



(b) 128x speedup



(c) 64x speedup

Fig. 8. Latency for web replay.

6.5 Performance

Latency. Figure 8 shows the CDFs of per-request latency, which measures the time from the last byte of the request sent from the client to the first byte of data received at the client.

As expected, when playing back the trace at high speed, PARAID spent more time at the high gear and used the original RAID-5 disk layout, and the latency CDFs matched closely. The average latency is

Table IV. PARAID Gear-Switching Statistics for Web Replay

	256x	128x	64x
Number of gear switches	15.2	8.0	2.0
% time spent in low gear	52%	88%	98%
% extra I/Os for update propagations	0.63%	0.37%	0.21%

within 2.7% (~840ms). The data path overhead of PARAID is negligible (see Section 8).

When the load was light at 64x, PARAID spent most of the time at the lower gear. PARAID-0 had to use 2 disks to consolidate requests for 5 disks. As a result, the average latency of PARAID-0 was 80ms compared to 33ms for RAID-0. However, a web end-user should not notice the response time difference during light loads.

Bandwidth. Figure 9 shows the bandwidth over time, which measures the number of bytes transferred in a 30-minute interval, divided by the time the client spent waiting for any request to complete within the same interval.


As expected, when PARAID operates mostly in low gear, having fewer active disks leads to lower bandwidth numbers during light loads (Figure 9(c)): 24MB/sec as opposed to 31MB/sec for RAID-0. However, during the time intervals when PARAID operates in high gear, the peak load bandwidth matches well with the original RAID (within 1.3% of 32MB/sec). Note that due to time-based data alignment and averaging effects, Figure 9(a) only shows a close bandwidth match when PARAID's high-gear performance dominates within a time bracket. Section 7 will also explore request-based alignment to demonstrate bandwidth matching.

Gear-switching statistics. Table IV summarizes various PARAID gear-switching statistics for the web replay experiment. Clearly, PARAID spends more time in the low gear as the intensity of workload decreases with the replay speed. Also, each gear switch introduces up to 0.1% extra system I/Os. Interestingly, frequent gear switches can reduce the per-switch cost down to 0.041%, since less time is available for updates to accumulate at a given gear.


7. HP CELLO99 REPLAY

The HP Cello99 trace [HP Labs 2007] is an SCSI-controller-level trace collected by the HP Storage Research Lab from January 14 to December 31, 1999. The Cello99 data represents IO-intensive activity with writes, which is in contrast to the read-mostly UCLA web trace with lighter traffic. The traced machine had 4 PA-RISC CPUs, and some devices were md devices, so we had to extract a trace that neither overwhelmed our system nor produced too little traffic. The spc formatted trace file was generated from the Cello99 data using SRT2txt, a program that comes with the HP Cello99 data. The generated trace file was further trimmed so that only the activity associated with lun 2 was used. Also, we looked for traces with cyclic behaviors. The extracted trace contains 50 hours beginning on September 12, 1999, consisting of ~1.5 million requests totaling 12GB (stored in 110K unique blocks).


13:20 • C. Weddle et al.



(a) 256x speedup



(b) 128x speedup




(c) 64x speedup


Fig. 9. Bandwidth for web replay.

7.1 Cello99 Experimental Settings


PARAID was compared this time with a RAID-5 device. We used a 3-disk gear and 5-disk gear, each reusing the RAID-5 disk layout and reliability mechanisms. The Cello99 trace was replayed on the server at 128x, 64x, and 32x speedup factors to vary the intensity of workloads, corresponding to 1020, 548, and 274 requests/second. The energy measurement framework is the same as



(a) 128x speedup



(b) 64x speedup



(c) 32x speedup

Fig. 10. Power consumption for Cello99.

depicted in Figure 4. The server was rebooted before each run, with PARAID configured to start in the lowest gear.

7.2 Power Savings

Figure 10 compares the power consumptions of PARAID and RAID-5. PARAID demonstrates a single-point-in-time savings of 30% at 128x speedup (~13 hours into the replay) and 13% overall power savings at 32x speedup. Eq. (2) suggests

13:22 • C. Weddle et al.

Table V. Percent Energy Saved for Cello99

Speed-up	Power savings
128x (1024 req/sec)	3.5%
64x (548 req/sec)	8.2%
32x (274 req/sec)	13%

a power-saving range of 22 to 28%. Adjusted by the time spent at high gear (no energy savings), PARAID should have saved 17 to 22% at 32x, 14 to 18% at 64x, and 10 to 13% at 128x. Based on Table V, PARAID gear switches, update propagations, and the additional parity computation incur about 4–10% of energy overhead, a future goal for optimization. Nevertheless, despite the heavy load of 270–1000 requests/second, PARAID can still conserve up to 13% of power.

Figure 11 shows how gears are shifted based on the current gear utilization, defined as the percentage of busy seconds within a 32-second window, and on adjusted utilization as if the workload were using the low gear. PARAID consolidates the load spread among 5 disks to 3 disks, so that disks 4 and 5 can be spun off while disks 1 to 3 can operate at 10–40% utilization. The graph also reconfirms the lack of opportunities to power-switch the entire RAID for power savings.

7.3 Performance

Completion time. Figure 12 shows the CDFs of completion time (from the time of PARAID forwarding a request to the moment the corresponding complete callback function is invoked). Latency is more difficult to measure, since blocks are served out of order and individual blocks from various disks need to be demultiplexed to the corresponding multiblock request to gather latency information. Therefore, completion time, which is also the worst-case bound for latency, is used.

Unlike the latency CDFs from the web trace, the completion time CDFs of Cello99 showed very similar trends between PARAID and RAID-5. Figure 12 only presents the high 90th percentile in order to see the details of the CDF curves. At 32x, since PARAID spends more time at the lower gear, its latency is 26% slower than RAID-5 (1.8ms versus 1.4ms).

To better understand PARAID's performance, we examined the composition of the I/O requests. Although only 51% of bytes are accessed at high gear, they account for 97% of unique bytes. During light-load periods such as that between the 6th and 27th hours, only 29Mbytes of unique data were referenced. Given that each powered disk can use 5 Mbytes of on-disk cache, the bandwidth degradation of PARAID at low gear is significantly dampened by low-level caches. Therefore, the shape of completion time CDFs is dominated by the high-gear operation, which uses the same RAID-5.

Figure 13 shows bandwidth comparisons between PARAID and RAID-5. Note that these graphs are aligned by request numbers to emphasize that 60% of requests that occur during the peak load have the same bandwidth. Whenever PARAID is at high gear, the peak bandwidth is within 1% of RAID-5 (23MB/sec). The low average bandwidth at high-load periods reflects small average request








Fig. 11. Gear utilization for Cello99 replay. Utilization measures the percentage of busy time of the current gear. Adjusted utilization measures the percentage of busy time of the low gear if the workload is applied to the low gear.

sizes. During periods of light loads, the high bandwidth of both PARAID and RAID-5 reaffirms the enhanced role of low-level caches during light loads. Since PARAID did not use the SCSI controller, which contains additional cache, the bandwidth degradation of PARAID at low gear is likely to be further dampened. When PARAID operates at low replay speed and spends most of its time in low gear, the average bandwidth degrades as expected (12MB/s versus 21MB/s for RAID-5).

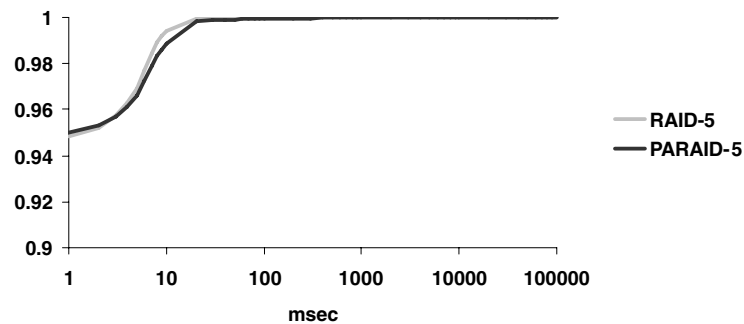
13:24 • C. Weddle et al.



(a) 128x speedup



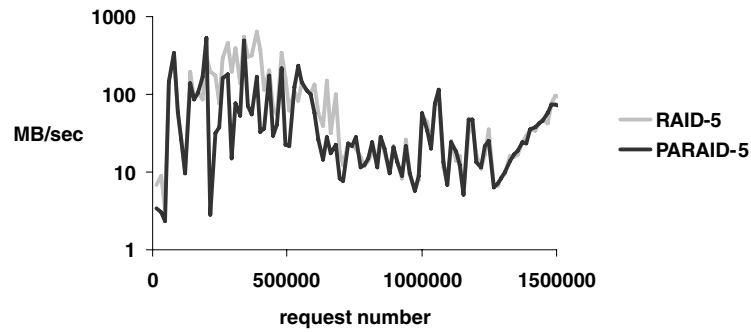
(b) 64x speedup



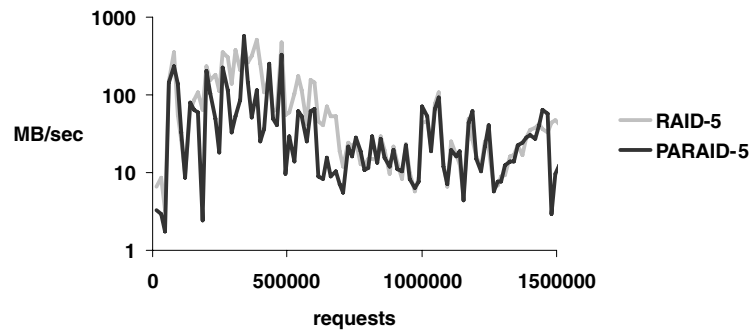
(c) 32x speedup

Fig. 12. Completion time for Cello99.


Gear-switching statistics. Table VI summarizes various PARAID gear-switching statistics for the Cello99 replay experiment. Again, PARAID spends more time in low gear with reduced workload with decreasing playback speed. Due to heavy updates, each gear switch needs to incur an extra 1.3% to 3.9% of system I/Os. Fortunately, gear shifting occurs either before the system becomes highly loaded or when it is about to downshift due to an upcoming period of



(a) 128x speedup



(b) 64x speedup



(c) 32x speedup

Fig. 13. Bandwidth for Cello99.

light loads. Therefore, these extra I/Os can be effectively absorbed by PARAID with spare I/O capacity which would otherwise be left unused.

8. POSTMARK BENCHMARK

The PostMark synthetic benchmark generates ISP-style workloads that stress a storage device's peak performance for its read- and write-intensive activity

13:26 • C. Weddle et al.

Table VI. PARAID Gear-Switching Statistics for Cello99

	128x	64x	32x
Number of gear switches	6.0	5.6	5.4
% time spent in low gear	47%	74%	88%
% extra I/Os for update propagations	8.0%	15%	21%

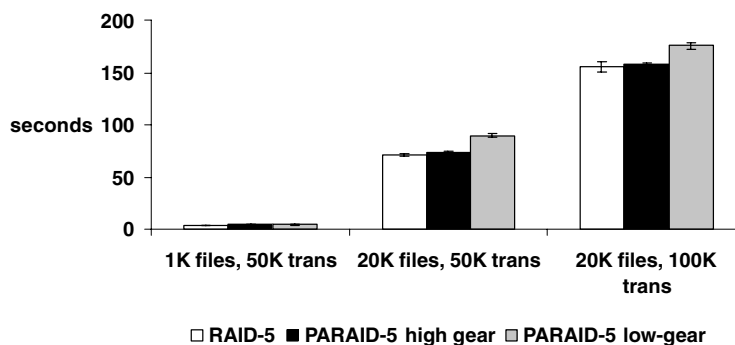


Fig. 14. PostMark results for a RAID-5 device compared to a PARAID device starting in the highest gear and starting in the lowest gear.

[Katcher 1997]. Running PostMark with PARAID starting at the lowest gear can be indicative of the overhead and latency of gear shifts during a request burst. The PostMark benchmark was run directly on the server. PARAID propagated updates synchronously during gear shifts.

Figure 14 presents PostMark results comparing the elapsed times of RAID 5, PARAID starting with the highest gear, and PARAID starting with the lowest gear under three different benchmark configurations.

For different PostMark configurations, PARAID starting with the highest gear demonstrates performance similar to RAID 5, which reflects the preserved layout of the underlying RAID and minimal disturbances to the md data path. Figure 15 shows that, as expected, PARAID does not save energy at the highest gear. PARAID's energy savings come primarily from low-gear settings.

Figure 14 also compares the performance of RAID-5 with PARAID starting in the lowest gear. It demonstrates how the current upshift policy prevents PARAID from being responsive to short bursts. The slowdown factor is about 13% due to upshift overhead. The most responsive approach is to upshift whenever a burst is encountered. However, this would cause too many gear shifts throughout a day. Our observations suggest that daily cyclic workloads cause few gear shifts, so this overhead is unnoticeable. We plan to explore online algorithms to improve PARAID's responsiveness to burst loads while minimizing the number of gear shifts.

Table VII demonstrates that PARAID in either configuration incurs similar CPU and system overhead when compared to RAID-5.

Table VII. PostMark CPU and System Overhead for RAID-5, PARAID Starting in Highest Gear and PARAID Starting in the Lowest Gear

	Mean% CPU	Mean% System
RAID-5	3.24%	41.18%
PARAID high gear	3.11%	41.60%
PARAID low gear	3.08%	41.93%

The experiment contains 20K files and 100K transactions.

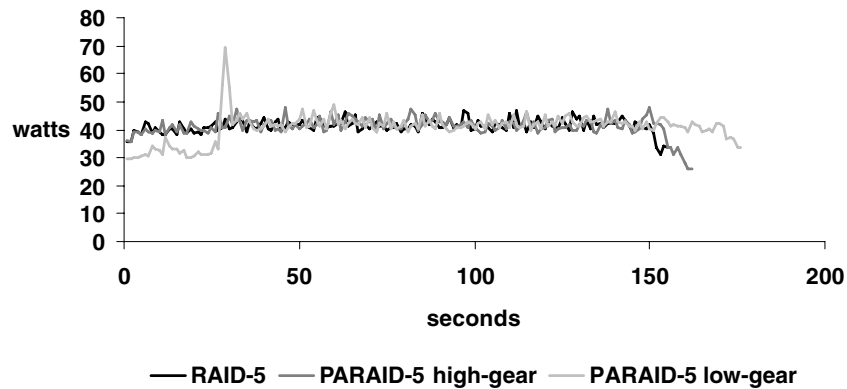


Fig. 15. The PostMark power consumption results for a RAID-5 device compared to a PARAID device starting in the highest gear and starting in the lowest gear. The experiment contains 20K files and 100K transactions.

9. RELATED WORK

Most energy reduction studies have addressed mobile computing [Douglass et al. 1995; Helmbold et al. 1996]. Recently, energy reduction for servers has also generated interest. Various approaches range from the hardware and RAID levels to the file system and server levels.

Reducing power consumption in hard disks. Carrera et al. [2003] suggest using hypothetical two-speed disks. During periods of high load, the disk runs at maximum speed and power. During periods of light load, the disk spins at a lower speed and possibly idles. The authors report simulated disk energy-savings between 15% to 22% for web servers and proxy servers, with throughput degradation of less than 5%. DRPM [Gurumurthi et al. 2003] assumes that disks with additional speed settings are available for power-performance tradeoffs.

FS2 [Huang et al. 2005] replicates blocks on a single disk to improve performance and reduce energy consumption via reducing seek time. FS2 reports up to 34% improvement in performance. The computed disk power consumption for per disk access also shows a 71% reduction. Since FS2 does not attempt to spin-down disks, and since PARAID has spare storage for disks in high gears due to skewed striping (see Figure 2), FS2 can be used on disks in high gears to extend PARAID's power savings.

13:28 • C. Weddle et al.

PCAP [Gniady et al. 2006] uses program counters to predict when a disk can be powered down. This approach mainly addresses laptop or desktop environments. In data server environments, aggregate workload fluctuations are a more suitable indicator of when to turn off disks.

Energy-efficient RAIDs. Hibernator [Zhu et al. 2005] aims to reduce the energy consumption of RAIDs without degrading performance through the use of multispeed disks. According to demand, data blocks are placed at different tiers of disks spinning at various speeds. A novel disk-block distribution scheme moves disk content among tiers to match disk speeds. When performance guarantees are violated, Hibernator spins disks at full speed to meet the demand. In simulation, Hibernator shows up to 65% energy savings.

Unlike Hibernator, PARAID is designed for existing server-class disks, and the minimum deployment granularity can be a small RAID on a typical server. Also, legacy systems can deploy PARAID via a software patch. As one consequence, some PARAID disks running at the lowest gear have few power-saving options. The future ubiquity of multispeed disks will allow PARAID to explore further energy savings when running at the lowest gear.

MAID [Colarelli and Grunwald 2002] assumes that the majority of data is being kept primarily for archival purposes, and its energy savings are based on the migration of this inactive majority to rarely used disks that fill a role similar to tape archives. PARAID, on the other hand, assumes that all data must be available at a high speed at all times. PARAID's techniques could be used on MAID's relatively few active disks to further improve the performance of that system.

Popular data concentration (PDC) [Pineiro and Bianchini 2004] saves energy by observing the relative popularity of data. PDC puts the most popular data on the first disk, the second most popular on the second disk, etc. Disks are powered off in PDC based on an idleness threshold. Without striping, PDC does not exploit disk parallelism.

In the absence of disk striping, the power-aware cache-management policy (PA-LRU) [Zhu et al. 2004a] saves power by caching data blocks from less active disks. Lengthening the access interval for less active disks allows them to be powered off for longer durations. Partition-based cache-management policy (PB-LRU) [Zhu et al. 2004b] divides the cache into separate partitions for each disk. Each partition is managed separately by a replacement algorithm such as LRU. Similar to PA-LRU, PB-LRU provides energy savings of 16%.

EERAID [Li et al. 2004] and its variant RIMAC [Yao and Wang 2006] assume the use of a nonvolatile cache at the disk-controller level and the knowledge of cache content to conserve energy in RAIDs. Both lengthen disk idle periods by using a nonvolatile disk-controller cache to delay writes and computing parity or data-block content on-the-fly. Both spin down at most one disk for RAID-5, which limits their power savings.

Pineiro et al. [2006] generalize RIMAC to erasure encoding schemes and demonstrate energy savings up to 61% in simulated tests. This approach defines and separates primary from redundant data and stores them on separate disks. Then, the system makes redundant data unavailable at times to save energy. Writes are buffered via nonvolatile RAM.

Energy-aware storage systems. BlueFS [Nightingale and Flinn 2004], a distributed file system, uses a flexible cache hierarchy to decide when and where to access data based on the energy characteristics of each device. Through empirical measurements, BlueFS achieved a 55% reduction in file system energy usage. Adding PARAID to BlueFS can further improve energy benefits.

The Conquest-2 file system [Xu et al. 2003] uses nonvolatile RAM to store small files to save the energy consumed by disks. PARAID can be readily used as a counterpart to serve large files while conserving energy.

Saving power in server clusters. Chase et al. [2001] and Pinheiro et al. [2001] have developed methods for energy-conscious server switching to improve the efficiency of server clusters at low request loads. They have reported energy reductions ranging from 29% to 43% for web server workloads.

PARAID can be combined with the server paradigm so that overprovisioned servers used to cushion highly bursty loads or prepowered to anticipate load increases can turn off many PARAID drives. Since powering-on disks is faster than booting servers, PARAID incurs less latency to respond to traffic bursts.

When traffic loads involve a mixture of reads and writes, disk switching in PARAID provides localized data movement and reduces stress on the network infrastructure. Also, PARAID can be deployed on individual machines without distributed coordination.

Other alternatives. Instead of implementing PARAID, one might use HP AutoRAID's reconfiguring ability to emulate PARAID's behavior [Wilkes et al. 1995]. However, one fundamental difference is that reconfiguring a RAID with D disks to $D - 1$ disks under AutoRAID requires restriping all content stored on D disks, while PARAID can restripe the content from a partial stripe, in this case 1 disk.

10. ONGOING WORK

PARAID is still a work in progress. First, although PARAID exploits cyclic fluctuations of workload to conserve energy, our experience with workloads suggests that it is difficult to predict the level of benefit based on traffic volume, the number of requests, number of unique bytes, peak-to-trough traffic ratios, and percentage of reads and writes. We are interested in measuring PARAID with diverse workloads to develop further understandings of PARAID's behavior. Also, we plan to test PARAID with other types of workloads, such as online transaction processing [UMass Trace Repository 2007].

Currently, PARAID is not optimized. The selection of the number of gears, number of disks in each gear, and gear-shifting policies are somewhat arbitrary. Since empirical measurement is unsuitable for exploring a large parameter space, we are constructing a PARAID-validated simulation for this purpose, which will allow more exploration of parameters. At the same time, we are investigating analytical approaches to develop online algorithms with provable optimality.

We will modify our disk synchronization scheme to explore asynchronous update propagation, allowing newly powered-on drives to serve requests immediately. We plan to implement selective recovery schemes for intermediate gears

13:30 • C. Weddle et al.

to speed-up cascaded recovery (currently, PARAID-5 as used in this article, recovers 2.7x slower than RAID-5), and also to incorporate the *SMART* tools [SANTools 2007] for disk health monitoring, allowing more informed decisions on rationing power cycles. Moreover, we will add rotation of the gear membership of disks. Finally, we plan to mirror a PARAID server on FSU's department server for live testing, deploy PARAID in a real-world environment, and compare PARAID with other energy-saving alternatives.

11. LEASONS LEARNED

The idea of PARAID was born as a simple concept to mimic the analogy of gear-shifting, which conserves fuel in vehicles. However, turning this concept into a kernel component for practical deployment has been much more difficult than we anticipated.

First, our early design and prototype of PARAID involved cloning and modifying RAID-0. As a result, we had to bear the burden of inventing replication-based reliability mechanisms to match different RAID levels. Our second-generation design can reuse the RAID encoding scheme, making the evolution of new RAID levels independent of PARAID. Although the resulting energy savings and performance characteristics were comparable in both implementations, PARAID's structural complexity, development time, and deployment potential improved in the new design.

Second, measuring energy consumption is difficult because of data-alignment problems and a lack of integrated tools. With continuous logging, aligning datasets is largely manual. For multithreaded experiments and physical disks, the alignment of datasets near the end of the experiment was significantly poorer than at the beginning. Early results obtained from averages were not explicable, since unaligned square waves can be averaged into nonsquare shapes.

Third, measuring systems under normal loads is harder than under peak loads. Replaying traces as quickly as possible was not an option, and we had to explore different speedup factors to see how PARAID reacts to loads changes. Since server loads have constant streams of requests, we cannot simply skip idle periods [Peek and Flinn 2005] because such opportunities are relatively infrequent. Worse, in our system consolidated workloads are handled by fewer powered-on components with less parallelism, further lengthening the measurement time.

Fourth, modern systems are complex. As modern hardware and operating systems use more complex optimizations, our perception of system behaviors increasingly deviates from their actual behaviors. Memory caching can reduce disk activity, while file systems can increase the burstiness of RAID traffic arrivals due to delayed write-back policies. Disks are powered with spikes of current, making it difficult to compute power consumption as the areas under the spike. Disk drives can still consume a significant amount of power even when they are spun down.

Fifth, matching the trace environment to our benchmarking environment is difficult. If we use a memory size larger than that of the traced machine, we may encounter very light disk activity. The opposite can saturate the disks

and achieve no power savings. Cyclic workload patterns before the cache may poorly reflect those after the cache. Additionally, traces might not have been made using RAIDs, some traces might be old, and the RAID geometry might not match our hardware settings. The base system might have multiple CPUs, which makes it difficult to judge whether a single modern CPU is powerful enough. Although the research community is well aware of these problems, the solutions still seem to be achieved largely by trial and error.

12. CONCLUSIONS

PARAID is a storage system designed to save energy for large computing installations that currently rely upon RAID systems to provide fast, reliable data storage. PARAID reuses standard RAID levels without special hardware, while decreasing their energy use by up to 34%. Since PARAID is not currently optimized, and since we measured only 5 drives (among which at least 2 are always powered), we believe that an optimized version of PARAID with many disks could achieve significantly greater energy savings.

A second important conclusion arises from the research described in this article. Actual implementation and measurement of energy-saving systems is vital, since many complex factors, such as caching policies, memory pressure, buffered writes, file-system-specific disk layouts, disk arm scheduling, and many physical characteristics of disk drives are difficult to capture fully and simultaneously validate using only simulation. Also, implementations need to address compatibility with legacy systems, the use of commodity hardware, and empirical evaluation techniques, all of which are necessary for practical deployments.

Unfortunately, our research also shows that there are considerable challenges to performing such experiments. We overcame several unforeseen difficulties in obtaining our test results, and had to invent techniques to do so. This experience suggests the value of developing standard methods of measuring the energy consumption of computer systems and their components under various conditions. We believe this is another fruitful area for study.

ACKNOWLEDGMENTS

We thank Keith Smith and the anonymous reviewers for their invaluable insights. Additionally, we thank Jason Flinn, Daniel Peek, Margo Seltzer, Daniel Ellard, Ningning Zhu, HP, and Sun Microsystems for providing accesses to various tools and traces. Christopher Meyers, Christine Savvakis, Sarah Diesburg, Cory Fox, Noriel Lu, Sean Toh, Daniel Beech, and Carl Owenby also have contributed to PARAID. This work is sponsored by NSF CNS-0410896. Opinions, findings, and conclusions or recommendations expressed in this document do not necessarily reflect the views of the NSF, FSU, UCLA, Harvey Mudd College, or the U.S. Government.

REFERENCES

- ADB-EL-MALEK, M., COURTRIGHT II, W. V., CRANOR, C., GRANGER, G. R., HENDRICKS, J., KLOSTERMAN, A. J., MESNIER, M., PRASAD, M., SALMON, B., SAMBASIVAN, R. R., SINNAMOHIDEEN, S., STRUNK, J. D., THERESKA, E., WACHS, M., AND WYLIE, J. J. 2005. URSA minor: Versatile cluster-based storage. In *Proceedings of the 4th USENIX Conference on File and Storage Technology*, San Francisco, CA.

13:32 • C. Weddle et al.

- ASARO, T. 2005. An introduction to thin provisioning. *Storage Technol. News*, http://searchstorage.techtarget.com/columnItem/0,294698,sid5_gci1134713,00.html.
- CARRERA, E., PINHEIRO, E., AND BIANCHINI, R. 2003. Conserving disk energy in network servers. In *Proceedings of the 17th Annual ACM International Conference on Super Computers*, San Francisco, CA.
- CHASE J., ANDERSON, D., THAKAR, P., VAHDAT, A., AND DOVAL, R. 2001. Managing energy and server resources in hosting centers. In *Proceedings of the 18th ACM Symposium on Operating System Principles*, Banff, Canada.
- COLARELLI, D. AND GRUNWALD, D. 2002. Massive arrays of idle disks For storage archives. In *Proceedings of the ACM/IEEE Conference on Supercomputing*, Baltimore, MD.
- DOUGLIS, F., KRISHNANA, P., AND BERSHAD, B. 1995. Adaptive disk spin-down policies for mobile computers. In *Proceedings of the 2nd USENIX Symposium on Mobile and Location-Independent Computing*, Ann Arbor, MI.
- FUJITSU, 2007. MAP 10K RPM. <http://www.fujitsu.com/us/services/computing/storage/hdd/discontinued/map-10k-rpm.html>.
- GNIADY, C., BUTT, A., HU, Y., AND LU, Y. H. 2006. Program counter-based prediction techniques for dynamic power management. *IEEE Trans. Comput.* 55, 6, 641–658.
- GRAY, J. 2005. Keynote address greetings from a filesystem user. In *the 4th USENIX Conference on File and Storage Technologies*, San Francisco, CA.
- GURUMURTHI, S., SIVASUBRAMANIAM, A., KANDEMIR, M., AND FRANKE, H. 2003. DRPM: Dynamic speed control for power management in server class disks. In *Proceedings of the International Symposium on Computer Architecture*, San Diego, CA.
- HELMBOLD, D. P., LONG, D. D. E., AND SHERROD, B. 1996. A dynamic disk spin-down technique for mobile computing. In *Proceedings of the 2nd Annual International Conference on Mobile Computing and Networking (MobiCom)*, Rye, NY.
- HERBST, G. 2007. Hitachi's drive temperature indicator processor (Drive-TIP) helps ensure high drive reliability. <http://www.hgst.com/hdd/technolo/drivetemp/drivetemp.htm>.
- HP Labs. 2007. Tools and traces. <http://www.hpl.hp.com/research/ssp/software/>.
- HUANG, H., HUNG, W., AND SHIN, K. G. 2005. FS2: dynamic data replication in free disk space for improving disk performance and energy consumption. In *Proceedings of the 20th Symposium on Operating Systems Principles*, Brighton, UK.
- HUANG, H., PILLAI, P., AND SHIN, K. G. 2003. Design and implementation of power aware virtual memory. In *Proceedings of the USENIX Annual Technical Conference*, San Antonio, TX.
- IBM. 2007. IBM hard disk drive ultrastar 36Z15. <http://www.hitachigst.com/hdd/ultra/ul36z15.htm>.
- IYENGAR, A., CHALLENGER, J., DIAS, D. AND DANTZIG, P. 2000. High-Performance web site design techniques. *IEEE Internet Comput.* 4, 2(Mar.), 17–26.
- KATCHER, J. 1997. PostMark: a new file system benchmark. Tech. Rep. TR3022, Network Appliance, Inc.
- KUENNING, G. H. AND POPEK, G. J. 1997. Automated hoarding for mobile computers. In *Proceedings of the 16th ACM Symposium on Operating Systems Principles*, Saint-Malo, France.
- LEVIN, M. 2006. Storage management disciplines are declining. *Comput. Econom.* <http://www.computereconomics.com/article.cfm?id=1129>.
- LI, D., GU, P., CAJ, H., AND WANG, J. 2004. EERAID: Energy efficient redundant and inexpensive disk array. In *Proceedings of the 11th ACM SIGOPS European Workshop*, Leuven, Belgium.
- MANLEY, S., SELTZER, M., AND COURAGE, M. 1998. A self-scaling and self-configuring benchmark for web servers. In *Proceedings of the ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*, Madison, WI.
- MILLER, E. AND KATZ. 1993. An analysis of file migration in a Unix supercomputing environment. In *Proceedings of the USENIX Winter Technical Conference*, San Diego, CA.
- MOORE, J., CHASE, J., RANGANATHAN, P., AND SHARMA R. 2005. Making scheduling “cool”: Temperature-Aware workload placement in data centers. In *Proceedings of the USENIX Annual Technical Conference*, Anaheim, CA.
- NIGHTINGALE, E. B. AND FLINN, J. 2004. Energy-Efficiency and storage flexibility in the blue file system. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation*, San Francisco, CA.

- PATTERSON, D. A., GIBSON, G., AND KATZ, R. H. 1988. A case for redundant arrays of inexpensive disks (RAID). In *Proceedings of the ACM SIGMOD International Conference on Management of Data 1*, 3(Jun.), 109–116.
- PEEK, D. AND FLINN, J. 2005. Drive-Thru: Fast, accurate evaluation of storage power management. In *Proceedings of the USENIX Annual Technical Conference*, Anaheim, CA.
- PINHEIRO, E., BIANCHINI, R., AND DUBNICKI, C. 2006. Exploiting redundancy to conserve energy in storage systems. In *Proceedings of the ACM SIGMETRICS Conference on Performance*, Saint-Malo, France.
- PINHEIRO, E. AND BIANCHINI, R. 2004. Energy conservation techniques for disk array-based servers. In *Proceedings of the 18th Annual ACM International Conference on Supercomputing*, Saint-Malo, France.
- PINHEIRO, E., BIANCHINI, R., CARRERA, E. V., AND HEATH, T. 2001. Load balancing and unbalancing for power and performance in cluster-based systems. In *Proceedings of the Workshop on Compilers and Operating Systems for Low Power*, Barcelona, Spain.
- RFC-3174. 2001. US secure hash algorithm 1, 2001. <http://www.faqs.org/rfcs/rfc3174.html>.
- SANTOOLS, INC. 2007. <http://www.santools.com/smartmon.html>.
- SANTRY, D. S., FEELEY, M. J., HUTCHINSON, N. C., VEITCH, A. C., CARTON, AND R. W., OFIR, J. 1999. Deciding when to forget in the elephant file system. In *Proceedings of the 17th ACM Symposium on Operating Systems Principles*, Charleston, SC.
- UMASS TRACE REPOSITORY. 2007. Storage traces. <http://traces.cs.umass.edu/index.php/Storage/Storage>.
- WILKES, J., GOLDING, R., STAELIN, C., AND SULLIVAN, T. 1995. The HP autoRAID hierarchical storage system. In *Proceedings of the 15th ACM Symposium on Operating Systems Principles*, Copper Mountain, CO.
- XU, R., WANG, A., KUENNING, G., REIHER, P., AND POPEK, G. 2003. Conquest: Combining battery-backed RAM and threshold-based storage scheme to conserve power. Work in progress report, in the *19th Symposium on Operating Systems Principles (SOSP)*, Bolton Landing, NY.
- YAO, X. AND WANG, J. 2006. RIMAC: A novel redundancy-based hierarchical cache architecture for energy efficient, high performance storage systems. In *Proceedings of the EuroSys*, Leuven, Belgium.
- YU, X., GUM, B., CHEN, Y., WANG, R., LI, K., KRISHNAMURTHY, A. AND ANDERSON, T. 2000. Trading capacity for performance in a disk array. In *Proceedings of the 4th Symposium on Operating Systems Design and Implementation*, San Diego, CA.
- ZHU, Q., CHEN, Z., TAN, L., ZHOU, Y., KEETON, K., AND WILKES, J. 2005. Hibernator: Helping disk arrays sleep through the winter. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*, Brighton, UK.
- ZHU, Q., DAVID, F.M., DEVARAJ, C., LI, Z., ZHOU, Y., AND CAO, P. 2004. Reducing energy consumption of disk storage using power-aware cache management. In *Proceedings of the 10th International Symposium on High Performance Computer Architecture*, Madrid, Spain.
- ZHU, Q., SHANKER, A., AND ZHOU, Y. 2004b. PB-LRU: A self-tuning power aware storage cache replacement algorithm for conserving disk energy. In *Proceedings of the 18th Annual ACM International Conference on Supercomputing*, Saint-Malo, France.

Received August 2007; accepted August 2007