Graceful Operation of Disk Drives under Thermal Emergencies

Youngjae Kim Jeonghwan Choi Sudhanva Gurumurthi† Anand Sivasubramaniam*

Dept. of Computer Science and Engineering
The Pennsylvania State University
University Park, PA 16802
{youkim, jechoi, anand}@cse.psu.edu

†Dept. of Computer Science
University of Virginia
Charlottesville, VA 22904
gurumurthi@cs.virginia.edu

ABSTRACT

Thermal-aware design of disk-drives is an important concern because high temperatures can cause reliability problems. Hence, Dynamic Thermal Management (DTM) has been proposed to operate the disk at the average case, rather than the worst case by modulating the activities to avoid thermal emergencies at run time while pushing the performance. A delay-based approach to adjust the disk seek activities is one DTM solution for normal disk-drives. Even if it could overcome thermal emergencies without stopping disk activity, it suffers from long delays when servicing the requests.

In this paper, we investigate the possibility of using a multi-speed disk-drive which dynamically modulates the rotational speed of the platter (called DRPM) for implementing DTM. Using a detailed performance and thermal simulator of storage system, we evaluate DTM policies and observe that the DRPM technique is the best to avoid thermal emergencies. However, we find that the time taken to transition between different rotational speeds of the disk is critical to the effectiveness of this DTM technique.

Keywords: Storage System, Disk Drives, Temperature Management.

1 INTRODUCTION

Thermal-awareness is becoming an integral aspect in the design of all computer system components, ranging from micro-architectural structures within processors to peripherals, server boxes, racks, and even entire machine rooms. This is increasingly important due to the growing power density at all the granularity of the system architecture. Deeper levels of integration, whether it be within a chip, or components within a server, or machines in a rack/room, cause a large amount of power to be dissipated in a much smaller footprint. Since the reliability of computing components is very sensitive to heat, it is crucial to drain away excess heat from this small footprint. At the same time, the design of cooling systems is becoming prohibitively expensive, especially for the commodity market [17][7]. Consequently, emerging technologies are attempting to instead build systems for the common case - which may not be subject to the peak power densities, and thereby operate at a lower cooling cost - and resort to Dynamic Thermal Management (DTM) solutions when temperatures exceed safe operational values. This paper explores one technique for implementing DTM for disk-drives.

Disk-drive performance is highly constrained by temperature. It can be improved by a combination of higher rotational speeds of the platters (called RPM), and higher recording densities. A higher RPM can provide a linear improvement in the data rate. However, the temperature rise in the drive enclosure can have nearly cubic relation to the RPM [3]. Such a rise in temperature can severely impact the reliable operation of the drive. Higher temperatures can cause instability in the recording media, thermal expansion of platters, and even out-gassing of spindle and voice-coil motor lubricants which can lead to head crashes [9]. One way of combating this generated heat is by reducing the platter sizes, which reduces the viscous dissipation by the fifth power. However, a smaller platter leads to a smaller disk capacity, unless more platters are added (in which case the viscous dissipation increases again by a linear factor). Moreover, a higher number of bits are necessary for storing Error Correcting Codes to maintain acceptable error rates due to lower signal-to-noise ratios in future disk-drives. All these factors make it difficult to sustain the continued 40% annual growth that we have been enjoying in the data rates until now [7]. This makes a strong case for building drives for the common case, with solutions built-in for dynamic thermal management when the need arises. DTM has been already implemented in Seagate Barracuda ES drive in the industry [14]. Using a feature of workload management, it helps prevent the temperature of disk-drive from rising because of a heavy workload.

There is one other important factor necessitating DTM. It is not enough to consider individual components of a computing system in isolation any more. These components are typically put together in servers, which are themselves densely packed in racks in machine rooms.

*Jeonghwan Choi visited PSU during his doctroal program at Korea Advanced Institute of Science and Technology (KAIST).

Provisioning a cooling system that can uniformly control the room so that all components are in an environment that matches the manufacturer specified “ambient” temperatures can be prohibitively expensive. With peak load surges on some components, parts of a room, etc., there could be localized thermal emergencies. Further, there could be events completely external to the computer systems - HVAC/fan break-down, machine room door left open, etc. - which can create thermal emergencies. Under such conditions, today’s disk-drives could overheat and fail, or some thermal monitor software could shut down the whole system. The disk is, thus, completely, unavailable during those periods. The need to sustain 24/7 availability, and growing power densities lead to the increased likelihood of thermal emergencies. This makes it necessary to provide a “graceful” operation mode for disk-drives. During this “graceful” mode, even if the disk is not performing as well as it would have when there was no such emergency, it would still continue to service requests, albeit slowly. This graceful mode would essentially be a period during which certain dynamic thermal management actions are carried out in the background, while continuing to service foreground requests.

Multi-speed disk operation [6][2] has been proposed as a solution to reduce disk-drive power, and can thus be a useful mechanism for thermal management as well. This mechanism is based on the observation that it is faster to change the rotational speed of a disk, rather than spinning it all the way down/up. DRPM allows the disk to service requests at a slower rate even at the lower RPM. During a thermal emergency, we can not only reduce the speed to reduce temperature, but we can continue to service requests at the lower speed. A commercial version of a multi-speed disk is sold by Hitachi [10], which provides two rotational speeds. Since the heat dissipated during the operation at a lower RPM is also much lower, the temperature within the drive can be lowered by employing this option during a thermal emergency. While a Hitachi’s multi-speed disk does provide a smaller window of time when the disk cannot service requests compared to a disk which only provides on-off modes, it still does not serve requests when it is at a lower RPM.

In this paper, we explore two options for temperature management during thermal emergency. We first consider disks which are tuned for maximum performance with the ideal/constant ambient temperature. We then introduce thermal emergencies - by adjusting the external ambient temperature of the drive - which pushes the drive into the emergency regions. We then investigate these two multi-speed drive options, and show that it is indeed possible for some regions of external ambient temperature variation to service disk requests even though such situations would have caused the drive to completely shut down in a non-multi-speed drive. As is to be expected, the performance during those periods is not as good as it would be when there are no emergencies. Between the two multi-speed options, not servicing the requests at the lower RPM causes frequent switches between RPMs, thus not faring as good as the DRPM disk in its availability. All these experiments are conducted using a detailed performance and thermal simulator of storage system, called STEAM [5].

The organization of the rest of this paper is as follows. Section 2 represents related work and Section 3 describes the micro-benchmark evaluations of a multi-speed disk under external ambient temperature variance. The experimental scenarios and results of DTM for real workloads are in Section 4. Finally Section 5 concludes this paper.

2 RELATED WORK

Temperature-aware design has been exploited for microprocessors [17], interconnection networks [15], storage systems [7] and even to the rack-mounted servers at machine rooms [16], because high temperature might cause unreliable operation to the system and require high cost of cooling. There have been various thermal simulation tools proposed to evaluate temperature-aware design. HotSpot [17] is a thermal simulator for microprocessors using thermal resistance and capacitance derived from the layout of microprocessor architecture. STEAM [5] is a performance and thermal simulator for disk-drives by using a finite difference method to calculate the heat flow and to capture the temperatures of different regions within the disk/storage system. A recent utility called Mercury [8] is a software suite to emulate temperatures at specific points of a server by using a simple flow equation. In addition, detailed simulations based on Computational Fluid Dynamics (CFD) provides more accurate thermal profiles by generating 3-dimensional thermal profiles within the system [12].

Dynamic Thermal Management has been adopted for individual components of the systems such as microprocessors [1][19] and disk-drives [11] or distributed environments such as distributed systems [23] and rack-mounted servers at data centers [16]. Of all these approaches, DTM for disk-drives has already been addressed [7]. One is from the observation that the seek activities might cause the temperature of disk-drive to be above thermal envelope (i.e., maximum operating temperature) because the voice-coil motor (VCM) (moving the disk arms) should be active for the seek activity. A delay-based DTM has been applied to prevent this situation from happening [4]. When the temperature of disk-drive reaches close to thermal envelope, DTM is invoked by stopping all the requests issued, where all the seek activities stop and the service resumes after the temperature is sufficiently reduced. However, even if this delay-based throttling by controlling the seek activities is feasible, many requests cannot be issued during thermal emergencies and thereby the performance is greatly affected by them. Today’s

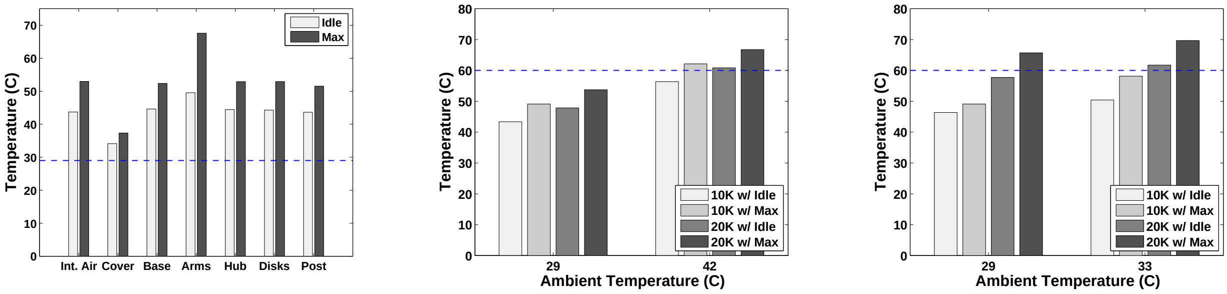

(a) Ultrastar 146Z10 (b) 3.3" 1 platter (c) 3.3" 4 platters

Figure 1. (a) shows temperature distribution of IBM Ultrastar 146Z10. Each bar denotes the average temperature of each component at the steady state for Max (where VCM is on all the times while the disk platters are spinning) and Idle (VCM just turns off). The dotted line denotes the ambient temperature (29°C). (b) and (c) are the steady-state base temperatures of disk for different disk dimensions (such as the size of platter and rotational speed of platter) and different power consumption modes under various ambient temperatures. The horizontal line in each graph is thermal envelope (60°C).

Seagate's Barracuda ES drives have a similar DTM feature by adjusting the workloads for thermal management [14]. The other possible approach is to modulate the RPM speed in a multi-speed disk. Since RPM has nearly cubic power relation to the viscous dissipation, it can be more effective to manage the temperature of disk-drive. This technique of dynamic RPM modulation for thermal management is discussed in the rest of this paper.

3 IMPACT OF THERMAL EMERGENCIES TO MULTI-SPEED DISK-DRIVE

In order to understand the heat distribution over all the components of a drive enclosure while it is in operation, we used STEAM to model Ultrastar 146Z10 disk-drive [21] installed in a 42U computer system rack. Ultrastar 146Z10 is composed of two 3.3" platters and rotates at 10K RPM. The power of spindle motor (SPM) (to rotate the platters) is derived from [3] and the VCM power (which is dependent on the platter dimensions) is obtained by applying the power-scaling curve from [18]. The power values of SPM and VCM are set to be 10W and 6.27W, and all other required parameters such as disk geometry were supplied as inputs into the model. From Figure 1-(a), we see that the hottest component over the drive enclosure is arm-assembly (which have the heads at their ends) whose temperature is around 68°C at Max (i.e., the disks are spinning and the arms are moving back and forth with VCM on all the time) and the lowest temperature is for the disk-cover surrounding disk-drive (around 37°C at Max). This is because the heat is directly drained away to the ambient air through the convection process. Idle (i.e., the disks are spinning but there are no arm movements) has a similar temperature distribution within disk-drive (though they are overall lower) than when it is at Max.

We performed a micro-benchmark evaluation to understand the impact of variation in ambient temperature to the disk and the feasibility of dynamic throttling by RPM modulation in a multi-speed disk. Figure 1-(a) shows that there is high thermal variation in temperature over the disk-drive, which means that thermal sensor location is very critical in applying DTM to a multi-speed disk. It is hard to decide the location that should be selected for detective emergencies. The base temperature of disk is chosen for DTM mechanism in this paper, because a thermal sensor is mounted on the back side of the electronics card close to the base of the actual disk-drive [9]. The highest RPM speed of a multi-speed disk is restricted to 20K and the baseline is 10K RPM. *Thermal Slack* is defined as temperature difference between current operating temperature and thermal envelope. We modeled two different disks for the experiments, where one is a 3.3" 1 platter disk-drive used in HPL Openmail and the other is a disk-drive with 3.3" 4 platters used for other workloads in Table 2.

We have measured the base temperature of disk with different ambient temperatures at the steady state in STEAM. We varied the ambient temperature from 29°C to 42°C for 3.3" 1 platter disk and from 29°C to 33°C for the disk of 3.3" 4 platters. In the experiment, the thermal envelope was set to be 60°C because the possible operating temperature range of disk-drive suggested in manuals is 5°C-60°C [21][13]. From Figure 1-(b), thermal emergencies never happen with even 20K rotational speed of platters and the VCM on all the time for 29°C ambient temperature. However, if it is increased further to 42°C ambient temperature, it could exceed thermal envelope (60°C). For example, a multi-speed disk operating at 20K under 42°C is above thermal envelope at both Idle and Max. However, if RPM drops down to 10K, it becomes below thermal envelope at Idle, while it exceeds 60°C at Max. This result shows that *Thermal Slack* would be around 4°C at the maximum. We also observe from Figure 1-(c) that, if the disk-drive has larger number of platters in a similar disk geometry, it is more prone to thermal emergencies even for small increase in ambient temperature because the heat dissipation inside the disk-drive is proportional to the number of platters. Figure 1-(c) shows that 33°C external ambient temperature introduces thermal emergency when it is operating at the higher speed. *Thermal Slack* becomes larger around 10°C at the maximum. Moreover, even if the disk-drive is operating at the lower speed, thermal emergency could exceed thermal envelope for even 29°C ambient temperature depending on the request patterns. Similarly, since the heat generated from the disk-drive is proportional to the fifth power of disk platter size, the disk-drive with the larger size of platters is more sensitive to variations in the ambient temperature.

Within these thermal slacks, a dynamic throttling mechanism can be applicable to avoid thermal emergencies. It can be achievable by pulling down the RPM speed of disk (when it reaches thermal emergencies). Once the temperature is lower than a given thermal envelope, it brings up the disk to full RPM after the cooling period.

4 DYNAMIC THERMAL MANAGEMENT

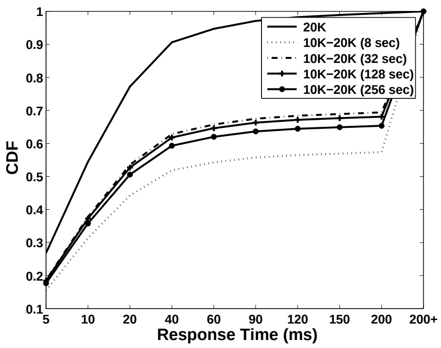

(a) HPL Openmail

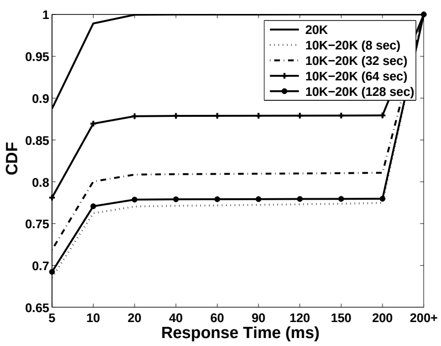

(b) TPC-C

Figure 2. Performance degradation of $DRPM_{simple}$ for the server workloads. Each graph shows the CDF of the average response time at I/O drivers across different disks. The solid curve at each graph shows when it is operating at the maximum speed of 20K RPM without DTM and others are applied by DTM. The value in parenthesis at each graph denotes a cooling unit time (which is given as a delay time, once it becomes close to thermal envelope (60°C)).

Thermal emergencies are generally caused by unexpected events, such as fan-breaks, increased inlet air temperature, etc. These unexpected events threaten the reliability of disk by causing data corruption of disk. However, unfortunately, it is hard to predict when such thermal emergencies happen in real time. Hence, in this experiment, by generating the scenarios of thermal emergencies for the workloads in Table 2, a multi-speed technique of disk-drive is evaluated as a thermal management solution. The detailed scenarios for thermal emergencies are in Table 1.

We have simulated the temperature behavior of RPM transitions in a multi-speed disk. The operable RPM for this disk are 10K and 20K RPMs and the maximum time taken between different rotational speeds of disk is assumed to be 7 seconds (from the lower to the higher and vice versa as in the commercial multi-speed disk-drive of Hitachi [10]). We used four commercial I/O traces for the experiment, whose characteristics are given in Table 2, and we consider two kinds of multi-speed disks as follows.

- $DRPM_{simple}$: This is the same approach as Hitachi's multi-speed disk, where the lower RPM is just used for cooling the hot disk, rather than servicing the requests.
- $DRPM_{opt}$: This is the technique that was proposed in [6], where the disk-drive still performs I/O at the lower RPM.

4.1 *Time-based Policy*

This policy is based on a pre-defined period for cooling time before resuming to service the requests under thermal emergencies. Once it reaches thermal emergency, the RPM drops down and the drive waits for a pre-defined period (called a cooling unit time) before resuming I/O by ramping up the RPM to full-speed. Since $DRPM_{simple}$ is not available to service during the cooling and transition times, the performance is constrained by these two values. Most server workloads generally have many requests issued with short inter-arrival times and they should be processed as quickly as possible. In addition, 7 seconds of delay for each RPM transition is not negligible to the performance of a multi-speed disk. Figure 2 shows the performance degradation by $DRPM_{simple}$, compared to the disk without DTM under thermal emergencies. As is to be expected, most requests suffer from delays more than 200 ms. Even if we varied cooling unit times to compensate for the performance degradation, it is not effective for both of the workloads. $DRPM_{simple}$ might be desirable for a DTM solution, because such a straightforward policy does not require significant additional complexity to the disk-controller design and after reasonable delays, it could still overcome thermal emergencies by resuming the service below thermal envelope.

$DRPM_{opt}$ has been experimented to minimize the performance drawback caused by long delays of $DRPM_{simple}$.

Workload	*Initial* T_{amb} (°C)	*Increased* T_{amb} (°C)	Thermal Emergency Start (Sec)	Thermal Emergency End (Sec)	Simulated Time (Sec)
HPL Openmail [20]	29	42	500.000	2,500.000	3,606.972
OLTP Application [22]	29	33	5,000.000	30,000.000	43,712.246
Search-Engine [22]	29	33	2,000.000	12,000.000	15,395.561
TPC-C	29	33	2,000.000	10,000.000	15,851.521

Table 1. Description of thermal emergencies for real workloads. T_{amb} denotes the ambient air temperature and *Increased* T_{amb} is the value when thermal emergencies happen.

Workload	# Requests	# Disks	Per-Disk Capacity (GB)	RPM	Platter Diameter (in)	Platters (#)
HPL Openmail [20]	3,053,745	8	9.29	10,000	3.3	1
OLTP Application [22]	5,334,945	24	19.07	10,000	3.3	4
Search-Engine [22]	4,579,809	6	19.07	10,000	3.3	4
TPC-C	6,155,547	4	37.17	10,000	3.3	4

Table 2. Description of workloads and storage systems used.

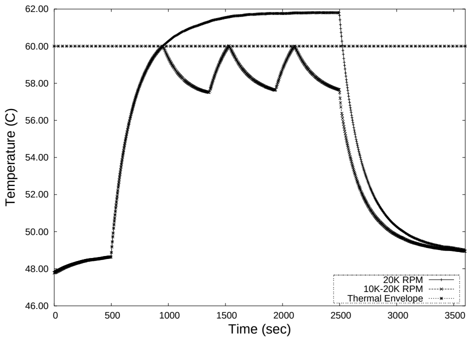

(a) HPL Openmail

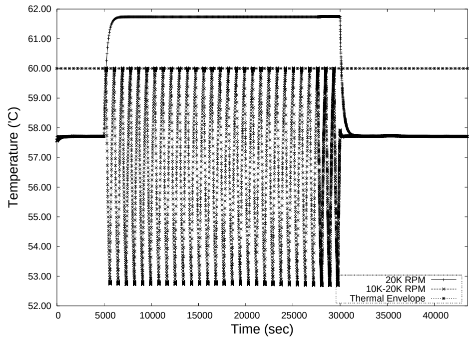

(b) OLTP

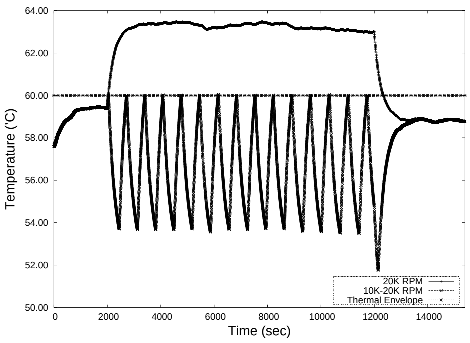

(c) Search-Engine

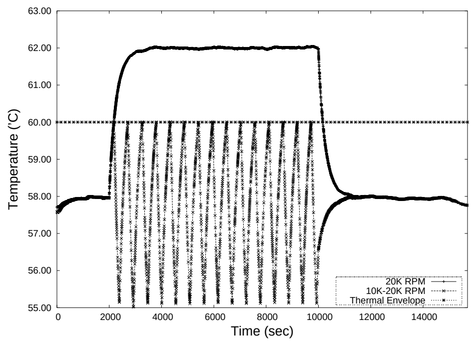

(d) TPC-C

Figure 3. Thermal profiles of the real workloads for $DRPM_{opt}$ under the scenarios of Table 1. They are all for the disk0 of disk arrays each of which is a 10K-20K multi-speed disk with 7 seconds of RPM transition time and 400 seconds of a cooling unit time.

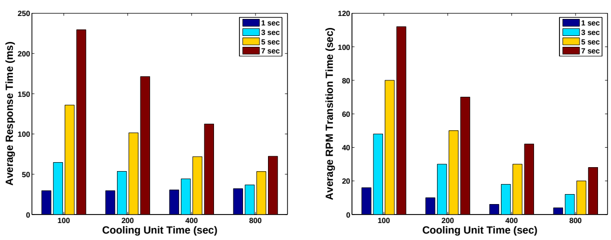

(a) HPL Openmail

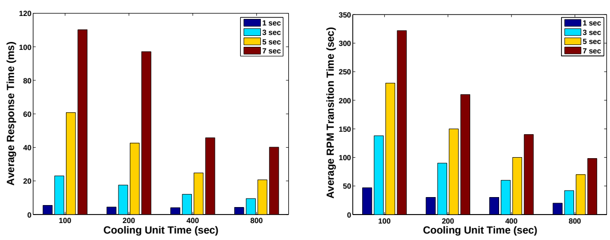

(b) TPC-C

Figure 4. The first column for each workload is to understand the correlation between cooling unit time and RPM transition time. Each bar denotes an average response time across the disks at disk array in the unit of millisecond. And the second column for each workload denotes the average total time taken for RPM transitions across the disks in the unit of second.

Figure 3 shows different thermal profiles for the workloads under thermal emergency situations described in Table 1. Since each disk-drive of the disk arrays of TPC-C, OLTP and Search-Engine has the same disk dimen-

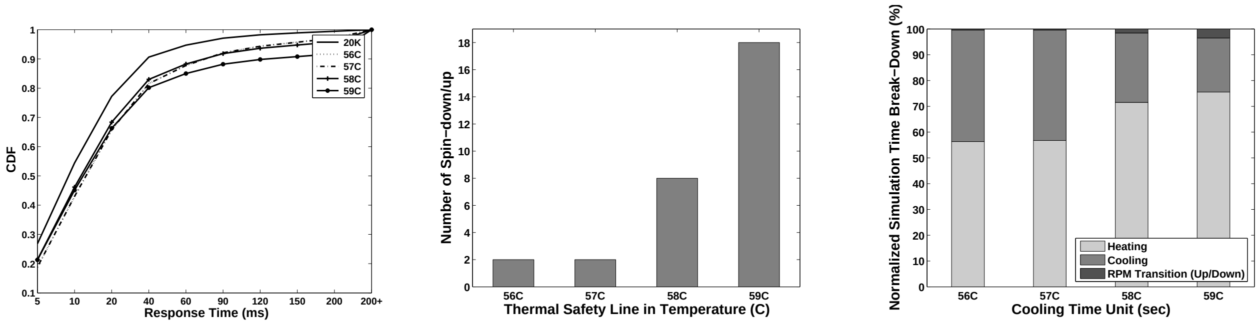

(a) HPL Openmail

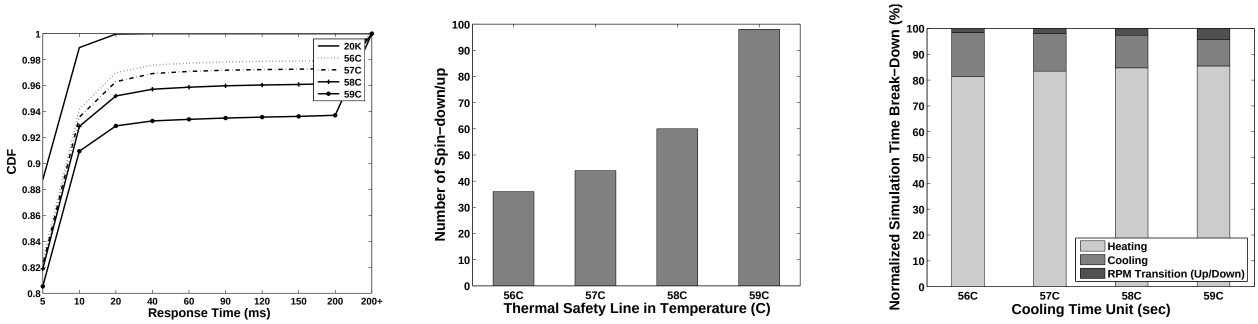

(b) TPC-C

Figure 5. Experimental results of $DRPM_{opt}$ using watermark-based policy for HPL Openmail and TPC-C

sion/characteristics and they have similar temperature profiles, we focus on the results for HPL Openmail and TPC-C. The upper curve for each graph is when DTM is not applied while operating at the maximum speed while the lower curve (going up and down) is the result from $DRPM_{opt}$ with 400 seconds of cooling unit time. As shown from Figure 3, operating at 20K RPM exceeds thermal envelope under thermal emergencies unless DTM is applied. $DRPM_{opt}$ avoids emergencies by dynamically modulating the rotational speed between high and low RPMs at need. However, many RPM transitions (for example, as shown from many transitions for OLTP in Figure 3) increase the overheads due to non-serviceable RPM transition time, even if $DRPM_{opt}$ could be available to service the requests during the cooling time.

The time taken for RPM transitions greatly affects the performance and thus, to study the impact of non-negligible RPM transition to the performance, we have experimented $DRPM_{opt}$ for different RPM transitions ranging from 1 second to 7 seconds in steps of 2 seconds. From Figure 4, we see that a small RPM transition time shows better performance for a given constant cooling unit time. In addition, high cooling times can hide the overhead of RPM transitions by reducing the number of RPM transitions and sparing more time for I/O disk operations. In the above performance equation of a multi-speed disk, T_{Cool} works positively in the performance of $DRPM_{opt}$, however, it still has an upper bound in performance improvement. This is because more cooling implies that more requests should be serviced at the lowest speed of RPMs in a multi-speed disk.

4.2 *Watermark-based Policy for $DRPM_{opt}$*

Watermark-based Policy uses two thresholds, T_{high} and T_{low}. In this policy, the thermal sensor of disk-drive periodically checks the temperature and if it detects that the temperature is close to thermal emergency (T_{high}), which is the temperature at which DTM is invoked, thermal management is applied to cool down the disk until the temperature gets down to the pre-determined threshold (T_{low}). After this point, the disk controller comes to know that the emergency has been resolved.

Figure 5 shows the experimental results of $DRPM_{opt}$ where T_{high} is 60°C (which is set to be the same as the thermal envelope in this experiment) and T_{low} is obtained by subtracting a few degree Celsius from T_{high}. The graphs in the first column of Figure 5 shows the CDF of the average response time across disk-drives. The solid curve in each graph represents the performance of the baseline system without any DTM and others are for DTM with different lower thresholds (T_{low}). In this experiment, the RPM transition time (up and down) is assumed to 7 seconds. The lower value of T_{low} helps the performance of a multi-speed disk. This is because the lower value of T_{low} allows more relaxation in throttling and reduces the number of RPM transitions. As shown from the graphs in the second column of Figure 5, the lower threshold of T_{low} has fewer number of RPM transitions. The graphs in the last column of Figure 5 show the time break-down for total simulated time composed of heating/cooling times and RPM transition time for different lower thresholds. The larger fraction of heating time implies better performance, because it implies that more requests could be serviced at the maximum speed. However, it is noted that it is not absolutely better than the small portion of heating time, because RPM transition time (in the order of seconds) offset this benefit.

5 CONCLUSIONS

This paper has presented graceful operation of multi-speed disk to handle thermal emergencies in large disk arrays. We studied several DTM policies to the disks executing real workloads and observed that DRPM technique is one of the best solutions to avoid thermal emergencies. $DRPM_{simple}$ technique overcomes thermal emergencies by dynamically modulating the rotational speed of disks and providing pre-defined delays. But such delays cause poor performance (such as response time), compared to a normal disk drive without DTM. However, $DRPM_{opt}$ technique further improves the performance by continuously servicing the requests at the lower speed. *Time* and *Watermark-based Policies* have been evaluated for thermal management and they showed that the time taken for RPM transition in a multi-speed disk is a crucial part in the performance for thermal management.

Acknowledgements

This research has been funded in part by NSF grants 0429500, 0325056, 0130143, 0509234 and 0103583, and an IBM Faculty Award.

6 REFERENCES

[1] D. Brooks and M. Martonosi. Dynamic Thermal Management for High-Performance Microprocessors. In *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, pages 171–182, January 2001.

[2] E.V. Carrera, E. Pinheiro, and R. Bianchini. Conserving Disk Energy in Network Servers. In *Proceedings of the International Conference on Supercomputing (ICS)*, June 2003.

[3] N.S. Clauss. A Computational Model of the Thermal Expansion Within a Fixed Disk Drive Storage System. Master's thesis, University of California, Berkeley, 1988.

[4] S. Gurumurthi. The Need for Temperature-Aware Storage Systems. In *Proceedings of the Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, pages 387–394, May 2006.

[5] S. Gurumurthi, Y. Kim, and A. Sivasubramaniam. Thermal Simulation of Storage Systems using STEAM. *IEEE Micro Special Issue on Computer Architecture Simulation and Modeling*, 2006.

[6] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke. DRPM: Dynamic Speed Control for Power Management in Server Class Disks. In *Proceedings of the International Symposium on Computer Architecture (ISCA)*, pages 169–179, June 2003.

[7] S. Gurumurthi, A. Sivasubramaniam, and V. Natarajan. Disk Drive Roadmap from the Thermal Perspective: A Case for Dynamic Thermal Management. In *Proceedings of the International Symposium on Computer Architecture (ISCA)*, pages 38–49, June 2005.

[8] T. Heath, A. P. Centeno, P. George, Y. Jaluria, and R. Bianchini. Mercury and Freon: Temperature Emulation and Management in Server Systems. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, October 2006.

[9] G. Herbst. IBM's Drive Temperature Indicator Processor (Drive-TIP) Helps Ensure High Drive Reliability. In *IBM Whitepaper*, October 1997.

[10] Hitachi Power and Acoustic Management - Quietly Cool. In *Hitachi Whitepaper*, March 2004. http://www.hitachigst.com/tech/techlib.nsf/productfamilies/White_Papers.

[11] Y. Kim, S. Gurumurthi, and A. Sivasubramaniam. Understanding the Performance-Temperature Interactions in Disk I/O of Server Workloads. In *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA)*, Febuary 2006.

[12] M. K. Patterson, X. Wei, and Y. Joshi. Use of computational fluid dynamics in the design and optimization of microchannel heat exchangers for microelectronics cooling. In *Proceedings of the ASME Summer Heat Transfer Conference*, 2005.

[13] Seagate Cheetah 15K.3 SCSI Disc Drive: ST3734553LW/LC Product Manual, Volume 1. http://www.seagate.com/support/disc/manuals/scsi/100148123b.pdf.

[14] Seagate Workload Management for Business-Critical Storage. http://www.seagate.com/docs/pdf/whitepaper/TP555_BarracudaES_Jun06.pdf.

[15] L. Shang, L-S. Peh, A. Kumar, and N.K. Jha. Thermal Modeling, Characterization and Management of On-Chip Networks. In *Proceedings of the International Symposium on Microarchitecture (MICRO)*, pages 67–78, December 2004.

[16] R.K. Sharma, C.E. Bash, C.D. Patel, R.J. Friedrich, and J.S. Chase. Balance of Power: Dynamic Thermal Management for Internet Data Centers. *IEEE Internet Computing*, 9(1):42–49, January 2005.

[17] K. Skadron, M.R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan. Temperature-Aware Microarchitecture. In *Proceedings of the International Symposium on Computer Architecture (ISCA)*, pages 1–13, June 2003.

[18] M. Sri-Jayantha. Trends in Mobile Storage Design. In *Proceedings of the International Symposium on Low Power Electronics*, pages 54–57, October 1995.

[19] J. Srinivasan and S.V. Adve. Predictive Dynamic Thermal Management for Multimedia Applications. In *Proceedings of the International Conference on Supercomputing (ICS)*, pages 109–120, June 2003.

[20] The Openmail Trace. http://tesla.hpl.hp.com/private software/.

[21] Ultrastar 146Z10 hard disk drives specifi cations. http://www.hitachigst.com/hdd/support/146z10/146z10.htm.

[22] UMass Trace Repository. http://traces.cs.umass.edu.

[23] A. Weissel and F. Bellosa. Dynamic Thermal Management for Distributed Systems. In *Proceedings of the First Workshop on Temperature-Aware Computer Systems (TACS)* , June 2004.