

Thermo-Mechanical Coupling Induced Performance Degradation in Storage Systems

Sanjeev Sondur and Kenny Gross
Oracle Corporation
Redwood Shores, CA 94065, USA
[sanjeev.sondur|kenny.gross]@oracle.com

Krishna Kant
Temple University
Philadelphia, PA 19122, USA
kkant@temple.edu

Abstract—This paper explores the coupling between power/thermal aspects of data center and the vibrations caused by chassis/server fans in environments dominated by mechanical disks. We show that the leakage power and cooling fan induced vibrations could substantially degrade the storage system performance which in turn could lead to a poor performance and high energy consumption. We then propose our Environment Aware Data Control (EADC) management policy that utilizes built-in instrumentation telemetry and fan-speed control to minimize overall energy consumption while providing a better overall performance. Based on actual experimentation with storage servers in Oracle labs, we show that the proposed mechanism can reduce the energy consumption by 63% while simultaneously reducing the IO delay by 60% as compared to the baseline system with no vibration aware controls.

Index Terms—Data backup and retrieval, Leakage current, Vibrations, Power Consumption, Thermo-Mechanical Coupling, Disk Performance

I. INTRODUCTION

Data storage is growing at an unprecedented rate, and large enterprises such as Facebook, Google, and YouTube are increasing their storage capacity by almost one Petabyte every day [1]. Storage systems continue to increase in complexity both in terms of data storage functionality and processing of this data. On the processing side, storage systems may do numerous compute intensive operations such as erasure coding, deduplication, compression, compaction, replication, etc. On the data storage side, a storage system may consist of multiple storage types and technologies, and data migration across them.

Storage systems are often classified as primary (for actively used data and applications), near-line (for infrequently used data and applications), backup (used for data protection and increasingly performed continuously), and archival (used for long-term preservation/record keeping). Each of these could potentially involve a tiering hierarchy using different storage technologies. For example, the primary storage may consist of SSD and fast-HDD tiers, and near-line storage may consist of normal and SMR HDD tiers. This paper is concerned with the HDD dominated parts of the storage since its primary aim is to study the *cumulative* impact of passive energy and fan induced vibrations on the mechanical drives. Although SSDs continue to increase their penetration in the primary storage, HDDs will likely stay put in other parts of storage

because of low cost, lack of endurance issues, and continuing technological developments. For concreteness, we focus on backup storage (Backup and Retrieval - *BaR*) in this paper, which are and will remain HDD dominated, but much of our analysis and observations apply to other situations that include HDD storage.

To guarantee data integrity, *BaR* systems comprise of additional hardware and software components like RAID disks, erasure coding, replication nodes, etc. For the backup storage, the primary SLA requirement concerns the time to recover the archived data and restore normal business operations. However, the backup performance is also increasingly important because of the “continuous backup” which needs to be done without any visible impact on the responsiveness of the system. Usually, storage systems use commodity grade hard disk drives (HDDs) [2], which are subject to rather high failure rates [1]. The delays caused by failures and recovery (e.g., rebuilding RAIDs) can impact the performance of such systems adversely.

Based on the data collected from servers, we show that the compute operations in a *BaR* system could result in significant CPU load which could, in turn, degrade the storage performance. For this, we define a *Environment Aware Data Control* (EADC) policy to improve the IO efficiency and minimizes the overall energy consumption. We do this by exploiting the instrumentation built into the modern servers to measure the low-level attributes [3]. By using environment and performance awareness in the form of disk and thermal characterization curves, EADC controller can make intelligent choices to mitigate the parameters affecting the poor disk performance and improve *BaR* efficiency. The results show that the proposed mechanism can reduce the energy consumption by 63% while simultaneously reducing the IO delay by 60% as compared to the baseline (i.e. no vibration aware controls).

The main contributions of this paper include:

- 1) Articulate the challenges in meeting the QoS/SLA requirements of a data backup/archival systems and highlight the need for studying the physical and thermo-mechanical properties together.
- 2) Present an environment aware management policy that mitigates the cumulative inefficiencies caused by multiple components and improves the overall performance of the *BaR* systems.

The rest of the paper is organized as follows. We start with a brief discussion of the related work on the topic in section II. We then discuss the thermo-mechanical coupling in storage systems in section III and provide a simple power consumption model in section IV. We describe the evaluation methodology and experiments conducted in section V. In section VI, we present the EADC management policies and the empirical results in section VII. We conclude the paper in section VIII.

II. RELATED WORK

Vibration-damping mechanisms do exist in data centers, but only target well-known sources such as the spinning hard drive motors themselves, physical drops, and HVAC building cooling systems [4]. Their success metrics are geared toward lowering hard drive failure rates, generally caused by head crashes (i.e., when the read-write head makes contact with the disk platters, causing irreversible damage) [5] and not towards performance/energy efficiency. Because of their focus on regular external vibrations, these mechanisms cannot react quickly enough to the changes introduced by variable fan speeds [6]. Wang et al. [7] mention fan-related resonance in their thermal and MIMO fan control model, but do not examine the root cause of such vibrations. Their fan control policy may show good power efficiency but ignores total runtime energy consumption. Joint energy thermal and cooling (JETC) management policy proposed by Ayoub [8] is the closest solution to our approach, but JETC policy operating at near-high thermal limits is counter-productive to performance since it degrades disk throughput.

Nachiappan [1] presents a summary of solutions, most of which are targeted towards reducing latencies, data redundancy techniques, geographical distribution of storage nodes, energy efficient storage, etc. A survey of thermal management is covered in Chaudhry [9] and discusses solutions that use multiple independent controllers to reduce thermal gradient, hotspots, and cooling magnitude., and it does discuss fan control and DVFS, both of which we consider as well, but not the coupling of vibrations and thermal management.

Wang, et.al. [10] develop a multi-controller model for server power capping and draw a relationship between power consumption, operating frequency, P-state and server utilization. Our proposed policy includes a oscillation mitigation control and considers P-state efficiency as one of the driving factors. Bartolini [11] support our study about higher total energy consumption at low frequency (i.e. high P-state) with their DVFS at the wall model and show power/energy and performance tradeoff. However, their multi-core model does not extend to include additional energy costs from fans and temperature related disk IO degradation. In contrast, we explore a holistic tradeoff that integrates DVFS, thermo-mechanical, physical and power characteristics.

III. THERMO-MECHANICAL COUPLING IN STORAGE SYSTEMS

In this section, we discuss both physical and architectural aspects of storage systems to bring out the thermo-mechanical

coupling that we explore in this paper.

A. Mechanical Aspects of Storage Systems

A typical storage server organization is shown in Fig. 1. The compute units (4), storage units (1) and cooling units (2) are mounted on the same physical and mechanical (3) plane. Air flow (shown by directional arrows) is controlled by on-board variable speed fans, that keep the component thermal footprint within safe operating limits. Modern servers typically have six or more fans with uniform voltage-RPM ratios and are operated with pulse width modulation (PWM), where the fan speed is adjusted by varying the duty cycle of a control signal. This density adds adversely to the thermal and environmental profile of the unit.




Fig. 1. Server Organization

Current HDDs have disk platters spinning up to 15,000 rotations per minute (RPMs) with the mechanical arm mounting the read/write head about 7 nanometers away from the data tracks (20 nanometers wide). At these extreme limits, the disk drive should be at a perfect standstill to have minimum read- and write-retries to achieve a maximum data transfer. Clearly any microscopic movement would disturb the head and data track alignment and degrade data throughput rates. It is well established that vibration induced errors disturb the hard disk component alignment resulting in read/write errors [12]. Thus, needing additional read and write operations increases the data transfer latency. HDD vendors acknowledge that rotational vibration from disk actuation and external forces as an area of concern for disk data transfer speeds [13].

B. Thermal Issues in Storage Systems

During data reconstruction (or backup), storage servers execute compute intensive tasks such as de-duplication, erasure coding, encryption logic, etc. Adding in either RAID or data-fault-tolerance makes computations more complex and adds to CPU-intensive IO coordination.

Storage providers build fault tolerance by keeping two or three copies of every piece of data. It is well documented that HDDs fail at a fairly high rate [1, 14]. Assuming a service provider's claim of 1% per year AFR (Annualized Failure Rate), with over 20,000 HDDs stacked in one section of a data lake, the data reconstruction rate for failed HDDs is about 4 HDDs per week. Data reconstruction is a CPU intensive task to re-distribute the many chunks of data so that there are once-again two or three copies of every piece of data. With current high density drives, *Oracle lab* studies have shown that this reconstruction operation can take about 36 hrs. Any unwanted disk IO latency will only make this time factor worse.

The compute overheads during reconstruction heat up the CPU, which has to be cooled by the on-board server fans. Modern servers have an array of PWM controlled fans [15] that are tasked to cool the CPU cores to acceptable operating limits. Power consumed by the array of fans is a significant

factor in the overall power consumed by the server (60W to 300W per server). Higher operating core temperatures also increase the leakage power in the system. Temperature dependent leakage power contributes up-to 40% of the total CPU power. Fan power, leakage current, undesired vibration induced disk latency add up to passive-power that does not directly contribute to work done, and decreases the overall energy efficiency of the system.

C. QoS Requirements in Storage Systems

The QoS requirements in storage systems vary depending on the type of storage such as primary, near-line, backup and recovery (BaR), etc. Implementing differentiated services in large primary storage systems involving high speed storage can be very challenging; however, in HDD dominated systems such as BaR, IO bandwidth and latency are the key QoS factors. IO retries can significantly affect the overall latency and hence the high layer components such as the RAID operation.

Enterprise-class users specify the backup and retrieval QoS/SLA with a recovery time objective and recovery point objective [16]. Recovery time objective (RTO) is the user defined policy, defined as the time required to recover from data loss and get back to normal operation. RTO is dictated by the end-user based on their business continuity needs, and ability to operate the businesses successfully without this data. Recovery Point Objective (RPO) is the desired amount of time between data protection events. It should be noted that both RTO and RPO are user defined policies based on the criticality and the business value of the data being protected - and the Cloud storage provider has to guarantee these SLA/QoS demands. The stringent nature of RTO/RPO is shown in Table I for different classes of user-data [16].

	Critical	Vital	Sensitive	Noncritical
Avg. data distribution	15%	20%	25%	40%
Availability index	99.999%	99.99%	99.9%	99%
Downtime min/year	5.256	52.56	525.6	5256
Typical RTO (RPO)	Device	Device	Device	Device
30 min (0 min.)	Disk	Disk	Disk	Disk
2 hrs (15 mins.)	Disk	Disk	Disk	Disk
12-24hrs (2-6hrs)	Disk	Disk	Tape	Tape
≤1 day (12-24hrs)	Disk/Tape	Disk/Tape	Tape	Tape
≤1 week (< 1 day)	Tape	Tape	Tape	Tape

TABLE I
TYPICAL RTO/ RPO OF DATA^[16]

Table I shows that mission critical applications can tolerate a downtime around 5.26 min/year and constitute about 15% of back-up data (user population). Such businesses can tolerate a maximum retrieval time of 30min; within which the data has to be restored for normal business operations. Typical mission-critical RTO is to restore the data within 0-6 hrs (1st three rows of RTO table).

IV. POWER CONSUMPTION MODEL OF STORAGE SYSTEMS

We explore the power consumption of a storage unit in distinct parts, (i) hard disk, (ii) CPU cores and (iii) fans - as these dominate the total energy consumed. The largest

Symbol	Description	Symbol	Description
E	Total Energy	t	Total time
P	Total power	t_{disk}	Disk IO time
t_{cpu}	CPU time	P_{other}	Power consumed by other components
P_{disk}	Disk power	P_{cpu}	CPU power
P_{fan}	Fan power	FS	Fan speed
I_{idle}	Leakage Current	d_{tr}	Disk transfer rate
ω	Vibration	T_{cpu}	Core temperature
T	Thermal Head Room Margin (THM)	T_{amb}	Ambient Temperature
T_N	Normal core temp.	T_D	Core temp. at DVFS
f_N	Normal Operation Frequency	f_D	Operation Freq. at DVFS
$t_{cpu;N}$	Normal CPU execution time	$t_{cpu;D}$	CPU execution time under DVFS
V_N	Normal Voltage	V_D	DVFS Voltage
$I_{idle;N}$	Normal I_{idle}	$I_{idle;D}$	DVFS I_{idle}
$E_{idle;N}$	Idle Energy under DVFS	$E_{dyn;D}$	Dynamic Energy under DVFS

TABLE II
NOTATIONS FOR POWER CONSUMPTION MODEL

power consumers are the processor and the cooling subsystem, consuming around 37% and 29%, respectively (See Fig. 2).

Though our discussion focuses on the main CPU cores, most electronic components on the motherboard including DIMMs, north/south bridges, various links and controllers, voltage regulators, etc. all exhibit a similar temperature dependent behaviour. In particular, the leakage current and hence idle power increase rapidly with the temperature. Also, the dynamic power is governed by the switching rate and voltage. For simplicity, their consolidated power is represented as P_{other} in our discussion.




Fig. 2. Server Power Consumption Breakdown^[17]

Li [15] show that the power consumption in the disk controller is independent of the workload. A study by Google [18] shows very little correlation between failure rates and activity levels (disk IO/workload). The power consumption of a HDD unit is dominated by the spindle motors, and remains almost *constant* with respect to time. HDD power during the read/write operations puts very little variation on this flat power consumption line.

Intel [19] and Sun Microsystem [20] have demonstrated the exponential relationship between temperature T and leakage current I_{idle} . Intel's work on impact of technology scaling on thermal behavior shows the relationship of core leakage current with respect to operating temperature. Using a 8 core processor, Sun Microsystem's work demonstrates that leakage power is exponentially related to temperature (Eq. 1), and a positive feedback loop exists between temperature and leakage, which can cause dramatic increases in temperature and damage the circuit if not controlled.

$$I_{idle} \propto e^T \quad (1)$$

Dynamic Voltage and Frequency Scaling (DVFS) control built into the CPU core regulates the operating temperature.

DVFS triggers core p-states to reduce the operating voltage and frequency to compensate for rising core temperatures.

Assuming that a workload has path-length (i.e., number of instructions) k , the time to execute a workload t_{cpu} at clock rate f is written as:

$$t_{cpu} = k \cdot (f) = f \quad (2)$$

where \bar{c} is the average cycles per instruction (CPI), and is itself a function of f . In particular, $(f) = \bar{c} + (f=f_0) \cdot \lambda$ where \bar{c} is the core CPI (assumed independent of f), $\lambda \ll 1$ is the relative impact of uncore latency, and λ is the uncore latency in cycles at some baseline frequency f_0 . Using notations in Table II, the total energy consumption E_D under DVFS can be written as follows:

$$E_D = [V_D I_{idle;D} + 0.5 C V_D^2 f_D] t_{cpu} \quad (3)$$

where the first part in brackets is the static power and second part dynamic power.

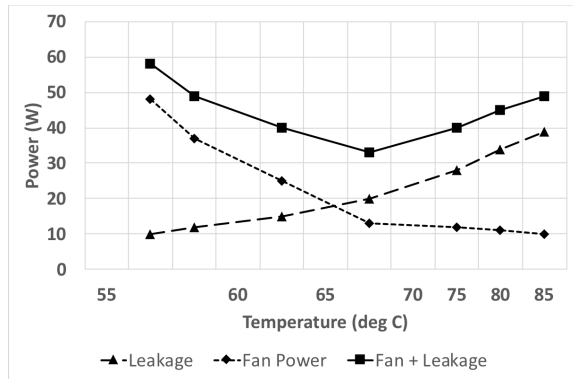


Fig. 3. Leakage Current & Fan Power vs. CPU Temperature

Operating at low frequency f_D increases the workload execution time $t_{D;cpu}$, but the dynamic energy reduces because of the lower voltage V_D . The combined E_D may be less or more than E_N depending on the voltage decrease and the processing time increase. However, passive energy from other sources ($P_{idle} + P_{fan} + P_{other}$) will increase because of longer execution time $t_{cpu;D}$. Intelligent fan controllers regulate the fan speed to keep the core temperature within the safe operational limits. Higher fan power increases air-flow, which reduces the temperature and thus keeps the leakage current low. On the other side, keeping the fan speed low conserves fan power, but increases the core operating temperature and in-turn increases the leakage current.

Using empirical data from Oracle labs, we captured the relationship between the cumulative fan power and leakage power with respect to temperature, as shown in Fig. 3. Server or chassis mounted cooling fans contribute to ambient vibrations that negatively affect the disk throughput as shown later. These insights help us decide the important attributes that contribute to the power consumption in a storage server.

In the following we write equations for the power consumption of a storage server based on the notation given in Table II:

$$E = P t; \quad P = P_{cpu} + P_{disk} + P_{fan} + P_{other} \quad (4)$$

$$t = t_{io} + t_{cpu}; \quad T = T_{cpu} - T_{amb}$$

We can observe the following properties from these equations (\Rightarrow indicates monotonic relationship):

$$t_{cpu} \Rightarrow T \text{ (effect of DVFS); } \quad t_{io} \Rightarrow d_{tr}^{-1} \Rightarrow !$$

$$! \Rightarrow FS; \quad FS \Rightarrow P_{fan} \Rightarrow T$$

$$P_{cpu} \Rightarrow I_{idle}; \quad I_{idle} \Rightarrow T \quad (5)$$

The above equation uncovers an important factor - that the temperature T plays an import role in defining both the disk access latency and energy efficiency. Keeping the ambient temperature constant T_{amb} , the allowable CPU operating temperature T_{cpu} drives most of the attributes that contribute to passive energy wastage results. One option is to keep the fan speed high (i.e high P_{fan}) to maintain a low T_{cpu} , thereby decreasing I_{idle} , but this option increases $!$ and t_{io} . On the other hand, conserving fan power (low P_{fan}) results in higher T_{cpu} and low $!$ and low t_{io} ; but it adds to higher I_{idle} and high t_{cpu} (because of DVFS).

Vendor	Power Rating	BTU/hr Rating	Oper. Temp.	Operating Elevation (ft)	HDD Slots
HP ^[21]	1000W	4500	10 to 35°C	0 to 10,000	6
Dell ^[22]	600W	2047	10 to 35°C	-50 to 10,000	12
Oracle ^[23]	550W	1877	5 to 35°C	0 to 10,000	16x20

TABLE III
BACKUP STORAGE SERVERS SPECIFICATIONS

Table III gives a sample of vendor specifications for *BaR* systems of different storage capacities. We consciously present additional data points relating to operating elevation, operating temperature and thermal BTU/hr rating. Operating elevation (the altitude of the data center) influences the cooling parameters and has a direct influence on fan speed, as fans have to operate at higher speeds to push the thinner air at higher elevations. Operating temperature is an influencing factor in multiple parameters - thermal headroom margin available for safe operation, leakage current, and fan speed (and indirectly disk throughput). About 40% of the power given under Power Rating is idle power and cannot be ignored while calculating energy efficiency. Assuming each HDD slot has a 10TB drive, we can envision the effect of poor disk IO performance in a single server. The seriousness of performance and energy consumption problem magnifies when we include the 100s of 1000s storage units in a typical data center.

A. Thermal Head-Room Margin

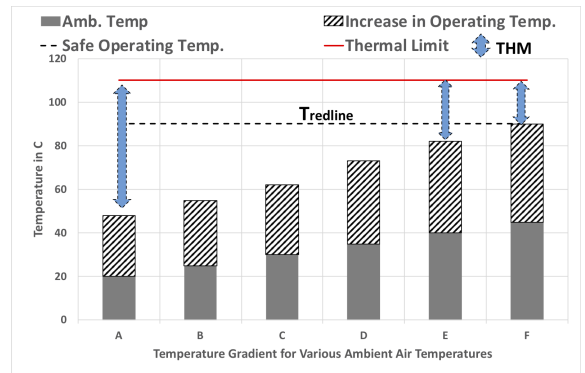


Fig. 4. Thermal Headroom vs. Temperature

Thermal headroom margin (THMs) is the difference between the allowable safe operating temperature for internal components and the actual real-time temperature of the components. A sample of THM for different operating temperatures of the servers T_{cpu} under various ambient temperatures T_{amb} is shown in Fig.4. Servers can be operational up-to a safe operational limit $T_{redline}$, beyond which it affects the component reliability and could trigger thermal trip events in the CPUs to automatically shut off the server. Fig.4 clearly shows that T_{amb} significantly defines the available THM, thereby restricts the operational thermal profile and the power consumption of a server.

Our work aims at finding an optimal operating point that maximize the performance gain and minimize the effects of thermo-mechanical coupling, while maintaining a safe operational point for the current workload.

V. EVALUATION METHODOLOGY

The events that lead to customer necessity for recovery from backup media are random and are not any more or less likely to occur depending on the nature of the workloads generating the data. In absence of *BaR* specific data, we used widely accepted TPC-H/SPEC-CPU 2006 workloads and present the research work that helped us capture the various metrics to establish the relationship in Eq. 5. Our evaluation, data collected and analysis support the discussion presented above.

A. System Configuration

First, we present the experimental setup to characterize the disk throughput degradation because of ambient vibrations. Our test server is comprised of 8 cores, 6 PWM controlled fan modules, and 8 disk drive slots populated with SATA drives. The server consumes between 330W-600W depending on utilization, while the CPU socket itself has maximum thermal design power (TDP) of 240W.

We run a parametric characterization suite of workloads (see Table IV) to measure the vibrational sensitivity on the disks. TPC-H is a decision support benchmark representing database requests. The queries comprise combinations of operations such as sequential scan, index scan, merge join, and hashing functions. The nomenclature in Table IV, workload W8 "3-B, 2-H, 1-Q10" refers to 3 instances of SPEC BZIP2 benchmark, 2 instances of SPEC HMMER benchmark, and one instance of TPC-H query no. 10. We choose queries 1, 3, 10, 13 and 19 to represent data-intensive service jobs for TPC-H. Of the SPEC suite, we chose *BZIP2*, a compression algorithm, and *HMMER*, a compiler workload. In each workload set (shown in Table IV), the processor is stressed up-to 75% utilization, running a single TPC-H query and five SPEC tasks.

Wkld	Benchmark
W1	5-B, 1-Q1
W2	5-B, 1-Q3
W3	5-B, 1-Q10
W4	5-B, 1-Q13
W5	5-B, 1-Q19
W6	3-B, 2-H, 1-Q1
W7	3-B, 2-H, 1-Q3
W8	3-B, 2-H, 1-Q10
W9	3-B, 2-H, 1-Q13
W10	3-B, 2-H, 1-Q19

TABLE IV
WORKLOADS FOR DISK
VIBRATION TEST

In this experiment, we disable the buffer cache that would have hidden disk access latency from the user. We run a pure IO generator which issues random writes to the disk, utilizing 100% of the I/O bus bandwidth to expose and isolate the effect that vibrations have on disk throughput. The impact of disk throughput degradation on the overall performance of the realistic database benchmarks vary depending on the behavior of each workload.

B. Vibration and Disk Throughput

Fan sweep test: For this test, the server is bolted to the stationary shake table, which is powered off and provides assurance the server is decoupled from any possible ambient vibration sources. To study the effect of internal vibrations, we initiate the pure IO generator and then manually sweep through the range of possible fan speeds while monitoring the average write throughput as a function of fan speeds.

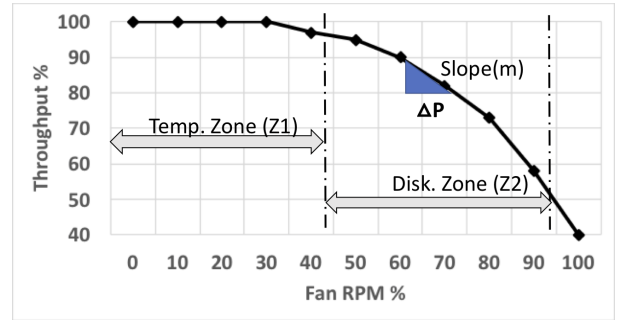


Fig. 5. Fan Speed vs. Disk Throughput

We step through fan speeds from 100% to 0% PWM at 10% step sizes, pausing for at least 20 seconds at each step-change in stimuli to allow the IO throughput metrics to equilibrate at the new fan speed and obtain stable results.. Figure 5 shows the average degradation of write throughput versus fan speeds, normalized to the maximum throughput measured for each disk. There are minimal vibrational effects below 40-50% PWM. Above 50% PWM, the degradation in IO throughput becomes pronounced and the average achievable throughput falls off to about 40% (i.e. 60% degradation) when the fan speeds are running at maximum PWM.

ID	CPU sockets	cores/CPU	Thermal Design Power	No. of Fans	Fan Power
Q	2	4	80W	2	60 W
E2	2	8	100W	2	60 W
E4	2	8	100W	4	200 W
E6	2	8	100W	6	300 W

TABLE V
SERVER SPECS USED FOR FAN SPEED TEST

Amplitude test with random frequencies: This experiment studies the effect of external vibrations by reproducing the range of frequencies and amplitudes obtained from a real operative data center. We used tri-axial accelerometers to measure vibrations in a server rack inside an operational data-center, then reproduce the environment with an Unholtz-Dickie model K170 electrodynamic programmable vibrational table. For this experiment, we are exploring the effects of

external rack vibration levels. Hence we lock the fan speeds to a constant value, which is 50% PWM for this sequence of measurements. The results are shown in Figure 6. Higher vibration amplitude results in lower disk throughput.

C. Fan Speed and Vibration

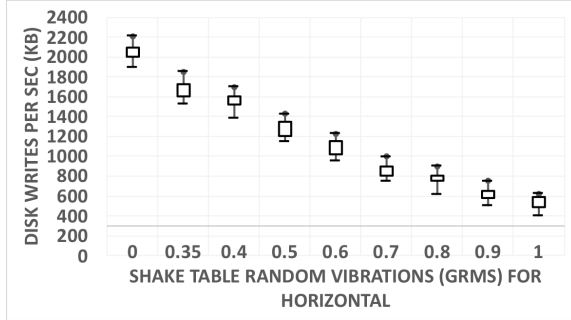


Fig. 6. Disk Throughput vs. Vibration

The next set of experiments establishes that the speed of the fan motors is the root cause of ambient vibration observed in the server. Our evaluation model consisted of four different server configurations, based on a current high-end enterprise servers, with specifications shown in Table V.

Workload execution causes an increase in CPU core junction temperatures and this heat has to be extracted to keep the components under safe operating limits. We conducted experiments to capture metrics related to core temperature, fan speed, fan power, and vibrations.

Wkld	Skt-A	Skt-B
Server Q		
W1	2-H + Q1	3-H
W2	2-H + Q13	3-H
W3	2-H + Q19	3-H
Server E2, E4, E6		
W4	5-H + Q1	6-H
W5	5-H + Q13	6-H
W6	5-H + Q19	6-H

TABLE VI
WORKLOAD COMBINATIONS FOR
FAN SPEED TEST

As expected, the results show that the aggregate fan motor power increases as the cube of the fan RPM. Through empirical means, we verified that the vibration amplitudes go up linearly with fan motor PWM.

Since the engineering interfaces to the internal fan control algorithms are not exposed to the end user (and therefore cannot be directly modified for testing purposes), we use simulations to investigate the combined effects of fine-grained fan control combined with workload scheduling. We characterized the effect of fan speeds on hard disk bandwidth through experiments on the E6 server, then separately characterized the effect of this hard disk bandwidth by modifying delays in the IO driver and monitoring the total execution time for a benchmark. With data from these experiments, we can estimate the dependence of application performance on fan speeds without an overhaul of the physical system design.

We executed workloads shown in Table VI for the two CPU sockets. For example, workload W4, characterized as "5-H, Q1, 6-H" refers to executing five SPEC HMMER tasks and one TPC-H query 1 on CPU socket A and executing six HMMER tasks on CPU socket B. We selected the queries that have intensive disk usage and relatively low CPU activity (query 1,

13 and 19) since they represent the range of disk bandwidth and CPU utilization across the benchmark set.

In Figure 7, we present the energy required to complete a fixed customer workload (e.g. updating a Multi-TB database) under different fan speeds (and hence vibration levels in x-axis). As fan speeds increase to keep cool the components to a safe operating limit, disk IO throughput performance drops (Fig. 6). Consequently, time taken to complete a given customer workload goes up (Fig. 7).

As expected, the energy required to complete the workload goes up as the ambient vibration levels go up. For example, if IO performance degrades by 20% due to an increase in fan speed (hence elevated ambient vibrations), then it takes 25% longer for the customer workload to complete. This means that all the other components inside the server (fan motors P_{fan} , HDDs P_{disk} , memory, ASICs, IO cards, and PSUs P_{other}) are consuming power for 25% longer period of time.

D. Cumulative Leakage and Fan Power




Fig. 7. Energy Consumption vs. vibration

We coupled the study of temperature dependent leakage current alongside fan power consumed in cooling the components (within operation limits). We used two enterprise class servers with CPUs in 2 sockets, 32 8GB memory DIMMs, 2 HDDs, and 6 fans distributed in 3 rows of 2 fans. To enable accuracy quantification of the leakage power component and fan power, we first characterize the fans by verifying their speed with vibration sensors and obtaining their power consumption values at each RPM setting. We used independent power supply units to directly control each pair of fans. Using built-in Continuous System Telemetry Harness (CSTH), we collected (i) 4 CPU temperature values (2 thermal sensors per die); (ii) 32 memory temperature values (1 per DIMM); (iii) per core voltage and current values; and, (iv) power consumed by the whole system. This data polled every 10 seconds provided sufficient visibility into the run time power and thermal behavior.

We explore all ranges of utilization scenarios using *LoadGen*, a customized dynamic load-synthesis tool that: (i) uses a core algorithm to fill maximal instruction pipes of the multi-threaded CPUs; and (ii) allows customized dynamic profiles that can meet any desired utilization level by very rapid duty-cycling between idle and 100%. While running *LoadGen*, the system is guaranteed to maintain the given CPU utilization and the workload is evenly spread among the cores. We plot the cumulative power using the power

measurements and temperature at various utilization levels. Figure 3 shows the fan power, the leakage power, and their sum for 100% utilization. The sum of leakage and fan power is a convex-like curve that reaches a minimum around $67^{\circ}C$, which corresponds to a fan speed of 2400RPM. This dataset confirms the influence of temperature T on two important parameters, I_{idle} and P_{fan} discussed in Eq. 5.

With the data collected from the above experiments, we proceed to define the management policy to locate an optimal operating point that minimizes the negative effects of (i) fan power, (ii) vibration (indirectly disk throughput degradation) and (iii) leakage current.

VI. PROPOSED MANAGEMENT POLICIES

The goal of *Environment Aware Data Control* (EADC) management policy is to find an optimal operation point that results in the best possible performance and power consumption. There should be a trade-off between achieving maximum performance vs. conserving maximum energy. Both maxima cannot be achieved because of the dynamic nature of the thermo-mechanical characteristics that counter one another.

A. Optimal Operational Point

In Eq. 5, T dictates most of the inefficiencies: I_{idle} , P_{fan} , $DVFS$, and d_{tr} . The cumulative effect of T on various parameters is depicted in Fig. 8, with the x-axis showing various operating temperatures and y-axis the normalized values for various parameters. The graph lines show the influence of the rise in temperature across various parameters (given in legend labels). The influence of temperature on THM as explained in section IV-A is given in Fig 4. The figure shows that THM shrinks as the operating temperature increases, negatively affecting component reliability.

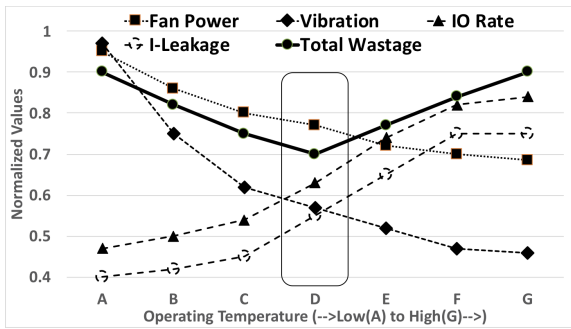


Fig. 8. Optimal Operating Temperature

To maintain low operating temperatures (left extreme, point A & B in Fig. 8), the fans have to operate at higher speeds and consume high P_{fan} , contributing to high vibrations, low disk transfer rates, and high IO time. The leakage current I_{idle} is lower at the low temperature. Cumulative power consumed at various temperatures is given with a thick solid line, showing higher total wasted power (see Fig. 8).

On the other hand, maintaining a high operating temperature (right extreme, point F & G in Fig. 8), needs less fan power P_{fan} but increases the leakage current I_{idle} . In addition, at

this point - the THM is very low and adversely affects the component reliability. The cumulative wasted power is again high owing to the high I_{idle} at higher temperatures.

The optimal operational point is a trade-off to minimize the wastage between the two extremes given above; and dynamically changes based on workload dependent core temperature (T_{cpu}). These optima shown as a boxed area in Fig. 8, refers to the point where these adverse reaches a minimum.

B. EADC Management Policy

Figure 9 shows the EADC system with controller, fan actuators, and sensors (thermal, power, fan speed, disk throughput). EADC aims to improve the server performance against (vibration induced) disk degradation and compensate for internally generated vibrations. To achieve this, EADC expands upon the concept of temperature estimation with band-limited predictors [24] and directly manages the internal fan vibrations. To quantify the effects of fan vibrations, the fan module monitors the fan PWM RPM tachometer sensors and uses those measurements in real time to estimate the corresponding disk throughput degradation using a trained inference model as shown earlier in Figure 5.

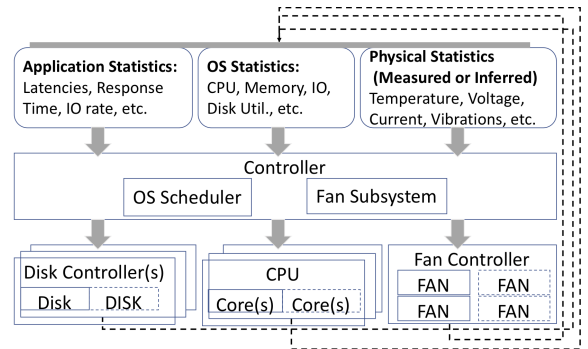


Fig. 9. EADC Controller Framework

EADC policy does not need prior I/O characteristics to be self-reported at the application level, or intrusive system calls to monitor I/O current and near-term utilization. The fan-disk characterization curve (such as Figure 5) is segmented into either temperature-driven zones (Z1) or disk-performance-driven zones (Z2) according to the slope (m) of the curve. Fan settings are discretized based on the disk sensitivity slope (m) in Figure 5. At each scheduling tick, the controller observes the fan speed, core power, and core temperatures. The difference between the thermal set point and the current measured temperature T , determines the desired change which is achieved with a proportional change in fan speed setting P . This change may be bounded either by physical constraints (the maximum speed of the fan) or a management policy (e.g. limit on fan speed to ensure good disk throughput), and is enforced by the fan actuator firmware.

If the workload results in moderate T_{cpu} temperatures and low fan speeds (zone Z1), the slope is flat ($m = 0$), indicating that workload and I/O throughput are independent of fan speeds in this operational regime. Consequently, in this operational regime, fan speeds are dictated only by the highest

measured core temperature and the preset Thermal Headroom Margin (THM).

If the workload results in high T_{cpu} (compute intensive workload), this triggers a cooling response with high fan speeds. At higher fan speeds (zone Z2), the slope becomes more negative and the policy enters a disk-driven zone, identifying times where disk sensitivity is particularly high relative to the fan speed. As we measure disk I/O throughput activity as part of our metric, the guideline for balancing fan speeds is relaxed in order to reduce vibrations. At thresholds between some fan step n and $n+1$, the controller assigns the minimum fan speed to obtain a higher gain in performance (reduced disk IO degradation), with lower cooling capabilities. EADC policy controls the fan speed while maintaining optimal operation performance (minimize vibrations) and maintaining a safe THM. The policy will throttle CPU activity via DVFS as the last resort if the cores reach the emergency temperature threshold.

EADC controller is implemented as a finite state machine triggered by thermal, power and tachometer feedback from the fans. The controller outputs the target power distributions for the CPUs and memory and a target temperature for the cooling system. The independent actuators in CPUs, memory and fans respond through workload scheduling, page migration and fan speed control respectively.

In summary, EADC targets variation in power consumption and achieves the end-goal as below: (\downarrow indicates decrease, \uparrow increase and \Rightarrow indicates monotonicity)

$$T \downarrow \Rightarrow P_{idle} \downarrow \quad FS \downarrow \Rightarrow P_{fan} \downarrow \quad ! \downarrow \Rightarrow d_{tr} \uparrow \Rightarrow t_{io} \downarrow$$

$$Application\ Speed \uparrow \Rightarrow Server\ Efficiency \uparrow \quad (6)$$

We track the disk intensity of workloads by monitoring IO throughput per second via the OS. The controller makes scheduling decisions at the granularity of 10ms while fan control and IO monitoring are done at 1 second intervals corresponding to server enclosure thermal time constants. At each fan control interval, there is a gap between heat dissipation of the cores T_{cpu} and the thermal set point (THM), which is usually closed by provisioning fan speed P_{fan} . During periods of significant disk activity, we only allow fans to spin as fast as the *specified* limit, while *progressively throttling* the workloads with lower disk access rates to meet thermal constraints. During periods of low disk activity, the limit on fan speed is lifted, and control returns to the default JETC policy.

Policy for Admission Control: Admission control policy schedules the tasks based on QoS class (a.k.a request priority) and the current thermal state. Under high core temperatures, EADC policy admits high priority requests that need disk IO thus satisfying their SLA first. Allowing CPU intensive requests will only worsen the thermal profile and triggers both the DVFS frequency and higher fan speeds - both of which are undesirable. During non-thermal peaks (Zone Z1 in Fig 5) the policy admits most requests within allowed THM, thereby satisfying maximum client SLA.

C. Mitigating Oscillations Through Stability Control


Time lag and quantization in temperature sensors (and other measured parameters) lead to stability concerns when multiple local controllers are running together. Oscillations occur when multiple controllers compete for control over fan speed, DVFS switching, thermal/power profile, etc. For example, when workload related T_{cpu_2} increases to $75^\circ C$ - fan speed would increase to speed (say S_2) to cool the components, additional air flow at S_2 can cool the core temperature to $65^\circ C$ causing the cooling system to decrease the fan speed to S_1 ($S_1 < S_2$). This reduced fan speed may heat the components again demanding higher air-flow (and in turn higher fan speed), and so on. Time lag measurement is caused by a lag in actual increase in metric to sensor detected metric. During our research, we measured power sensor measurements following workload changes in a CPU and observed that the change in temperature T_{cpu} suffers a 10 second lag. We observed that both fan-speeds FS and p - states in the CPUs can go into oscillations under various conditions. Further, our research shows that multiple controllers working together (without coordination) can create "competitive oscillations".

To avoid oscillations, we built a global hierarchical stability control model inside the proposed controller to jointly determine the optimal fan speed FS , and maximum allowable CPU utilization (a.k.a CPU cap), to maintain the CPU operating temperature T_{cpu} within a safe operating region (THM) (e.g., under $80^\circ C$). The controller consists of: (i) multiple local controllers and (ii) a global controller. In the target architecture, we have two independent local controllers for fan speed and CPU cap controls. The details of the mechanism are contained in a paper by Oracle and EPFL [25] and are omitted here for brevity.


VII. RESULTS

To help understand the results, we first give an overview of work closest to our research, and then present the results in terms of performance gains and energy savings.

We compare our policies with: (i) Dynamic load balancing (DLB) policy used in modern operating systems, which uses thread migration to minimize the length difference between task queues, but does not monitor temperatures. It uses a proportional-integral (PI) fan controller and keeps all modules active (hence max P). We consider this as our default baseline. (ii) Dynamic Thermal Management (DTM) policy proposed by Choi [26], wherein the authors migrate workloads proactively to alleviate the maximum core temperature and the memory activity is throttled when the temperatures approach the emergency threshold. They use the default workload-independent PI based fan actuator. (iii) Joint Energy, Thermal and Cooling (JETC): Ayoub [8] proposed a JETC combined solution for core scheduling, memory module activation and fan control. JETC is not explicitly aware of disk throughput degradation, and tends to operate closer to peak thermal bounds than the default, so disk-bound workloads suffer due to high fan speeds and vibrations, while CPU-bound workloads suffer if core temperatures hit the thermal threshold and emergency (DVFS)



(a) Hard Disk Wait Time (Delay Reduction)



(b) Energy Saving (Total Work Done)

Fig. 10. Comparative Results for EADC Performance and Energy

core throttling happens. Also, by operating at peak thermal bounds, JETC threatens the component reliability.

A. Performance Gains

We compare results for three thermal/cooling management policies, DTM, JETC and EADC, against the default DLB policy across four server configurations (Table V) and workload is shown in Table VI. The core thermal set point and fan trigger point is set at 85°C and the emergency threshold 90°C . The benefits of EADC policy in terms of the reduction of IO wait time is shown in Fig. 10(a). In server E4, we observed that the fan causes some vibrations, but DLB still performs well because thermal management through core thread migration mitigates a significant portion of wasted power. A larger cooling system in E6 drives up IO delays, requiring policies to mitigate that effect. With EADC, we mitigate fan-induced IO degradation (in E6), so the maximum fan PWM is limited to 40%, thus reducing the IO delay by 60%. This is the near minimum level possible in the system, as it detects disk access patterns and ensures that the fan speed stays low enough to minimize the effect on disk throughput.

Server E2 (not shown in figure) has the highest level of DVFS throttling (0.77%) due to a severely under-provisioned cooling system, wherein EADC proactively throttles jobs to meet thermal constraints even with maximum fan speed. In server E6, the average performance hit due to throttling is 0.4%, with a maximum of 0.66% for workload W4. DTM and JETC slow down the cores reactively at emergency temperatures, causing instability and lower efficiency of heat removal. JETC as explained earlier, allows average temperatures to rise above the thermal set point in favor of lower fan power (hence higher vibrations, low throughput and longer execution time). EADC proactively controls temperatures before they reach the emergency threshold (thus avoiding triggering DVFS), thus keeping fan speeds balanced and ultimately reducing disk vibrations.

For example, while running queries on E6 server with default DLB policy, the disk bandwidth degradation is around 60%. EADC policy reduces this overhead and achieves an overall application speed up of 1.35x for TPCHQ-19. On an average (across workloads and servers), EADC achieves on average of 60% speedup by optimizing disk throughput only based on the disk sensitivity to fan speeds.

B. Energy Savings

We show the energy saving in Figure 10(b) comparing different management policies. Server E2 has a higher TDP, but the cooling system is under-provisioned, leading to higher core temperatures, need for DVFS throttling (and longer execution time), and thus smaller energy saving opportunities. Although E4 and E6 cooling subsystems are more capable, core temperatures are still high because of the workload characteristics and non-linear increase of cooling capability. JETCs goal of reducing fan speeds and overall energy consumption also leads to lower fan vibrations, indirectly boosting disk throughput. However, JETC balances the fans at the expense of an elevated average core temperature. Energy savings in JETC are not as high because it hits emergency throttling (again DVFS throttling) extending the run-time of an application, and thus consuming more energy. In the E6 system, because of controlled fan power (and reduce disk vibrations) policy, it improves energy savings by 1.8x over JETC, and 4.5x as compared to DTM. EADC achieves an average of 63% energy saving by reducing the peak and average core temperatures, *controlled* low fan speeds and minimizes job execution times through fan optimization.

VIII. CONCLUSION

Complex thermo-mechanical coupling in storage servers create undesired penalties that hinder a data center provider's ability to satisfy the end-user's QoS/SLA requirement. The temperature dependent leakage power, DVFS triggered low operating frequency, fan power, fan induced vibration, and disk throughput degradation – *all collectively* contribute negatively towards performance and energy. We model the impact of all these temperature dependent parameters and propose an Environment Aware Data Control (EADC) policy that locates an optimal operating point to minimize the passive energy consumption. EADC monitors the workload dependent core temperature and adjusts the fan speeds to balance the trade-off between waste energy (fan power, leakage, etc.) and performance (high disk throughput, high operation frequency, etc.).

In the future, we plan to use this work as a foundation to study component reliability with respect to thermal head room margin (THM). In another on-going work, we are studying

the negative effect of high ambient temperature T_{amb} (i.e. low THM) on the total energy consumed.

ACKNOWLEDGMENT

This research was supported by Oracle Labs and NSF grant CNS-1422921.

REFERENCES

- [1] R. Nachiappan, B. Javadi, R. N. Calheiros, and K. M. Matawie, "Cloud storage reliability for big data applications: A state of the art survey," *Journal of Network and Computer Applications*, vol. 97, pp. 35–47, 2017.
- [2] E. Brewer, L. Ying, L. Greenfield, R. Cypher, and T. T'so, "Disks for data centers," *Google@FAST'2016*, 2016.
- [3] K. Vaidyanathan and K. C. Gross, "Intelligent power monitoring and management for enterprise servers," in *Energy Efficient Electronic Systems (E3S), 2013 Third Berkeley Symposium on*. IEEE, 2013, pp. 1–2.
- [4] S. Sankar, M. Shaw, K. Vaid, and S. Gurusurthi, "Data-center scale evaluation of the impact of temperature on hard disk drive failures," *ACM Transactions on Storage (TOS)*, vol. 9, no. 2, p. 6, 2013.
- [5] C. S. Chan, A. S. Akyurek, B. Aksanli, and T. S. Rosing, "Optimal performance-aware cooling on enterprise servers," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Sep. 2019.
- [6] C. Chan, B. Pan, K. Gross, K. Vaidyanathan, and T. Rosing, "Correcting vibration-induced performance degradation in enterprise servers," *ACM SIGMETRICS Performance Evaluation Review*, pp. 83–88, 01 2014.
- [7] Z. Wang, C. Bash, N. Tolia, M. Marwah, X. Zhu, and P. Ranganathan, "Optimal fan speed control for thermal management of servers," in *ASME 2009 InterPACK Conf.* ASME Digital Collection, 2009, pp. 709–719.
- [8] R. Ayoub, R. Nath, and T. Rosing, "JETC: Joint energy thermal and cooling management for memory and cpu subsystems in servers," in *IEEE Intl. Symp. on High-Performance Comp Architecture*, 2012.
- [9] M. T. Chaudhry, T. C. Ling, and et.al., "Thermal-aware scheduling in green data centers," *ACM Computing Surveys (CSUR)*, vol. 47, no. 3, p. 39, 2015.
- [10] Z. Wang, X. Zhu, C. McCarthy, P. Ranganathan, and V. Talwar, "Feedback control algorithms for power management of servers," in *Third International Workshop on Feedback Control Implementation and Design in Computing Systems and Networks*, 2008.
- [11] A. Bartolini, M. Cacciari, A. Tilli, and L. Benini, "Thermal and energy management of high-performance multicores: Distributed and self-calibrating model-predictive controller," *IEEE TPDS*, pp. 170–183, Jan 2013.
- [12] J. Turner, "Effects of data center vibration on compute system performance," in *Proceedings of the First USENIX Conference on Sustainable Information Technology*, Berkeley, CA, USA, 2010, pp. 5–5.
- [13] D. Szabados, "Are All Hard Drives Created Equal? Examining Rotational Vibration in Desktop vs. Enterprise." 2010. [Online]. Available: <http://enterprise.media.seagate.com>
- [14] H. S. Gunawi, T. Do, P. Joshi, J. M. Hellerstein, A. C. Arpaci-Dusseau, R. H. Arpaci-Dusseau, and K. Sen, "Towards automatically checking thousands of failures with micro-specifications," in *HotDep*, 2010.
- [15] Z. Li, K. M. Greenan, A. W. Leung, and E. Zadok, "Power consumption in enterprise-scale backup storage systems," *Power*, vol. 2, no. 2, 2012.
- [16] I. Horison, "Implementing a Modern Backup Architecture," 2018. [Online]. Available: <https://www.oracle.com/us/products/servers-storage/storage/tape-storage/modern-backup-arch-1933021.pdf>
- [17] L. A. Barroso, J. Clidaras, and U. Hölzle, "The data-center as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 8, no. 3, pp. 1–154, 2013.
- [18] E. Pinheiro, W.-D. Weber, and L. A. Barroso, "Failure trends in a large disk drive population," in *5th USENIX FAST 07*, San Jose, CA, Feb. 2007.
- [19] O. Semenov, A. Vassighi, and M. Sachdev, "Impact of technology scaling on thermal behavior of leakage current in sub-quarter micron mosfets: perspective of low temperature current testing," *Microelectronics Journal*, vol. 33, no. 11, pp. 985–994, 2002.
- [20] A. K. Coskun, T. S. Rosing, and K. C. Gross, "Temperature management in multiprocessor socs using on-line learning," in *Design Automation Conference, 45th ACM/IEEE*, 2008, pp. 890–893.
- [21] HPE, "HP DL380 G5 Specifications," 2018. [Online]. Available: https://support.hpe.com/hpsc/doc/public/display?docId=emr_na-c01604169
- [22] Dell, "Dell MD1420 Specifications," 2018. [Online]. Available: <https://www.dell.com/en-us/work/shop/povw/storage-md1420>
- [23] Oracle, "Oracle ZFS Storage Systems," 2018. [Online]. Available: <https://www.oracle.com/storage/nas/zs7-2/>
- [24] R. Z. Ayoub and T. S. Rosing, "Predict and act: dynamic thermal management for multi-core processors," in *Proceedings of the 2009 ACM/IEEE Intl. Symp. on Low power electronics and design*, 2009.
- [25] J. Kim, M. M. Sabry, and et.al, "Global fan speed control considering non-ideal temperature measurements in enterprise servers," in *Proceedings of the conference on Design, Automation & Test in Europe*. European Design and Automation Association, 2014, p. 276.
- [26] J. Choi, C.-Y. Cher, H. Franke, H. Hamann, A. Weger, and P. Bose, "Thermal-aware task scheduling at the system software level," in *Proceedings of the 2007 Intl. Symp. on Low power electronics and design*. ACM, 2007, pp. 213–218.