USING STEAM FOR THERMAL SIMULATION OF STORAGE SYSTEMS

AS DEMANDS FOR HIGHER DATA TRANSFER RATES MOUNT, TEMPERATURE IS JOINING BANDWIDTH AND LATENCY AS A KEY CONSIDERATION IN STORAGE SYSTEM DESIGN. THE STEAM SIMULATOR LETS USERS ABSTRACT THE DETAILS OF LOW-LEVEL RECORDING PHYSICS AND HEAT TRANSFER PHENOMENA AND FACILITATES EXPLORATION OF A LARGE DESIGN SPACE OF STORAGE CONFIGURATIONS UNDER REALISTIC WORKLOADS.

Sudhanva Gurumurthi
University of Virginia

Youngjae Kim
Anand Sivasubramaniam
Pennsylvania State University

Computing today is highly data-driven. Many applications, especially in the business and scientific domains, must move massive amounts of data between the computing nodes and the storage system at high transfer speeds. In scientific applications such as high-energy physics, data from instruments must move continuously to the storage system at transfer speeds of terabytes or even petabytes per second.[1] A significant determinant of these applications' overall performance is the storage system—particularly the data rates of the individual disk drives.[1,2]

Designing disks involves trade-offs between capacity, speed, and power. One way to increase disk drive capacity is to use larger platters or several of them. The number of platters and their size both affect the heat generated by viscous dissipation inside the disk drive: the number by a linear factor, and the size by the 4.6th power. Improvements in linear density, expressed in bits per inch (bpi), or increases in revolutions per minute (rpm) can increase the disk drive data rate. Increases in rpm increase the heat generated by nearly a cubic factor.

One disk drive design requirement is to ensure that the disk's operating temperature is always below a particular threshold—the *thermal envelope*. Given a particular maximum external ambient temperature, the design must ensure that under worst-case operating conditions, the disk's temperature does not exceed the thermal envelope. Designers can achieve performance improvements within this thermally constrained design space through a combination of improvements in the magnetic-recording technology and structural changes to the disk drive. The structural modifications involve shrinking the platters, which reduces power by the 4.6th power, and exploiting this slack to ramp up the rpm.

The main objective in designing disks to operate within the thermal envelope is reliability. High temperatures cause reliability problems ranging from data corruption to complete device failure. For example, a 15°C rise in ambient temperature can nearly double a disk

0272-1732/06/$20.00 © 2006 IEEE Published by the IEEE Computer Society

drive's failure rate.[3] Moreover, in a rack-based server system, disks can be close to processors and memory cards, which have similar thermal constraints. Excess heat from the disks can preheat the air around the other components, and vice versa. With the high costs of cooling modern electronic systems, it is important that disks not increase the burden.

For nearly two decades, the thermally constrained design methodology has scaled disk drive data rates along a 40 percent annual growth curve, but we expect the future to be different. Disk drives have achieved a good portion of performance through growth in density, supported by periodic increases in drive rpm. However, various factors will slow the density growth rate, such as difficulty in lowering the head's height above the platter, limitations on shrinking a magnetic bit cell, and limitations on closely packing tracks on the surface. This means that the only way we can stay on the performance curve is by increasing rpm to a greater degree than we have in the past. But this will make operating the disk below the thermal envelope very challenging.[2] For instance, in 2009, a single-platter 2.6-inch disk—a common platter size in SCSI disk drives today—will operate at around 85°C if its rpm is set high enough to meet the 40 percent target that year. This is nearly as hot as a high-performance microprocessor. Ironically, the scientific community strongly demands that disk drive data rates scale up even faster than 40 percent per year.[1]

As temperature emerges as an important constraint in designing and deploying storage systems, we need good simulation tools to provide insight into their thermal behavior. In this article, we present such a tool—the Storage Thermal Exploration and Modeling (STEAM) simulator. We built STEAM to address two key goals. First, we wanted a simulator that computer architects and systems researchers can use to abstract the details of low-level recording physics and heat transfer phenomena. Previous researchers have performed thermal analysis of storage systems using computational fluid dynamics models,[4] but these are not easy for computer systems designers to use. Second, we wanted the simulator to be flexible enough to facilitate exploration of a large design space of storage configurations and also to be capable of running realistic workloads.

Simulator design

STEAM consists of two components: a performance model and a thermal model. The performance model simulates all storage system activities that potentially can affect workload performance, such as disk and interconnect latencies, RAID (redundant arrays of inexpensive disks) organizations, and so forth. The thermal model captures heat transfer phenomena in the storage system and all physical parameters that potentially affect temperature. For the performance model, STEAM uses the DiskSim simulator, which models the performance aspects of disk drives, controllers, caches, and interconnects in a fairly detailed manner. It is an event-driven simulator, with simulated time updated on discrete events such as request arrival and seek completion. Researchers have used DiskSim extensively in several storage system studies and thoroughly validated it with several disk models.

Thermal model

The thermal model is an extension of Eibeck and Cohen's work at the University of California, Berkeley.[4] This model evaluates the drive's temperature distribution by calculating the amount of heat generated by components such as the spindle motor (SPM)—which rotates platters—and the voice-coil motor (VCM)—which moves disk arms; heat conduction along solid components; and convection of heat to air.

The model assumes that the drive is completely enclosed and that the only interaction with external air is heat conduction through the base and the cover and convection to the outside air, assumed to maintain a constant temperature via a cooling system such as fans.

The model divides the hard disk into four components: internal drive air, SPM assembly consisting of motor hub and platters, base and cover, and VCM and disk arms.

Newton's law of cooling gives the heat transfer rate over time interval t, dQ/dt (in watts), through cross-sectional area A, as

$$dQ/dt = hA\Delta T$$

where h is the heat transfer coefficient and ΔT is the temperature difference between the two entities. For solids, which transfer heat via conduction, heat transfer coefficient h

depends on thermal conductivity *k* and the material's thickness and is given as *k*/*thickness*.

Heat exchange between solids and fluids takes place via convection, in which the heat transfer coefficient depends on whether the fluid flow is laminar or turbulent and also on the solid components' exact geometry. STEAM's thermal model uses empirical correlations to calculate the heat transfer coefficient of the disk drive's various solid components. These correlations provide equations for the heat transfer coefficient by approximating the complex geometry of the internal drive structures to simpler shapes. (Researchers have validated these correlations for real disk drives in previous work.[4]) The heat of the internal drive air is the sum of heat energy convected to it by each solid component and viscous dissipation (internal friction) in the air itself, minus heat lost through the cover to the outside. The following equation[5] gives the viscous dissipation:

$$\begin{aligned} \textit{viscous dissipation} \propto & (\textit{no. of platters}) \\ & \times (\textit{rpm})^{2.8} \\ & \times (\textit{platter diameter})^{4.6} \end{aligned}$$

The model uses the finite-difference method to solve the heat equations for the various components. At each time step, it calculates the temperatures of all the components and the air; it iteratively revises this calculation at each subsequent time step until the temperatures converge to a steady state. The model's accuracy depends on the size of the time steps: Using coarse-grained time steps lets the model execute faster, at the cost of accuracy. Using fine-grained time steps slows execution but improves accuracy.

The thermal model requires several input parameters. The first set of parameters describes disk geometry such as platter dimensions, drive base and cover dimensions including those of the disk drive's inner cavities, disk arms' length, and so forth. We have created parameterized geometry models based on physical measurements of different types of disk drives. These models let users specify high-level disk drive parameters such as the number of platters and rpm, which they use to automatically generate the geometry for the simulation. We determined the properties of materials used in drives by talking to engineers in disk drive companies.

Even as a stand-alone tool, STEAM's thermal model is useful for studying the effects of recording technology and structural design on the disk drive's temperature.[2] Integrated with the performance model, it lets us study entire storage systems.

Integrating performance and thermal models

The performance model and the thermal model each use some parameters and events that don't affect the other model. For instance, the performance model doesn't need to know the external air temperature or the material composition of the disk platters. Similarly, the thermal model is relatively unaffected by disk cache size (although that parameter can alter the time between requests that must access the platters). However, several shared parameters (or state information) need to flow from one model to the other. For example, the thermal model must know when seeks start and end, since that directly affects VCM power. First, we must make sure that the relevant state information flows from one model to the other. Second, we must ensure that simulation time is reconciled between the event-driven performance model and the time-step-based thermal model for each disk.

Integration of the two models relies on the observation that only two governing factors from the performance model affect the thermal model: seek activities (particularly VCM on and off events) and changes in rpm. At these points, the performance model invokes the thermal model to iteratively (that is, in time steps) compute heat flow until the thermal model's simulated time reaches that of the next such point in the performance model. In other words, we normally run the performance model for the sequence of incoming I/O requests. Whenever this model incurs a VCM switch from its prior state (on from off, or vice versa), it invokes the thermal model with the appropriate VCM state information so that the thermal model can catch up in time to the performance model, at which point control flows back to the performance model. In the case of a multispeed dynamic rpm (DRPM) disk,[6] this invocation also occurs at rpm change events.

Modeling physical behavior of disk seeks

A key event that the simulator must model accurately is a disk seek, which has significant

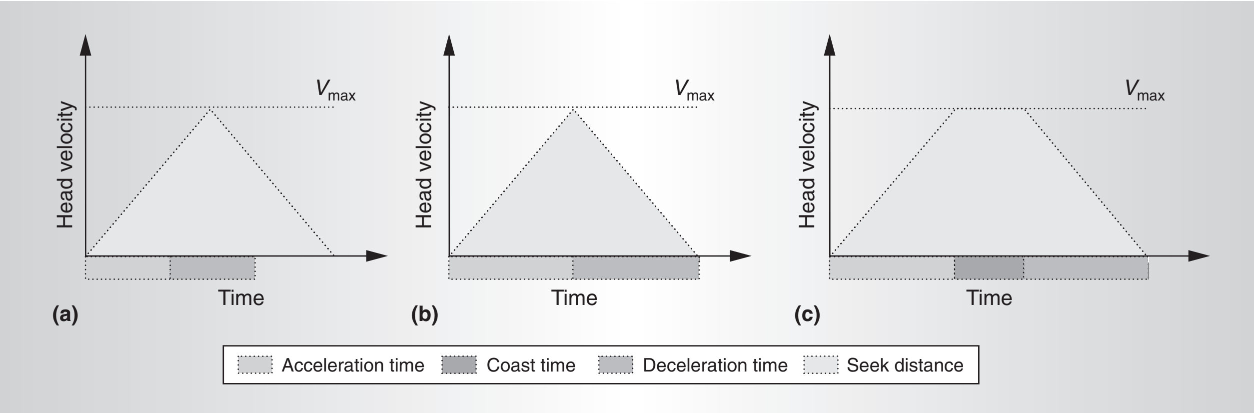

Figure 1. Physical seek operation possibilities: seek distance less than average (a); average seek distance (b); and seek distance greater than average (c).

impact on disk temperature. Disk seeks are induced by I/O requests, which are performance events. The DiskSim performance model already accounts for time taken by a seek operation. The thermal model must account for the mechanical work involved in physically doing the seek. (Another publication gives complete details of how we model seek operations.[7])

The time a seek takes depends on two factors: the VCM assembly's inertial power and the radial length along the platter that the arm must move. Physically, a seek consists of

- an acceleration phase, in which the VCM is powered,
- a coast phase of constant velocity, in which the VCM is off,
- a deceleration phase to stop the arms near the desired track by turning the VCM on again but reversing the current to generate a braking effect, and
- settling of the head.

Figure 1 depicts these phases.

We capture the physical behavior of seeks using a bang-bang triangular model, in which acceleration and deceleration times are equal. The head's maximum velocity (V_{max}), dictated by the VCM assembly's characteristics and the underlying servo system's bandwidth, is assumed to be 120 inches per second, a value that reflects many modern disk drive implementations. For a large number of random seeks, we refer to the distance across one third of the platter's data zone as the average seek distance (D_{avg}). We set up the seek time model such that an average seek's coast time is zero because that yields the lowest seek time on average. We compute the power of the VCM itself using the model presented by Sri-Jayantha.[8]

Validation

We have validated STEAM throughout its development cycle. The Berkeley model's original authors validated it using disk drives that are now more than 15 years old. To make sure that our extensions are valid and applicable to modern disk drives, we validated our model using a Hitachi Deskstar 7K500 500-Gbyte ATA disk drive (http://www.hitachigst.com/hdd/support/7k500/7k500_ov.htm). Figure 2 shows a photograph of the laboratory setup for this validation experiment.

To ensure that we could precisely control external ambient temperature, we placed the disk in an isothermal oven. Since STEAM assumes that external ambient temperature is constant, we wanted to create the same scenario for the validation experiment. We drew the IDE cables out through the oven door and linked them to a workstation, where we collected the experimental data. We equipped the disk with self-monitoring, analysis, and reporting technology (SMART) sensors that report diagnostic information about the drive, including the enclosure temperature. We set the oven temperature at 68°C and let the disk run on idle (platters rotating but no I/O activity) for six hours to stabilize to its steady-state temperature, which was 81°C.

We then replicated the experiment parameters in STEAM and simulated it for the same duration of (simulated) time. STEAM reported the steady-state temperature as 83.4°C, a difference of 3 percent. A key reason for this small inaccuracy is that the thermal sensor measures temperature only at the disk cover, whereas STEAM reports internal air temperature.

To validate the seek model, we calculated the acceleration that STEAM computes, under all the stated assumptions for a Fujitsu AL-7LX disk drive, and compared this calculated acceleration with the disk drive's measured mechanical seek characteristics.[9] Using the drive characteristics, we found the D_{avg} for this disk was 0.22 inch. The reported acceleration value satisfying the seek time requirement is 220 G (2,150 m/s^2); using the same D_{avg}, our model calculates the acceleration as 253.5 G (2,488.1 m/s^2), which is within 15 percent of the reported value.

Figure 2. Experimental setup for thermal model validation. The box to the right of the workstation is the oven, with integrated drive electronics (IDE) cables coming out the door.

Temperature-aware storage research

We have used STEAM for two pieces of research: studying the effects of changes in recording technology and structural design on temperature phenomena in a single-disk drive, and analyzing the thermal behavior of applications that run on a multidisk storage system.

Temperature phenomena in a single-disk drive

For many years, areal density—a product of bpi and track density, expressed in tracks per inch (tpi)—grew at a rate of nearly 100 percent per year. This brisk growth in density allowed disks' internal data rate (IDR) to grow by 40 percent each year. These growth rates profoundly affected the price of storage and paved the way for us to be able to store several gigabytes of data in a disk drive. However, this trend is slowing down, and we shall briefly explain why.

The bpi growth rate is expected to slow for several reasons. For one thing, it is difficult to lower the head's flying height; the gap between it and the platter surface is only a few nanometers. Second, a higher bpi requires a more coercive recording medium—meaning a medium that requires more energy to change the magnetic state (that is, to write to disk). Since writes take place via magnetic induction from the write head, increasing this energy requires more powerful electric fields; these are difficult to create with currently known head materials. Finally, to achieve a higher bpi, the standard scaling approach has been to shrink the size of the magnetic grains that compose a bit cell. However, if the grain size becomes smaller than a particular threshold, called the superparamagnetic limit, external thermal energy can overwhelm the grain's stored signal energy, essentially causing the bit to flip. We have recently hit this point in the curve.

Higher densities also cause problems with tpi. Narrower tracks are more susceptible to media noise and intertrack interference. Although alternative approaches exist for working around some of these problems, industry projections for long-term growth in areal density indicate a slowdown to roughly 50 percent per year.[10] Moreover, compensating for reduced signal-to-noise ratios at high areal densities requires a significant amount of error-correcting code (ECC). ECC bits are stored alongside regular data bits, so they eat away the disk drive's effective capacity and data rate. Therefore, maintaining the status quo in performance growth requires a more aggressive scaling of drive rpm to compensate for the slowdown in recording technology.

Given that the underlying magnetic technology is undergoing all these changes, we were interested in seeing what, if any, impact such changes will have on the established thermally constrained design methodology. To ascertain this, we developed a detailed predictive technology model that takes into account the salient technological and organi-

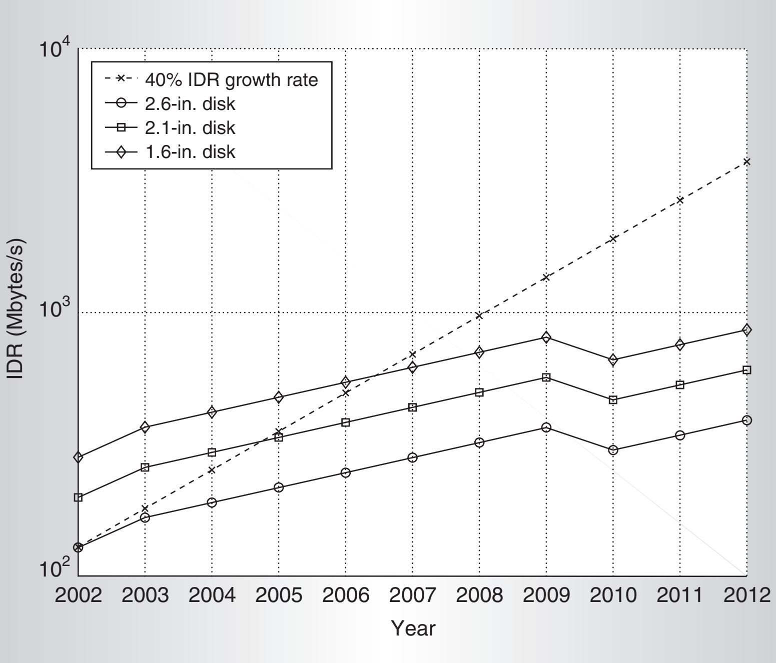

Figure 3. One-platter disk drive roadmap: Each solid curve gives maximum attainable IDR for that platter size within the thermal envelope. Dotted line indicates the 40 percent target IDR growth rate. The *y*-axis is in log scale.

zational parameters to compute bpi and tpi for different calendar years. Once we had these fundamental parameters, we generated a roadmap for the years 2002 to 2012. Our basic premise was to sustain the expected IDR growth rate of at least 40 percent a year over the 11-year period. Since disks must be designed to adhere to their thermal envelopes, we needed to know the temperature resulting from operating each disk configuration. We used STEAM's stand-alone thermal model for this purpose.

Figure 3 shows results for a one-platter disk drive. The graph shows data rates for disk drives with three different platter sizes housed in a 3.5-inch form-factor drive enclosure. The straight dotted line corresponds to the 40 percent IDR growth rate.

For a brief analysis of this roadmap, consider the 2.6-inch platter size. As we move along the 40 percent growth rate curve, the IDR requirements increase nearly 29 times from 2002 to 2012. A portion of the required increase comes from growth in linear density alone. Any demands beyond that must come from an increase in rpm. To determine the roadmap points where significant rpm changes are required, it is useful to subdivide the timeline into three regions. The first covers the years before 2004, where the bpi and tpi growth rates are 30 percent and 50 percent, respectively.[11] The second region covers the years from 2004 to 2009, which correspond to subterabit areal densities. In this region, bpi and tpi growth rates slow, but ECC requirements are still moderate. To compensate for the slowdown in recording density, we would need to scale up rpm more aggressively to meet data rate requirements. Third, the years from 2010 to 2012 correspond to the terabit areal density region; this will mean steep growth in ECC requirements, so that we will have to scale up rpm even more aggressively.

Analyzing these design points in terms of thermal behavior, we find that in the second region of the roadmap, ramping up the rpm increases viscous dissipation from 2 W in 2004 to more than 35.55 W in 2009, when the disk runs as hot as 85.04°C (almost as hot as a high-performance microprocessor). This is a significant rise in temperature—well above the thermal envelope. Overall, tracking the 40 percent IDR curve will not be possible, from the thermal perspective, from 2007 onward. Starting in 2010, viscous dissipation will take another leap upward, reaching a value of 499.73 W in 2012, raising internal air temperature to a scorching 602.98°C.

This study showed that it will not be possible to neglect temperature issues when deploying disk drives and building storage systems. Temperature must be treated as a first-class constraint along with traditional performance metrics such as bandwidth and latency. A possible approach to tackling this performance problem is to relax design constraints, requiring operation below the thermal envelope only in the average case, in which the I/O load is moderate. To prevent thermal emergencies, we would need to continuously monitor the disk drive's temperature and control its activities dynamically. This form of dynamic thermal management (DTM) has already been appearing in microprocessors,[12] and similar techniques can be applied in disk drives as well. Another article gives detailed information about our roadmap study and a preliminary evaluation of DTM techniques.[2]

Thermal behavior of entire storage systems

One of STEAM's most powerful features is that it enables designers to compose entire storage systems, consisting of disks, controllers, and associated interconnect, and study their temperature behavior while running a workload. Our roadmap study indicated that disk drive performance will experience a significant slowdown if we keep designing disk drives conservatively—to operate below the thermal envelope in the worst case. However, it is possible that real workloads rarely perform I/O in a sustained manner that embodies worst-case characteristics. If this is so, we might be able to safely relax the thermally constrained design rules without having to trigger DTM countermeasures. To learn whether this is indeed the case, we used STEAM's full-fledged integrated-performance thermal simulator.

We conducted this study using five commercial I/O traces. We configured STEAM to model the storage system from which each I/O trace was collected as closely as possible. Figures 4 and 5 show temperature profiles of two workloads used in the study. The graphs show results for disks used in the original storage system (which run at 10,000 rpm for both workloads) and two higher rpms. The physical organization of the higher-rpm disks are identical in all respects to those in the original storage system except for the rpm. We give complete details of this study elsewhere.[7]

Figures 4 and 5 give temperature profiles for the two workloads at three time granularities. Figures 4a and 5a show profiles across the entire workload simulation time; Figures 4b and 5b show a five-second window starting from the 50th minute of simulation; and Figures 4c and 5c zoom into a one-second window at the beginning of the five-second period. For clarity, Figures 4a and 5a show the temperature response curves for just one disk, whereas parts b and c of both figures show trends for all the disks in their respective storage systems.

For both workloads, the disk drives could accommodate a 5,000-rpm increase from the baseline within the thermal envelope without increasing the cooling requirements. This was due to the nature of the seek times, which were relatively small. As a result, acceleration and deceleration times were small, generat-

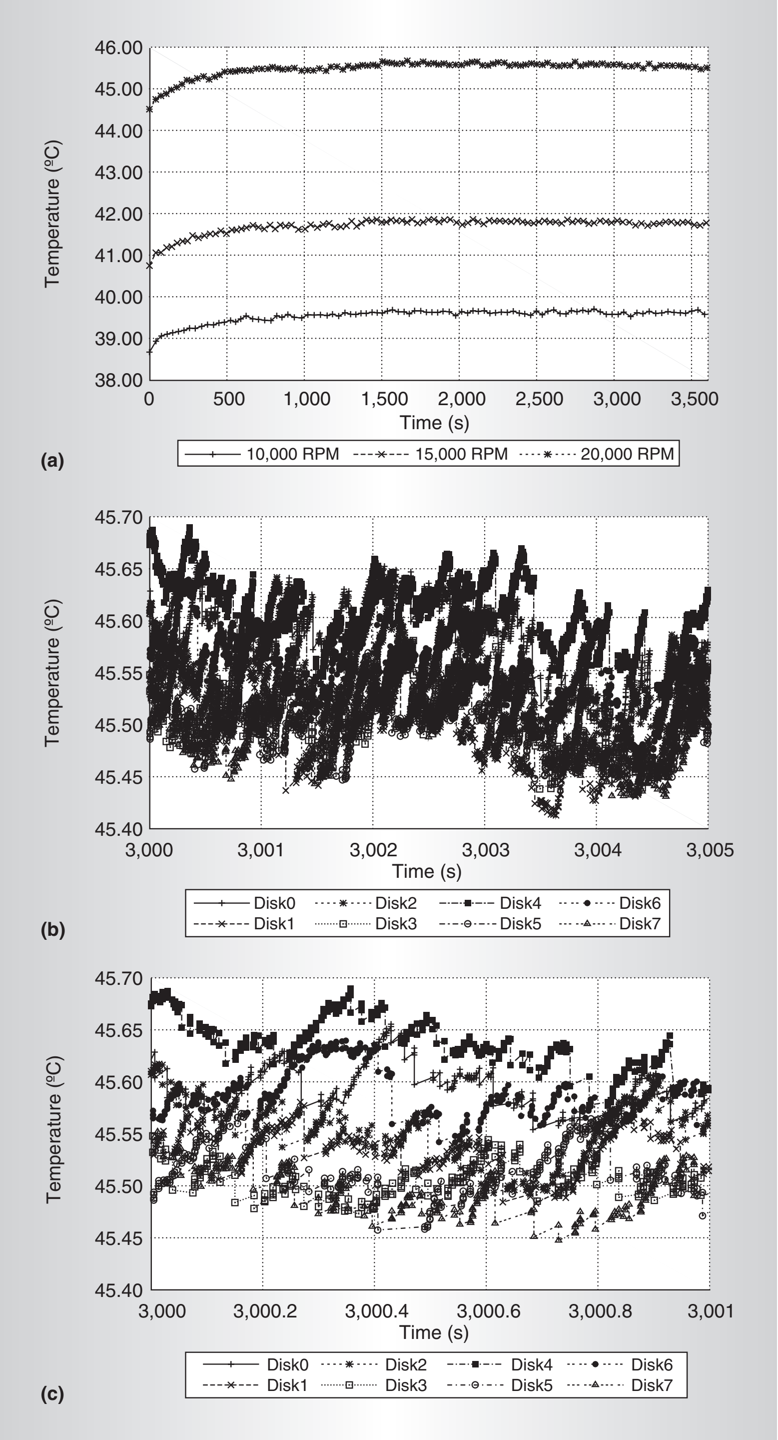

Figure 4. Thermal profile of the Openmail workload at three different rpm and three time ranges: the entire workload simulation time for 10,000-, 15,000-, and 20,000-rpm disks (a); a five-second window starting from the 50th minute of simulation at 20,000 rpm (b); and a one-second window at the beginning of the five-second period at 20,000 rpm (c). The thermal envelope is 45.22°C. To make temperature variations as detailed as possible, each graph has a different *y*-axis scale.

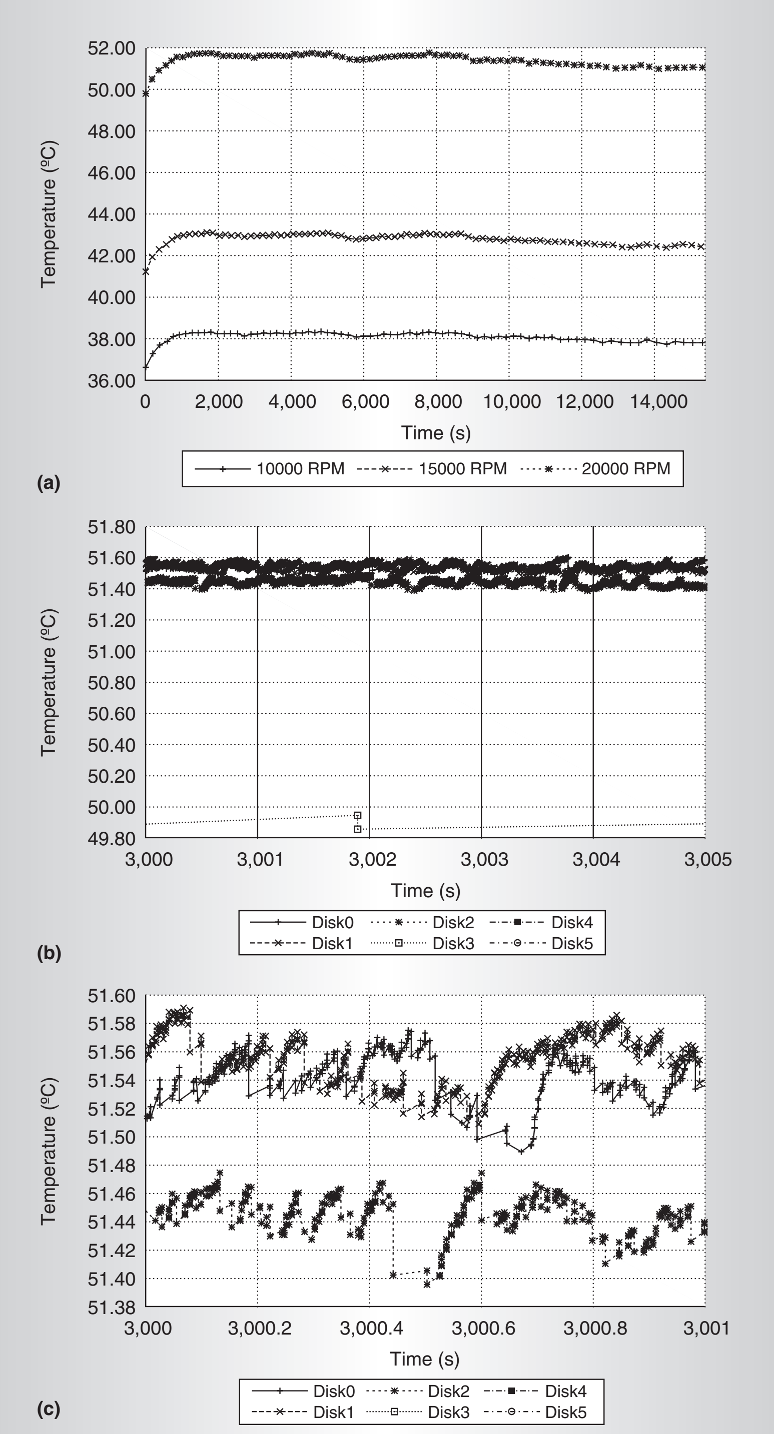

Figure 5. Thermal profile of the Search-Engine workload: the entire workload simulation time for 10,000-, 15,000-, and 20,000-rpm disks (a); a five-second window starting from the 50th minute of simulation at 20,000 rpm (b); and a one-second window at the beginning of the five-second period at 20,000 rpm (c). The thermal envelope is 45.22°C. Each graph has a different *y*-axis scale.

ing very little heat. However, when we increased speed by another 5,000 rpm, to 20,000 rpm, disk temperatures exceeded the thermal envelope. The disk used in Search-Engine (Figure 5) shows a more significant excursion above the thermal envelope at 20,000 rpm than the disk in Openmail (Figure 4). This is because Search-Engine's storage system uses four-platter disks, which dissipate more heat at the higher rpm than the one-platter disks in Openmail. Because the cooling budget was established to satisfy the thermal envelope at 10,000 rpm, at 20,000 rpm the cooling system does not sufficiently extract excess heat, causing the drive temperature to rise.

Our study shows that disk drive designers do have some flexibility to relax the worst-case thermal design requirements and still avoid thermal emergencies. However, we have only 5,000 rpm of slack before we hit the "thermal wall," at which time we would be forced to invoke DTM mechanisms more frequently to keep the temperature in check. Therefore, we need to devise DTM techniques that are both lightweight in terms of performance impact and effective at keeping the operating temperature below the thermal envelope.

The STEAM simulator is a key piece of our Thermal-Aware Storage Systems Project. It is an excellent tool for collaborative research in computer architecture, computer systems, thermal engineering, and magnetic-recording physics. Because many challenges in computer architecture today are linked to the physical nature of devices, having tools and formalisms that enable such cross-disciplinary research is imperative. We have used STEAM to study the thermal problems facing disk drive designers and explain them to the computer architecture community. We have also used the simulator to understand relationships between system-level I/O behavior and low-level physical behavior. Such studies are important because they facilitate the establishment of metrics for designing and evaluating storage architectures. MICRO

Acknowledgments

We thank Nathanael Paul of the University of Virginia for carrying out the thermal model validation experiment.

References

1. Office of Science Data-Management Challenge, http://www.sc.doe.gov/ascr/Final-report-v26.pdf.
2. S. Gurumurthi, A. Sivasubramaniam, and V. Natarajan, "Disk Drive Roadmap from the Thermal Perspective: A Case for Dynamic Thermal Management," *Proc. 32nd Int'l Symp. Computer Architecture* (ISCA 05), IEEE Press, 2005, pp. 38-49.
3. D. Anderson, J. Dykes, and E. Riedel, "More than an Interface—SCSI vs. ATA," *Proc. 2nd Ann. Conf. File and Storage Tech.* (FAST 03), Usenix, 2003, pp. 245-256.
4. P.A. Eibeck and D.J. Cohen, "Modeling Thermal Characteristics of a Fixed Disk Drive," *IEEE Trans. Components, Hybrids, and Manufacturing Tech.*, vol. 11, no. 4, Dec. 1988, pp. 566-570.
5. S.-O. Cho, S.-Y. Lee, and Y.-C. Rhim, "Aerodynamically Induced Power Loss in Hard Disk Drives," *Microsystem Technologies*, vol. 11 nos. 8-10, Aug. 2005, pp. 741-746.
6. S. Gurumurthi et al., DRPM: Dynamic Speed Control for Power Management in Server Class Disks," *Proc. 30th Int'l Symp. Computer Architecture* (ISCA 03), IEEE Press, 2003, pp. 169-179.
7. Y. Kim, S. Gurumurthi, and A. Sivasubramaniam, "Understanding the Performance-Temperature Interactions in Disk I/O of Server Workloads," *Proc. 12th Int'l Symp. High-Performance Computer Architecture* (HPCA 06), IEEE Press, 2006, pp. 179-189.
8. M. Sri-Jayantha, "Trends in Mobile Storage Design," *Proc. Int'l Symp. Low-Power Electronics*, IEEE Press, 1995, pp. 54-57.
9. K. Aruga, "3.5-Inch High-Performance Disk Drives for Enterprise Applications: AL-7 Series," *Fujitsu Science and Tech. J.*, vol. 37, no. 2, Dec. 2001, pp. 126-139.
10. E. Riedel, "Wheels Keep Spinning—Disc Drive Research," *Proc. HEC-IWG File Systems and I/O R&D Workshop*, Nat'l Coordination Office for Networking and Information Tech. Research and Development (NITRD), 2005; http://www.nitrd.gov/subcommittee/hec/workshop/20050816_storage/industry/riedel.pdf
11. Hitachi Global Storage Technologies—HDD Technology Overview Charts; http://www.hitachigst.com/hdd/technolo/overview/storagetechchart.html.
12. S. Gunther et al., "Managing the Impact of Increasing Microprocessor Power Consumption," *Intel Technology J.*, vol. 5, no. 1, Feb. 2001; http://www.intel.com/technology/itj/archive/2001.htm.

Sudhanva Gurumurthi is an assistant professor in the Computer Science Department of the University of Virginia. His research interests include computer architecture and storage systems. Gurumurthi has a PhD from Pennsylvania State University and a BE from the College of Engineering, Guindy, Anna University, India, both in computer science and engineering. He is a member of the ACM.

Youngjae Kim is a graduate student in the Computer Science and Engineering Department of Pennsylvania State University. His research interests include operating systems and power and thermal management of storage systems. Kim has a BS from Sogang University and an MS from Korea Advanced Institute of Science and Technology.

Anand Sivasubramaniam is a professor in the Computer Science and Engineering Department of Pennsylvania State University. His research interests include computer architecture, operating systems, parallel and distributed computing, and performance evaluation. Sivasubramaniam has a BTech from the Indian Institute of Technology, Madras, and an MS and a PhD from Georgia Institute of Technology, all in computer science. He is a member of the IEEE and the ACM.

Direct questions and comments about this article to Sudhanva Gurumurthi, Dept. of Computer Science, University of Virginia, Charlottesville, VA 22904; gurumurthi@cs.virginia.edu.