



Environmental Conditions and Disk Reliability in Free-cooled Datacenters

Ioannis Manousakis, *Rutgers University*; Sriram Sankar, *GoDaddy*;
Gregg McKnight, *Microsoft*; Thu D. Nguyen, *Rutgers University*; Ricardo Bianchini, *Microsoft*

<https://www.usenix.org/conference/fast16/technical-sessions/presentation/manousakis>

This paper is included in the Proceedings of the
14th USENIX Conference on
File and Storage Technologies (FAST '16).

February 22–25, 2016 • Santa Clara, CA, USA

ISBN 978-1-931971-28-7

Open access to the Proceedings of the
14th USENIX Conference on
File and Storage Technologies
is sponsored by USENIX

Environmental Conditions and Disk Reliability in Free-Cooled Datacenters

Ioannis Manousakis,^{*‡} Sriram Sankar,[‡] Gregg McKnight,^δ Thu D. Nguyen,[†] Ricardo Bianchini^δ

[†]Rutgers University [‡]GoDaddy ^δMicrosoft

Abstract

Free cooling lowers datacenter costs significantly, but may also expose servers to higher and more variable temperatures and relative humidities. It is currently unclear whether these environmental conditions have a significant impact on hardware component reliability. Thus, in this paper, we use data from nine hyperscale datacenters to study the impact of environmental conditions on the reliability of server hardware, with a particular focus on disk drives and free cooling. Based on this study, we derive and validate a new model of disk lifetime as a function of environmental conditions. Furthermore, we quantify the tradeoffs between energy consumption, environmental conditions, component reliability, and datacenter costs. Finally, based on our analyses and model, we derive server and datacenter design lessons.

We draw many interesting observations, including (1) relative humidity seems to have a dominant impact on component failures; (2) disk failures increase significantly when operating at high relative humidity, due to controller/adaptor malfunction; and (3) though higher relative humidity increases component failures, software availability techniques can mask them and enable free-cooled operation, resulting in significantly lower infrastructure and energy costs that far outweigh the cost of the extra component failures.

1 Introduction

Datacenters consume a massive amount of energy. A recent study [18] estimates that they consume roughly 2% and 1.5% of the electricity in the United States and world-wide, respectively. In fact, a single hyperscale datacenter may consume more than 30MW [7].

These staggering numbers have prompted many efforts to reduce datacenter energy consumption. Perhaps the most successful of these efforts have involved reducing the energy consumption of the datacenter cooling infrastructure. In particular, three important techniques

have helped reduce the cooling energy: (1) increasing the hardware operating temperature to reduce the need for cool air inside the datacenter; (2) building datacenters where their cooling can directly leverage the outside air, reducing the need for energy-hungry (and expensive) water chillers; and (3) eliminating the hot air recirculation within the datacenter by isolating the cold air from the hot air. By using these and other techniques, large datacenter operators today can report yearly Power Usage Effectiveness (PUE) numbers in the 1.1 to 1.2 range, meaning that only 10% to 20% of the total energy goes into non-IT activities, including cooling. The low PUEs of these modern (“direct-evaporative-cooled” or simply “free-cooled”¹) datacenters are substantially lower than those of older generation datacenters [14, 15].

Although lowering cooling costs and PUEs would seem like a clear win, increasing the operating temperature and bringing the outside air into the datacenter may have unwanted consequences. Most intriguingly, these techniques may decrease hardware component reliability, as they expose the components to aggressive environmental conditions (e.g., higher temperature and/or higher relative humidity). A significant decrease in hardware reliability could actually increase rather than decrease the total cost of ownership (TCO).

Researchers have not yet addressed the tradeoffs between cooling energy, datacenter environmental conditions, hardware component reliability, and overall costs in modern free-cooled datacenters. For example, the prior work on the impact of environmental conditions on hardware reliability [10, 25, 27] has focused on older (non-free-cooled) datacenters that maintain lower and more stable temperature and relative humidity at each spatial spot in the datacenter. Because of their focus on these datacenters, researchers have not addressed the reliability impact of relative humidity in energy-efficient, free-cooled datacenters at all.

¹Throughout the paper, we refer to free cooling as the direct use of outside air to cool the servers. Some authors use a broader definition.

*This work was done while Ioannis was at Microsoft.

Understanding these tradeoffs and impacts is the topic of this paper. First, we use data collected from the operations of nine world-wide Microsoft datacenters for 1.5 years to 4 years to study the impact of environmental conditions (absolute temperature, temperature variation, relative humidity, and relative humidity variation) on the reliability of server hardware components. Based on this study and the dominance of disk failures, we then derive and validate a new model of disk lifetime as a function of both temperature and relative humidity. The model leverages data on the impact of relative humidity on corrosion rates. Next, we quantify the tradeoffs between energy consumption, environmental conditions, component reliability, and costs. Finally, based on our dataset and model, we derive server and datacenter design lessons.

We draw many observations from our dataset and analyses, including (1) disks account for the vast majority (89% on average) of the component failures regardless of the environmental conditions; (2) relative humidity seems to have a *much stronger* impact on disk failures than absolute temperature in current datacenter operating conditions, even when datacenters operate within ASHRAE’s “allowable” conditions [4] (i.e., 10-35°C inlet air temperature and 20–80% relative humidity); (3) temperature variation and relative humidity variation are *negatively* correlated with disk failures, but this is a consequence of these variations tending to be strongest when relative humidity is low; (4) disk failure rates increase significantly during periods of high relative humidity, i.e. these periods exhibit temporal clustering of failures; (5) disk controller/connectivity failures increase significantly when operating at high relative humidity (the controller and the adaptor are the only parts that are exposed to the ambient conditions); (6) in high relative humidity datacenters, server designs that place disks in the back of enclosures can reduce the disk failure rate significantly; and (7) though higher relative humidity increases component failures, relying on software techniques to mask them and operate in this mode also significantly reduces infrastructure and energy costs, and *more than compensates* for the cost of the additional failures.

Note that, unlike disk vendors, we do not have access to a large isolated chamber where thousands of disks can be exposed to different environmental conditions in a controlled and repeatable manner.² Instead, we derive the above observations from multiple statistical analyses of large-scale commercial datacenters under their real operating conditions, as in prior works [10, 25, 27]. To increase confidence in our analyses and inferences, we

²From these experiments, vendors derive recommendations for the ideal operating conditions for their parts. Unfortunately, it is very difficult in practice to guarantee consistent operation within those conditions, as server layouts vary and the environment inside servers is difficult to control exactly, especially in free-cooled datacenters.

personally inspected two free-cooled datacenters that experience humid environments and observed many samples of corroded parts.

In summary, we make the following contributions:

- We study the impact of relative humidity and relative humidity variation on hardware reliability (with a strong focus on disk drives) in datacenters.
- We study the tradeoffs between cooling energy, environmental conditions, hardware reliability, and cost in datacenters.
- Using data from nine datacenters, more than 1M disks, and 1.5–4 years, we draw many interesting observations. Our data suggests that the impact of temperature and temperature variation on disk reliability (the focus of the prior works) is *much less* significant than that of relative humidity in modern cooling setups.
- Using our disk data and a corrosion model, we also derive and validate a new model of disk lifetime as a function of environment conditions.
- From our observations and disk lifetime model, we draw a few server and datacenter design lessons.

2 Related Work

Environmentals and their impact on reliability. Several works have considered the impact of the cooling infrastructure on datacenter temperatures and humidities, e.g. [3, 22, 23]. However, they did not address the hardware reliability implications of these environmental conditions. The reason is that meaningful reliability studies require large server populations in datacenters that are monitored for multiple years. Our paper presents the largest study of these issues to date.

Other authors have had access to such large real datasets for long periods of time: [5, 10, 20, 25, 27, 28, 29, 35]. A few of these works [5, 28, 29, 35] considered the impact of age and other factors on hardware reliability, but did not address environmental conditions and their potential effects. The other prior works [10, 25, 27] have considered the impact of absolute temperature and temperature variation on the reliability of hardware components (with a significant emphasis on disk drives) in datacenters with fairly stable temperature and relative humidity at each spatial spot, i.e. non-air-cooled datacenters. Unfortunately, these prior works are inconclusive when it comes to the impact of absolute temperature and temperature variations on hardware reliability. Specifically, El-Sayed *et al.* [10] and Pinheiro *et al.* [25] found a smaller impact of absolute temperature on disk lifetime than previously expected, whereas Sankar *et al.* [27] found a significant impact. El-Sayed *et al.* also found temperature variations to have a more significant impact than absolute temperature on Latent Sector Errors

(LSEs), a common type of disk failure that renders sectors inaccessible. None of the prior studies considered relative humidity or relative humidity variations.

Our paper adds to the debate about the impact of absolute temperature and temperature variation. However, our paper suggests that this debate may actually be moot in modern (air-cooled) datacenters. In particular, our results show that relative humidity is a more significant factor than temperature. For this reason, we also extend an existing disk lifetime model with a relative humidity term, and validate it against our real disk reliability data.

Other tradeoffs. Prior works have also considered the impact of the cooling infrastructure and workload placement on cooling energy and costs, e.g. [1, 6, 8, 17, 19, 21, 23]. However, they did not address the impact of environmental conditions on hardware reliability (and the associated replacement costs). A more complete understanding of these tradeoffs requires a more comprehensive study, like the one we present in this paper. Specifically, we investigate a broader spectrum of tradeoffs, including cooling energy, environmental conditions, hardware reliability, and costs. Importantly, we show that the increased hardware replacement cost in free-cooled datacenters is *far outweighed* by their infrastructure and operating costs savings.

However, we do not address effects that our dataset does not capture. In particular, techniques devised to reduce cooling energy (increasing operating temperature and using outside air) may increase the energy consumption of the IT equipment, if server fans react by spinning faster. They may also reduce performance, if servers throttle their speed as a result of the higher operating temperature. Prior research [10, 32] has considered these effects, and found that the cooling energy benefits of these techniques outweigh the downsides.

3 Background

Datacenter cooling and environmentals. The cooling infrastructure of hyperscale datacenters has evolved over time. The first datacenters used water chillers with computer room air handlers (CRAHs). CRAHs do not feature the integrated compressors of traditional computer room air conditioners (CRACs). Rather, they circulate the air carrying heat from the servers to cooling coils carrying chilled water. The heat is then transferred via the water back to the chillers, which transfer the heat to another water loop directed to a cooling tower, before returning the chilled water back inside the datacenter. The cooling tower helps some of the water to evaporate (dissipating heat), before it loops back to the chillers. Chillers are expensive and consume a large amount of energy. However, the environmental conditions inside the datacenter

Technology	Temp/RH Control	CAPEX	PUE
Chillers	Precise / Precise	\$2.5/W	1.7
Water-side	Precise / Precise	\$2.8/W	1.19
Free-cooled	Medium / Low	\$0.7/W	1.12

Table 1: Typical temperature and humidity control, CAPEX [11], and PUEs of the cooling types [13, 34].

can be precisely controlled (except for hot spots that may develop due to poor air flow design). Moreover, these datacenters do not mix outside and inside air. We refer to these datacenters as *chiller-based*.

An improvement over this setup allows the chillers to be bypassed (and turned off) when the cooling towers alone are sufficient to cool the water. Turning the chillers off significantly reduces energy consumption. When the cooling towers cannot lower the temperature enough, the chillers come back on. These datacenters tightly control the internal temperature and relative humidity, like their chiller-based counterparts. Likewise, there is still no mixing of outside and inside air. We refer to these datacenters as *water-side economized*.

A more recent advance has been to use large fans to blow cool outside air into the datacenter, while filtering out dust and other air pollutants. Again using fans, the warm return air is guided back out of the datacenter. When the outside temperature is high, these datacenters apply an evaporative cooling process that adds water vapor into the airstream to lower the temperature of the outside air, before letting it reach the servers. To increase temperature (during excessively cold periods) and/or reduce relative humidity, these datacenters intentionally recirculate some of the warm return air. This type of control is crucial because rapid reductions in temperature (more than 20°C per hour, according to ASHRAE [4]) may cause condensation inside the datacenter. This cooling setup enables the forgoing of chillers and cooling towers altogether, thus is the cheapest to build. However, these datacenters may also expose the servers to warmer and more variable temperatures, and higher and more variable relative humidities than other datacenter types. We refer to these datacenters as direct-evaporative-cooled or simply *free-cooled*.

A survey of the popularity of these cooling infrastructures can be found in [16]. Table 1 summarizes the main characteristics of the cooling infrastructures in terms of their ability to control temperature and relative humidity, and their estimated cooling infrastructure costs [11] and PUEs. For the PUE estimates, we assume Uptime Institute’s surveyed average PUE of 1.7 [34] for chiller-based cooling. For the water-side economization PUE, we assume that the chiller only needs to be active 12.5% of the year, i.e. during the day time in the summer. The PUE of free-cooled datacenters depends on their locations, but we assume a single value (1.12) for simplicity. This value

is in line with those reported by hyperscale datacenter operators. For example, Facebook’s free-cooled datacenter in Prineville, Oregon reports an yearly average PUE of 1.08 with peaks around 1.14 [13]. All PUE estimates assume 4% overheads due to factors other than cooling.

Hardware lifetime models. Many prior reliability models associated component lifetime with temperature. For example, [30] considered several CPU failure modes that result from high temperatures. CPU manufacturers use high temperatures and voltages to accelerate the onset of early-life failures [30]. Disk and other electronics vendors do the same to estimate mean times to failure (mean lifetimes). The Arrhenius model is often used to calculate an *acceleration factor* (AF_T) for the lifetime [9].

$$AF_T = e^{\frac{E_a}{k} \cdot (\frac{1}{T_b} - \frac{1}{T_e})} \quad (1)$$

where E_a is the activation energy (in eV) for the device, k is Boltzmann’s constant ($8.62 \cdot 10^{-5}$ eV/K), T_b is the average baseline operating temperature (in K) of the device, and T_e is the average elevated temperature (in K).

The acceleration factor can be used to estimate how much higher the failure rate will be during a certain period. For example, if the failure rate is typically 2% over a year (i.e., 2% of the devices fail in a year) at a baseline temperature, and the acceleration factor is 2 at a higher temperature, the estimate for the accelerated rate will be 4% ($2\% \times 2$) for the year. In other words, $\overline{FR}_T = AF_T \times \overline{FR}_{T_b}$, where \overline{FR}_T is the average failure rate due to elevated temperature, and \overline{FR}_{T_b} is the average failure rate at the baseline temperature. Prior works [10, 27] have found the Arrhenius model to approximate disk failure rates accurately, though El-Sayed *et al.* [10] also found accurate linear fits to their failure data.

The Arrhenius model computes the acceleration factor assuming steady-state operation. To extrapolate the model to periods of changing temperature, existing models compute a weighted acceleration factor, where the weights are proportional to the length of the temperature excursions [31]. We take this approach when proposing our extension of the model to relative humidity and free-cooled datacenters. Our validation of the extended model (Section 6) shows very good accuracy for our dataset.

4 Methodology

In this section, we describe the main characteristics of our dataset and the analyses it enables. We purposely omit certain sensitive information about the datacenters, such as their locations, numbers of servers, and hardware vendors, due to commercial and contractual reasons. Nevertheless, the data we do present is plenty to make our points, as shall become clear in later sections.

DC Tag	Cooling	Months	Refresh Cycles	Disk Popul.
CD1	Chiller	48	2	117 K
CD2	Water-Side	48	2	146 K
CD3	Free-Cooled	27	1	24 K
HD1	Chiller	24	1	16 K
HD2	Water-Side	48	2	100 K
HH1	Free-Cooled	24	1	168 K
HH2	Free-Cooled	22	1	213 K
HH3	Free-Cooled	24	1	124 K
HH4	Free-Cooled	18	1	161 K
Total				1.07 M

Table 2: Main datacenter characteristics. The “C” and “D” tags mean cool and dry. An “H” as the first letter of the tag means hot, whereas an “H” as the second letter means humid.

Data sources. We collect data from nine hyperscale Microsoft datacenters spread around the world for periods from 1.5 to 4 years. The data includes component health and failure reports, traces of environmental conditions, traces of component utilizations, cooling energy data, and asset information.

The datacenters use a variety of cooling infrastructures, exhibiting different environmental conditions, hardware component reliabilities, energy efficiencies, and costs. The three first columns from the left of Table 2 show each datacenter’s tag, its cooling technology, and the length of data we have for it. The tags correspond to the environmental conditions inside the datacenters (see caption for details), not their cooling technology or location. We classify a datacenter as “hot” (“H” as the first letter of its tag) if at least 10% of its internal temperatures over a year are above 24°C , whereas we classify it as “humid” (“H” as the second letter of the tag) if at least 10% of its internal relative humidities over a year are above 60%. We classify a datacenter that is not “hot” as “cool”, and one that is not “humid” as “dry”. Although admittedly arbitrary, our naming convention and thresholds reflect the general environmental conditions in the datacenters accurately. For example, HD1 (hot and dry) is a state-of-the-art chiller-based datacenter that precisely controls temperature at a high setpoint. More interestingly, CD3 (cool and dry) is a free-cooled datacenter so its internal temperatures and relative humidities vary more than in chiller-based datacenters. However, because it is located in a cold region, the temperatures and relative humidities can be kept fairly low the vast majority of the time.

To study hardware component reliability, we gather failure data for CPUs, memory modules (DIMMs), power supply units (PSUs), and hard disk drives. The two rightmost columns of Table 2 list the number of disks we consider from each datacenter, and the number of “refresh” cycles (servers are replaced every 3 years)

in each datacenter. We filter the failure data for entries with the following properties: (1) the entry was classified with a maintenance tag; (2) the component is in either *Failing* or *Dead* state according to the datacenter-wide health monitoring system; and (3) the entry’s error message names the failing component. For example, a disk error will generate either a SMART (Self-Monitoring, Analysis, and Reporting Technology) report or a failure to detect the disk on its known SATA port. The nature of the error allows further classification of the underlying failure mode.

Defining exactly when a component has failed permanently is challenging in large datacenters [25, 28]. However, since many components (most importantly, hard disks) exhibit recurring errors before failing permanently and we do not want to double-count failures, we consider a component to have failed on the first failure reported to the datacenter-wide health monitoring system. This failure triggers manual intervention from a datacenter technician. After this first failure and manual repair, we count no other failure against the component. For example, we consider a disk to have failed on the first LSE that gets reported to the health system and requires manual intervention; this report occurs after the disk controller itself has already silently reallocated many sectors (e.g., 2000+ sectors for many disks in our dataset). Though this failure counting may seem aggressive at first blush (a component may survive a failure report and manual repair), note that others [20, 25] have shown high correlations of several types of SMART errors, like LSEs, with permanent failures. Moreover, the disk Annualized Failure Rates (AFRs) that we observe for chiller-based datacenters are in line with previous works [25, 28].

Detailed analyses. To correlate the disk failures with their environmental conditions, we use detailed data from one of the hot and humid datacenters (HH1). The data includes server inlet air temperature and relative humidity values, as well as outside air conditions with a granularity of 15 minutes. The dataset does not contain the temperature of all the individual components inside each server, or the relative humidity inside each box. However, we can accurately use the inlet values as the environmental conditions at the disks, because the disks are placed at the front of the servers (right at their air inlets) in HH1. For certain analyses, we use CD3 and HD1 as bases for comparison against HH1. Although we do not have information on the disks’ manufacturing batch, our cross-datacenter comparisons focus on disks that differ mainly in their environmental conditions.

To investigate potential links between the components’ utilizations and their failures, we collect historical average utilizations for the processors and disks in a granularity of 2 hours, and then aggregate them into lifetime average utilization for each component.

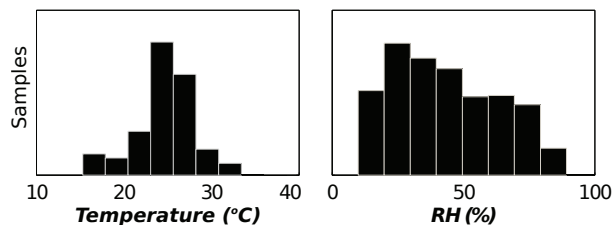


Figure 1: HH1 temperature and relative humidity distributions. Both of these environmentals vary widely.

We built a tool to process all these failure, environmental, and utilization data. After collecting, filtering, and deduplicating the data, the tool computes the AFRs and timelines for the component failures. With the timelines, it also computes daily and monthly failure rates. For disks, the tool also breaks the failure data across models and server configurations, and does disk error classification. Finally, the tool derives linear and exponential reliability models (via curve fitting) as a function of environmental conditions, and checks their accuracy versus the observed failure rates.

5 Results and Analyses

In this section, we first characterize the temperatures and relative humidities in a free-cooled datacenter, and the hardware component failures in the nine datacenters. We then perform a detailed study of the impact of environmental conditions on the reliability of disks. We close the section with an analysis of the hardware reliability, cooling energy, and cost tradeoffs.

5.1 Environmentals in free-cooled DCs

Chiller-based datacenters precisely control temperature and relative humidity, and keep them stable at each spatial spot in the datacenter. For example, HD1 exhibits a stable 27°C temperature and a 50% average relative humidity. In contrast, the temperature and relative humidity at each spatial spot in the datacenter vary under free cooling. For example, Figure 1 shows the temperature and relative humidity distributions measured at a spatial spot in HH1. The average temperature is 24.5°C (the standard deviation is 3.2°C) and the average relative humidity is 43% (the standard deviation is 20.3%). Clearly, HH1 exhibits wide ranges, including a large fraction of high temperatures (greater than 24°C) and a large fraction of high relative humidities (greater than 60%).

5.2 Hardware component failures

In light of the above differences in environmental conditions, an important question is whether hardware component failures are distributed differently in different types

DC Tag	Cooling	AFR	Increase wrt AFR = 1.5%
CD1	Chiller	1.5%	0%
CD2	Water-Side	2.1%	40%
CD3	Free-Cooled	1.8%	20%
HD1	Chiller	2.0%	33%
HD2	Water-Side	2.3%	53%
HH1	Free-Cooled	3.1%	107%
HH2	Free-Cooled	5.1%	240%
HH3	Free-Cooled	5.1%	240%
HH4	Free-Cooled	5.4%	260%

Table 3: Disk AFRs. HH1-HH4 incur the highest rates.

of datacenters. We omit the full results due to space limitations, but highlight that disk drive failures dominate with 76%–95% (89% on average) of all hardware component failures, regardless of the component models, environmental conditions, and cooling technologies. As an example, disks, DIMMs, CPUs, and PSUs correspond to 83%, 10%, 5%, and 2% of the total failures, respectively, in HH1. The other datacenters exhibit a similar pattern. Disks also dominate in terms of failure rates, with AFRs ranging from 1.5% to 5.4% (average 3.16%) in our dataset. In comparison, the AFRs of DIMMs, CPUs, and PSUs were 0.17%, 0.23%, and 0.59%, respectively. Prior works had shown that disk failures are the most common for stable environmental conditions, e.g. [27]. Our data shows that they also dominate in modern, hotter and more humid datacenters. Interestingly, as we discuss in Section 7, the placement of the components (e.g., disks) inside each server affects their failure rates, since the temperature and relative humidity vary as air flows through the server.

Given the dominance of disk failures and rates, we focus on them in the remainder of the paper.

5.3 Impact of environmentals

Disk failure rates. Table 3 presents the disk AFRs for the datacenters we study, and how much they differ relative to the AFR of one of the datacenters with stable temperature and relative humidity at each spatial spot (CD1). We repeat the cooling technology information from Table 2 for clarity. The data includes a small number of disk models in each datacenter. For example, CD3 and HD1 have two disk models, whereas HH1 has two disk models that account for 95% of its disks. More importantly, the most popular model in CD3 (55% of the total) and HD1 (85%) are the same. The most popular model in HH1 (82%) is from the same disk manufacturer and series as the most popular model in CD3 and HD1, and has the same rotational speed, bus interface, and form factor; the only differences between the models are their storage and cache capacities. In more detailed studies below,

we compare the impact of environmentals on these two models directly.

We make several observations from these data:

1. The datacenters with consistently or frequently dry internal environments exhibit the lowest AFRs, regardless of cooling technologies. For example, CD1 and HD1 keep relative humidities stable at 50%.
2. High internal relative humidity increases AFRs by 107% (HH1) to 260% (HH4), compared to CD1. Compared to HD1, the increases range from 55% to 170%. For example, HH1 exhibits a wide range of relative humidities, with a large percentage of them higher than 50% (Figure 1).
3. Free cooling does not necessarily lead to high AFR, as CD3 shows. Depending on the local climate (and with careful humidity control), free-cooled datacenters can have AFRs as low as those of chiller-based and water-side economized datacenters.
4. High internal temperature does not directly correlate to the high range of AFRs (greater than 3%), as suggested by datacenters HD1 and HD2.

The first two of these observations are indications that relative humidity may have a significant impact on disk failures. We cannot state a stronger result based solely on Table 3, because there are many differences between the datacenters, their servers and environmental conditions. Thus, in the next few pages, we provide more detailed evidence that consistently points in the same direction.

Causes of disk failures. The first question then becomes why would relative humidity affect disks if they are encapsulated in a sealed package? Classifying the disk failures in terms of their causes provides insights into this question. To perform the classification, we divide the failures into three categories [2]: (1) mechanical (pre-fail) issues; (2) age-related issues; and (3) controller and connectivity issues. In Table 4, we list the most common errors in our dataset. Pre-fail and old-age errors are reported by SMART. In contrast, IOCTL ATA PASS THROUGH (inability to issue an ioctl command to the controller) and SMART RCV DRIVE DATA (inability to read the SMART data from the controller) are generated in the event of an unresponsive disk controller.

In Figure 2, we present the failure breakdown for the popular disk model in HD1 (top), CD3 (middle), and HH1 (bottom). We observe that 67% of disk failures are associated with SMART errors in HD1. The vast majority (65%) of these errors are pre-fail, while just 2% are old-age errors. The remaining 33% correspond to controller and connectivity issues. CD3 also exhibits a substantially smaller percentage (42%) of controller and connectivity errors than SMART errors (58%). In contrast, HH1 experiences a much higher fraction of controller and connectivity errors (66%). Given that HH1 runs its servers cooler than HD1 most of the time, its two-fold increase in

Error Name	Type
IOCTL ATA PASS THROUGH	Controller/Connectivity
SMART RCV DRIVE DATA	Controller/Connectivity
Raw_Read_Error_Rate	Pre-fail
Spin_Up_Time	Pre-fail
Start_Stop_Count	Old age
Reallocated_Sectors_Count	Pre-fail
Seek_Error_Rate	Pre-fail
Power_On_Hours	Old age
Spin_Retry_Count	Pre-fail
Power_Cycle_Count	Old age
Runtime_Bad_Block	Old age
End-to-End_Error	Old age
Airflow_Temperature	Old age
G-Sense_Error_Rate	Old age

Table 4: Controller/connectivity and SMART errors [2].

controller and connectivity errors (66% vs 33% in HD1) seems to result from its higher relative humidity. To understand the reason for this effect, consider the physical design of the disks. The mechanical parts are sealed, but the disk controller and the disk adaptor are directly exposed to the ambient, allowing possible condensation and corrosion agents to damage them.

The key observation here is:

5. High relative humidity seems to increase the incidence of disk controller and connectivity errors, as the controller board and the disk adaptor are exposed to condensation and corrosion effects.

Temporal disk failure clustering. Several of the above observations point to high relative humidity as an important contributor to disk failures. Next, we present even more striking evidence of this effect by considering the temporal clustering of failures of disks of the same characteristics and age.

Figure 3 presents the number of daily disk failures at HD1 (in red) and HH1 (in black) for the same two-year span. We normalize the failure numbers to the size of HD1’s disk population to account for the large difference in population sizes. The green-to-red gradient band across the graph shows the temperature and humidity within HH1. The figure shows significant temporal clustering in the summer of 2013, when relative humidity at HH1 was frequently very high. Technicians found corrosion on the failed disks. The vast majority of the clustered HH1 failures were from disks of the same popular model; the disks had been installed within a few months of each other, 1 year earlier. Moreover, the figure shows increased numbers of daily failures in HH1 after the summer of 2013, compared to before it.

In more detail, we find that the HD1 failures were roughly evenly distributed with occasional spikes. The exact cause of the spikes is unclear. In contrast, HH1 shows a slightly lower failure rate in the first 12 months

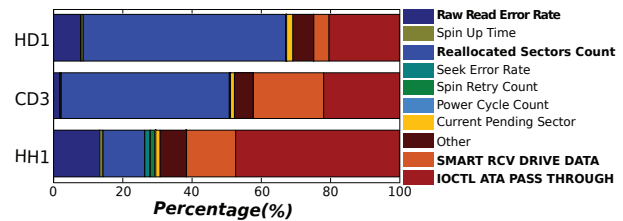


Figure 2: Failure classification in HD1 (top), CD3 (middle), and HH1 (bottom). Controller/connectivity failures in HH1 are double compared to HD1.

of its life, followed by a 3-month period of approximately 4x–5x higher daily failure rate. These months correspond to the summer, when the outside air temperature at HH1’s location is often high. When the outside air is hot but dry, HH1 adds moisture to the incoming air to lower its temperature via evaporation. When the outside air is hot and humid, this humidity enters the datacenter directly, increasing relative humidity, as the datacenter’s ability to recirculate heat is limited by the already high air temperature. As a result, relative humidity may be high frequently during hot periods.

We make three observations from these HH1 results:

6. High relative humidity may cause significant temporal clustering of disk failures, with potential consequences in terms of the needed frequency of manual maintenance/replacement and automatic data durability repairs during these periods.
7. The temporal clustering occurred in the *second* summer of the disks’ lifetimes, suggesting that they did not fail simply because they were first exposed to high relative humidity. Rather, this temporal behavior suggests a disk lifetime model where high relative humidity excursions consume lifetime at a rate corresponding to their duration and magnitude. The amount of lifetime consumption during the summer of 2012 was not enough to cause an increased rate of failures; the additional consumption during the summer of 2013 was. This makes intuitive sense: it takes time (at a corrosion rate we model) for relative humidity (and temperature) to produce enough corrosion to cause hardware misbehavior. Our relative humidity modeling in Section 6 embodies the notion of lifetime consumption and matches our data well. A similar lifetime model has been proposed for high disk temperature [31].
8. The increased daily failures after the second summer provide extra evidence for the lifetime consumption model and the long-term effect of high relative humidity.

Correlating disk failures and environmentals. So far, our observations have listed several indications that relative humidity is a key disk reliability factor. However, one of our key goals is to determine the relative impact of

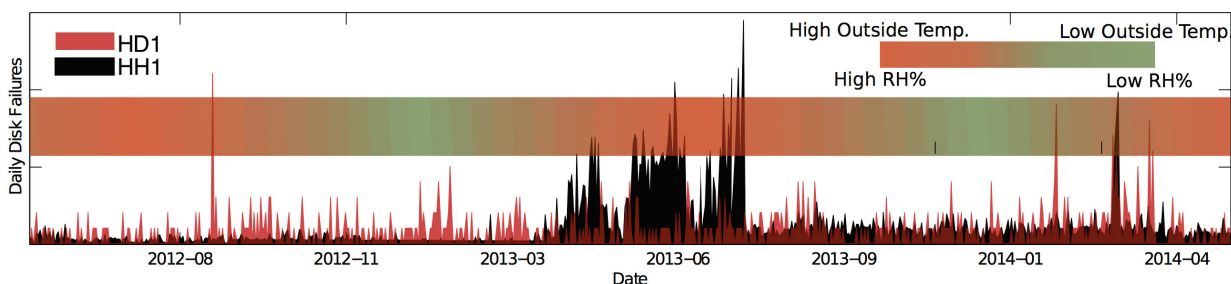


Figure 3: Two-year comparison between HD1 and HH1 daily failures. The horizontal band shows the outside temperature and relative humidity at HH1. HH1 experienced temporally clustered failures during summer’13, when the relative humidities were high. We normalized the data to the HD1 failures due to the different numbers of disks in these datacenters.

the four related environmental conditions: absolute temperature, relative humidity, temperature variation, and relative humidity variation.

In Table 5, we present both linear and exponential fits of the monthly failure rates for HH1, as a function of our four environmental condition metrics. For example, when seeking a fit for relative humidity, we use x = the average relative humidity experienced by the failed disks during a month, and y = the failure rate for that month. Besides the parameter fits, we also present R^2 as an estimate of the fit error. The exponential fit corresponds to a model similar to the Arrhenius equation (Equation 1). We use the disks’ temperature and relative humidity Coefficient of Variation (CoV) to represent temporal temperature and relative humidity variations. (The advantage of the CoV over the variance or the standard deviation is that it is normalized by the average.) El-Sayed *et al.* [10] studied the same types of fits for their data, and also used CoVs to represent variations. Note that we split HH1’s disk population into four groups (P1–P4) roughly corresponding to the spatial location of the disks in the datacenter; this accounts for potential environmental differences across cold aisles.

We can see that the relative humidity consistently exhibits positive correlations with failure rates, by considering the a parameter in the linear fit and the b parameter in the exponential fit. This positive correlation means that the failure rate increases when the relative humidity increases. In contrast, the other environmental conditions either show some or almost all negative parameters, suggesting much weaker correlations.

Interestingly, the temperature CoVs and the relative humidity CoVs suggest mostly *negative* correlations with the failure rate. This is antithetical to the physical stresses that these variations would be expected to produce. To explain this effect most clearly, we plot two correlation matrices in Figure 4 for the most popular disk model in HH1. The figure illustrates the correlations between all environmental condition metrics and

the monthly failure rate. To derive these correlations, we use (1) the Pearson product-moment correlation coefficient, as a measure of linear dependence; and (2) Spearman’s rank correlation coefficient, as a measure of non-linear dependence. The color scheme indicates the correlation between each pair of metrics. The diagonal elements have a correlation equal to one as they represent dependence between the same metric. As before, a positive correlation indicates an analogous relationship, whereas a negative correlation indicates an inverse relationship. These matrices show that the CoVs are *strongly negatively* correlated (correlations close to -1) with average relative humidity. This implies that the periods with higher temperature and relative humidity variations are also the periods with low relative humidity, providing extended lifetime.

For completeness, we also considered the impact of average disk utilization on lifetime, but found no correlation. This is consistent with prior work, e.g. [25].

From the figures and table above, we observe that:

9. Average relative humidity seems to have the strongest impact on disk lifetime of all the environmental condition metrics we consider.
10. Average temperature seems to have a substantially lower (but still non-negligible) impact than average relative humidity on disk lifetime.
11. Our dataset includes no evidence that temperature variation or relative humidity variation has an effect on disk reliability in free-cooled datacenters.

5.4 Trading off reliability, energy, and cost

The previous section indicates that higher relative humidity appears to produce shorter lifetimes and, consequently, higher equipment costs (maintenance/repair costs tend to be small compared to the other TCO factors, so we do not consider them. However, the set of tradeoffs is broader. Higher humidity results from free cooling in certain geographies, but this type of cooling

Popul.	%	Linear Fit $a \cdot x + b$											
		Temperature			RH			CoV - Temperature			CoV - RH		
		a	b	R^2	a	b	R^2	a	b	R^2	a	b	R^2
P1	30.1	$5.17 \cdot 10^{-5}$	$-2.43 \cdot 10^{-3}$	0.81	$1.20 \cdot 10^{-4}$	$-4.88 \cdot 10^{-3}$	0.83	$-7.90 \cdot 10^{-3}$	$1.14 \cdot 10^{-3}$	0.83	$-6.56 \cdot 10^{-3}$	$3.29 \cdot 10^{-3}$	0.84
P2	25.6	$-1.91 \cdot 10^{-5}$	$2.46 \cdot 10^{-3}$	0.83	$1.03 \cdot 10^{-4}$	$-1.73 \cdot 10^{-3}$	0.84	$-9.06 \cdot 10^{-3}$	$1.28 \cdot 10^{-3}$	0.84	$-3.71 \cdot 10^{-3}$	$1.98 \cdot 10^{-3}$	0.83
P3	23.3	$1.41 \cdot 10^{-3}$	$-1.04 \cdot 10^{-1}$	0.75	$2.11 \cdot 10^{-4}$	$-5.59 \cdot 10^{-3}$	0.71	$-4.91 \cdot 10^{-2}$	$7.26 \cdot 10^{-2}$	0.77	$-4.46 \cdot 10^{-2}$	$2.42 \cdot 10^{-2}$	0.78
P4	19.6	$1.73 \cdot 10^{-3}$	$-1.07 \cdot 10^{-1}$	0.36	$4.45 \cdot 10^{-4}$	$-16.4 \cdot 10^{-3}$	0.44	$-1.36 \cdot 10^{-1}$	$1.33 \cdot 10^{-2}$	0.47	$-8.02 \cdot 10^{-2}$	$4.13 \cdot 10^{-2}$	0.55

Popul.	%	Exponential Fit $a \cdot e^{b \cdot x}$											
		Temperature			RH			CoV - Temperature			CoV - RH		
		a	b	R^2	a	b	R^2	a	b	R^2	a	b	R^2
P1	30.1	$2.38 \cdot 10^{-3}$	$-6.45 \cdot 10^{-3}$	0.84	$2.67 \cdot 10^{-4}$	$5.11 \cdot 10^{-2}$	0.89	$2.06 \cdot 10^{-3}$	$-1.74 \cdot 10^0$	0.72	$1.74 \cdot 10^{-2}$	$-9.01 \cdot 10^0$	0.81
P2	25.6	$3.08 \cdot 10^{-3}$	$-1.64 \cdot 10^{-2}$	0.85	$5.37 \cdot 10^{-4}$	$5.38 \cdot 10^{-2}$	0.88	$1.84 \cdot 10^{-3}$	$-1.09 \cdot 10^1$	0.79	$1.37 \cdot 10^{-2}$	$-1.63 \cdot 10^1$	0.81
P3	23.3	$2.57 \cdot 10^{-3}$	$6.35 \cdot 10^{-3}$	0.76	$5.31 \cdot 10^{-5}$	$9.93 \cdot 10^{-2}$	0.69	$5.13 \cdot 10^{-3}$	$-8.45 \cdot 10^0$	0.58	$1.66 \cdot 10^{-3}$	$3.57 \cdot 10^0$	0.57
P4	19.6	$3.62 \cdot 10^{-4}$	$3.36 \cdot 10^{-2}$	0.43	$1.31 \cdot 10^{-5}$	$11.54 \cdot 10^{-2}$	0.59	$6.10 \cdot 10^{-3}$	$-3.91 \cdot 10^0$	0.23	$7.17 \cdot 10^{-3}$	$-1.39 \cdot 10^0$	0.21

Table 5: Linear ($a \cdot x + b$) and nonlinear ($a \cdot e^{b \cdot x}$) fits for the monthly failure rates of four disk populations in HH1, as a function of the absolute temperature, relative humidity, temperature variation, and relative humidity variation.

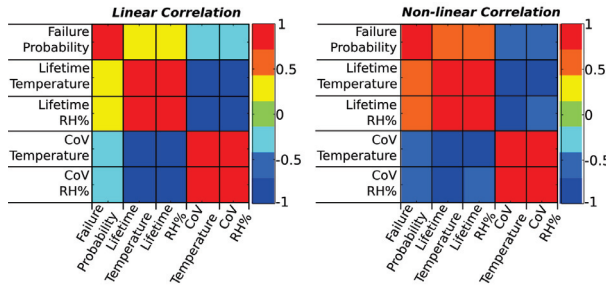


Figure 4: Linear and non-linear correlation between the monthly failure rate and the environmental conditions we consider, for the most popular disk model in HH1. red = high correlation, dark blue = low correlation.

also improves energy efficiency and lowers both capital and operating costs.

Putting all these effects together, in Figure 5 we compare the cooling-related (capital + operational + disk replacement) costs for a chiller-based (CD1), a water-side economized (HD2), and a free-cooled datacenter (HH4). For the disk replacement costs, we consider only the cost above that of a 1.5% AFR (the AFR of CD1), so the chiller-based bars do not include these costs. We use HD2 and HH4 datacenters for this figure because they exhibit the highest disk AFRs in their categories; they represent the worst-case disk replacement costs for their respective datacenter classes. We use the PUEs and CAPEX estimates from Table 1, and estimate the cooling energy cost assuming \$0.07/kWh for the electricity price (the average industrial electricity price in the United States). We also assume that the datacenter operator brunts the cost of replacing each failed disk, which we assume to be \$100, by having to buy the extra disk (as opposed to simply paying a slightly more expensive warranty). We perform the cost comparison for 10, 15, and 20 years, as the datacenter lifetime.

The figure shows that, for a lifetime of 10 years, the cooling cost of the chiller-based datacenter is roughly balanced between capital and operating expenses. For a lifetime of 15 years, the operational cooling cost becomes the larger fraction, whereas for 20 years it be-

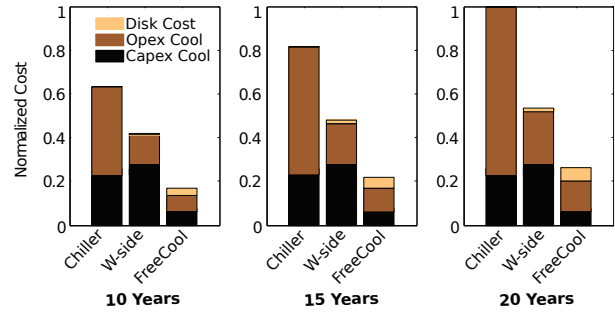


Figure 5: 10, 15, and 20-year cost comparison for chiller-based, water-side, and free-cooled datacenters including the cost replacing disks.

comes roughly 75% of the total cooling-related costs. In comparison, water-side economized datacenters have slightly higher capital cooling costs than chiller-based ones. However, their operational cooling costs are substantially lower, because of the periods when the chillers can be bypassed (and turned off). Though these datacenters may lead to slightly higher AFRs, the savings from lower operational cooling costs more than compensate for the slight disk replacement costs. In fact, the fully burdened cost of replacing a disk would have to be several fold higher than \$100 (which is unlikely) for replacement costs to dominate. Finally, the free-cooled datacenter exhibits lower capital cooling costs (by 3.6x) and operational cooling costs (by 8.3x) than the chiller-based datacenter. Because of these datacenters' sometimes higher AFRs, their disk replacement costs may be non-trivial, but the overall cost tradeoff is still clearly in their favor. For 20 years, the free-cooled datacenter exhibits overall cooling-related costs that are roughly 74% and 45% lower than the chiller-based and water-side economized datacenters, respectively.

Based on this figure, we observe that:

- Although operating at higher humidity may entail substantially higher component AFRs in certain geographies, the cost of this increase is small compared to the savings from reducing energy consumption and infrastructure costs via free cooling, especially for longer lifetimes.

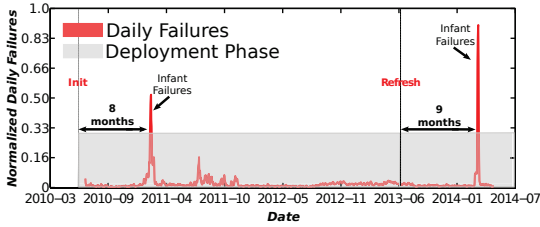


Figure 6: Normalized daily failures in CD1 over 4 years. Infant mortality seems clear.

6 Modeling Lifetime in Modern DCs

Models of hardware component lifetime have been used extensively in industry and academia to understand and predict failures, e.g. [9, 29, 30]. In the context of a datacenter, which hosts hundreds of thousands of components, modeling lifetimes is important to select the conditions under which datacenters should operate and prepare for the unavoidable failures. Because disk drives fail much more frequently than other components, they typically receive the most modeling attention. The two most commonly used disk lifetime models are the *bath-tub model* and the *failure acceleration model based on the Arrhenius equation* (Equation 1). The latter acceleration model does not consider the impact of relative humidity, which has only become a significant factor after the advent of free cooling (Section 5.3).

Bath-tub model. This model abstracts the lifetime of large component populations as three consecutive phases: (1) in the first phase, a significant number of components fail prematurely (*infant mortality*); (2) the second phase is a stable period of lower failure rate; and (3) in the third phase, component failure increase again due to wear-out effects. We clearly observe infant mortality in some datacenters but not others, depending on disk model/vendor. For example, Figure 3 does not show any obvious infant mortality in HH1. The same is the case of CD3. (Prior works have also observed the absence of infant mortality in certain cases, e.g. [28].) In contrast, CD1 uses disks from a different vendor and does show clear infant mortality. Figure 6 shows the normalized number of daily failures over a period of 4 years in CD1. This length of time implies 1 server refresh (which happened after 3 years of operation). The figure clearly shows 2 spikes of temporally clustered failures: one group 8 months after the first server deployment, and another 9 months after the first server refresh. A full deployment/refresh may take roughly 3 months, so the first group of disks may have failed 5-8 months after deployment, and the second 6-9 months after deployment.

A new disk lifetime model for free-cooled datacenters. Section 5.3 uses multiple sources of evidence to argue that relative humidity has a significant impact on hard-

ware reliability in humid free-cooled datacenters. The section also classifies the disk failures into two main groups: SMART and controller/connectivity. Thus, we extend the acceleration model above to include relative humidity, and recognize that there are two main disk lifetime processes in free-cooled datacenters: (1) one that affects the disk mechanics and SMART-related errors; and (2) another that affects its controller/connector.

We model process #1 as the Arrhenius acceleration factor, i.e. $AF_1 = AF_T$ (we define AF_T in Equation 1), as has been done in the past [10, 27]. For process #2, we model the corrosion rate due to high relative humidity and temperature, as both of them are known to affect corrosion [33]. Prior works have devised models allowing for more than one accelerating variable [12]. A general such model extends the Arrhenius failure rate to account for relative humidity, and compute a corrosion rate CR :

$$CR(\bar{T}, \overline{RH}) = const \cdot e^{\left(\frac{-E_a}{k\bar{T}}\right)} \cdot e^{(b\overline{RH}) + \left(\frac{c\overline{RH}}{k\bar{T}}\right)} \quad (2)$$

where \bar{T} is the average temperature, \overline{RH} is the average relative humidity, E_a is the temperature activation energy, k is Boltzmann's constant, and $const$, b , and c are other constants. Peck empirically found an accurate model that assumes $c = 0$ [24], and we make the same assumption. Intuitively, Equation 2 exponentially relates the corrosion rate with both temperature and relative humidity.

One can now compute the acceleration factor AF_2 by dividing the corrosion rate at the elevated temperature and relative humidity $CR(\bar{T}_e, \overline{RH}_e)$ by the same rate at the baseline temperature and relative humidity $CR(\bar{T}_b, \overline{RH}_b)$. Essentially, AF_2 calculates how much faster disks will fail due to the combined effects of these environmental. This division produces:

$$AF_2 = AF_T \cdot AF_{RH} \quad (3)$$

where $AF_{RH} = e^{b \cdot (\overline{RH}_e - \overline{RH}_b)}$ and AF_T is from Equation 1.

Now, we can compute the compound average failure rate \overline{FR} as $AF_1 \cdot \overline{FR}_{1b} + AF_2 \cdot \overline{FR}_{2b}$, where \overline{FR}_{1b} is the average mechanical failure rate at the baseline temperature, and \overline{FR}_{2b} is the average controller/connector failure rate at the baseline relative humidity and temperature. The rationale for this formulation is that the two failure processes proceed in parallel, and a disk's controller/connector would not fail at exactly the same time as its other components; AF_1 estimates the extra failures due to mechanical/SMART issues, and AF_2 estimates the extra failures due to controller/connectivity issues.

To account for varying temperature and relative humidity, we also weight the factors based on the duration of those temperatures and relative humidities. Other works have used weighted acceleration factors, e.g. [31]. For simplicity in the monitoring of these environmental, we can use the average temperature and average relative

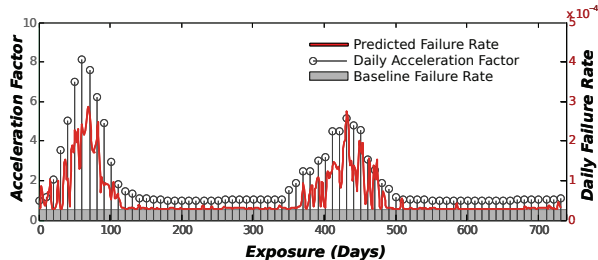


Figure 7: AFR prediction using our new disk lifetime model and the environmental conditions from HH1.

humidity per unit of time (e.g., hour, day) and compute the compound acceleration as a function of time t :

$$\overline{FR}(t) = AF_1(t) \cdot \overline{FR}_{1b}(t) + AF_2(t) \cdot \overline{FR}_{2b}(t) \quad (4)$$

Model validation. We validate our disk lifetime model using real temperature, relative humidity, and daily failure rate data from HH1. (In the next section, we also validate it for CD3.) We use daily failure rates as they are finer grained than annualized failure rates and capture the underlying processes well.

Since ours is an acceleration model, we need to build it on a baseline failure rate. Unfortunately, our dataset does not include data that we could use as such a baseline (e.g., daily failure rates for a datacenter that keeps temperature and relative humidity consistently low *and* uses the same disk model as HH1). Instead of this data, we use daily failure rates derived from Equation 4 with temperature and relative humidity fixed at 20°C and 30%, respectively. For this baseline, we use $E_a = 0.46$ as computed by Sankar *et al.* [27], and $b = 0.0455$ as we derive from the copper corrosion rate measured by Rice *et al.* [26]. These values produce an AFR = 1.5%, and an average daily failure rate of 1.5%/365.

Using the actual temperatures and relative humidities of HH1, we apply our acceleration model to the baseline average daily failure rate. We train the model with data for populations P2 and P3 (for these data, $b = 0.0652$), and validate the results with data from P1 and P4. We do not include any infant mortality effects, because HH1 does not exhibit them. If desired, infant mortality can be modeled using a Weibull distribution, as other authors have done, e.g. [28].

Figure 7 shows the predicted acceleration factors, predicted daily failure rates, and the baseline failure rate over 2 years, starting in the beginning of the summer. Based on the daily baseline and acceleration factors, we compute the predicted daily failure rates and depict them with the red curve. These results show that our model is accurate: from the red curve, we compute the predicted disk AFR for P1 and P4 to be 3.04%, whereas the real AFR for these populations is 3.18%.

As one would expect from Section 5.3, the relative humidity contributes the most to the accuracy of the model.

Removing the temperature components from the acceleration factors shows that it accounts for only 8.12% of the predicted AFR. Previous models (which are based solely on temperature) predict no increase in failure rates with respect to a baseline of 1.5%. Thus, models that do not account for relative humidity severely underestimate the corresponding failure rates.

7 Design Lessons

We next derive design lessons for servers and datacenters. We start by using our acceleration model to discuss different server layouts, and then discuss the implications of our findings to the placement, design, and management of datacenters.

Server design lessons. Our discussion of server designs relies on two main observations: (1) high relative humidity tends to produce more disk failures, as we have seen so far; and (2) higher temperature leads to lower relative humidity, given a constant amount of moisture in the air. These observations suggest that the layout of the disks within a server blade or enclosure may have a significant impact on its reliability in free-cooled datacenters.

To see this, consider Figure 8, where we present three possible layouts for a two-socket server blade. Our dataset has examples of all these layouts. In Figures 8(a) and (c), the disks are not exposed to the heat generated by the processors and memory DIMMs. This means that the relative humidity to which the disks will be exposed is roughly the same as that in the server’s air inlet. In contrast, in Figure 8(b), the disks will be exposed to lower relative humidity, as they are placed downstream from the processors and DIMMs. The difference in relative humidity in this layout can be significant.

To demonstrate this difference, we consider the inlet air temperature and relative humidity data from CD3, in which the server blades have the layout of Figure 8(b). This datacenter exhibits average inlet air temperatures of 19.9°C, leading to a much higher air temperature of roughly 42°C at the disks. Given these temperatures and an average inlet relative humidity of 44.1%, psychometrics calculations show that the average relative humidity at the disks would be only 13%. This is one of the reasons that CD3 exhibits such a low AFR. In fact, given this adjusted temperature and relative humidity, our acceleration model produces an accurate prediction of the AFR: 1.75% versus the real AFR of 1.8%.

In contrast with CD3, HH1 uses server blades with the layout in Figure 8(a). As HH1 exhibits a wider range of temperatures and relative humidities than CD3, we compute what the maximum relative humidity at any disk across the entire range of air temperatures would be at the back to the servers. The maximum values would be

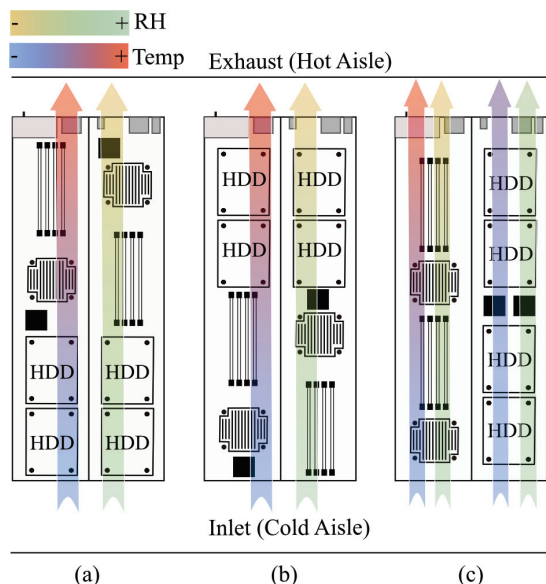


Figure 8: The server blade designs. Server designs with disks near the air exhaust lower their AFRs.

27% and 54.4°C at the disks. These values suggest that relative humidity should no longer have an impact on the AFR, while temperature should have a stronger impact. Overall, our model predicts AFR = 2.7%, instead of the actual AFR = 3.1% with disks in the front. Thus, changing the server layout in HH1 would decrease the importance of the relative humidity, increase the importance of temperature, but still produce a substantially lower AFR.

Obviously, placing disks at the back of servers could cause other components to experience higher relative humidity. Memory DIMMs would perhaps be the most relevant concern, as they also have exposed connectors. To prevent increases in AFR, DIMMs may be placed downstream from CPUs, which are the components most responsible for lowering relative humidity. Clearly, the full AFR and cost implications of the different layouts need to be understood, especially in free-cooled datacenters.

Datacenter lessons. As we argue in Section 5.4, free cooling reduces capital and operational cooling expenses significantly compared to other approaches. These savings may come at the cost of higher AFRs, depending on the outside environmental conditions at the datacenter’s location. Fortunately, the server design lessons above suggest that organizations can lower the disk AFR of their hot and humid datacenters by simply using server designs in which the disks are placed in the back.

In fact, since high relative humidity seems much more harmful than high temperature, operators may consider running their datacenters somewhat hotter in the summer, instead of increasing the relative humidity to keep the temperature lower. Obviously, this has to be done carefully, since our observations and inferences apply only to the conditions that our datacenters have experienced.

Another important lesson involves data availability/reliability. Clearly, a higher disk AFR could require the datacenter software (i.e., its online services and/or management systems) to manage data redundancy more aggressively. Fortunately, this does not pose a serious problem for at least two reasons: (1) in large datacenters with hundreds of thousands of hardware components, failures are a common occurrence, so software is already capable of tolerating them via data redundancy within or across datacenters; (2) the increases in disk AFR that may result from using free cooling at challenging locations (roughly 3x in our dataset) are not large enough that they would threaten the reliability of data stored by these datacenters. In the worst case, the software would add slightly more redundancy to the data (disk space has negligible cost per bit, compared to other TCO factors). Nevertheless, if software cannot manage the impact of the higher AFRs, datacenter operators must tightly control the relative humidity and site their datacenters in locations where this is more easily accomplished.

8 Conclusions

In this paper, we studied the impact of environmental conditions on the reliability of disk drives from nine datacenters. We also explored the tradeoffs between environmental conditions, energy consumption, and datacenter costs. Based on these analyses, we proposed and validated a new disk lifetime model that is particularly useful for free-cooled datacenters. Using the model and our data, we derived server and datacenter design lessons.

Based on our experience and observations, we conclude that high relative humidity degrades reliability significantly, having a much more substantial impact than temperature or temperature variation. Thus, the design of free-cooled datacenters and their servers must consider the relative humidity to which components are exposed as a first-class issue. Organizations that operate datacenters with different cooling technologies in multiple geographies can select the conditions under which their services strike the right tradeoff between energy consumption, hardware reliability, cost, and quality of service.

Acknowledgements: We thank Darren Carroll, Eoin Doherty, Matthew Faist, Inigo Gouri, John Hardy, Xin Liu, Sunil Panghal, Ronan Murphy, Brandon Rubenstein, and Steve Solomon for providing some of the data we use in this paper. We are also indebted to Annemarie Callahan, Matthew Faist, Dave Gauthier, Dawn Larsen, Bikash Sharma, Mark Shaw, Steve Solomon, Kushagra Vaid, and Di Wang for comments that helped improve the paper significantly. Finally, we thank our shepherd, Daniel Peek, and the anonymous referees for their extensive comments and support.

References

- [1] AHMAD, F., AND VIJAYKUMAR, T. Joint Optimization of Idle and Cooling Power in Data Centers while Maintaining Response Time. In *Proceedings of the 15th International Conference on Architectural Support for Programming Languages and Operating Systems* (2010).
- [2] ALLEN, B. Monitoring Hard Disks with SMART. *Linux Journal*, 117 (2004).
- [3] ANUBHAV, K., AND YOGENDRA, J. Use of Airside Economizer for Data Center Thermal Management. In *Proceedings of the 2nd International Conference on Thermal Issues in Emerging Technologies* (2008).
- [4] ASHRAE TECHNICAL COMMITTEE TC 9.9. 2011 Thermal Guidelines for Data Processing Environments – Expanded Data Center Classes and Usage Guidance, 2011.
- [5] BAIRAVASUNDARAM, L. N., GOODSON, G. R., PASUPATHY, S., AND SCHINDLER, J. An Analysis of Latent Sector Errors in Disk Drives. *ACM SIGMETRICS Performance Evaluation Review* 35, 1 (2007).
- [6] BANERJEE, A., MUKHERJEE, T., VARSAMOPOULOS, G., AND GUPTA, S. Cooling-Aware and Thermal-Aware Workload Placement for Green HPC Data Centers. In *Proceedings of the International Green Computing Conference* (2010).
- [7] BARROSO, L. A., CLIDARAS, J., AND HOLZLE, U. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, second ed. 2013.
- [8] BASH, C., PATEL, C., AND SHARMA, R. Dynamic Thermal Management of Air Cooled Data Centers. In *Proceedings of the 10th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems* (2006).
- [9] COLE, G. Estimating Drive Reliability in Desktop Computers and Consumer Electronics Systems. *Seagate Technology Paper*, TP-338.1 (2000).
- [10] EL-SAYED, N., STEFANOVICI, I., AMVROSIADIS, G., HWANG, A., AND SCHROEDER, B. Temperature Management in Data Centers: Why Some (Might) Like it Hot. In *Proceedings of the 12th International Conference on Measurement and Modeling of Computer Systems* (2012).
- [11] EMERSON NETWORK POWER. Energy Efficient Cooling Solutions for Data Centers, 2007.
- [12] ESCOBAR, L. A., AND MEEKER, W. Q. A Review of Accelerated Test Models. *Statistical Science* 21, 4 (2006).
- [13] FACEBOOK. Prineville Data Center – PUE/WUE, 2014.
- [14] GREENBERG, A., HAMILTON, J., MALTZ, D. A., AND PATEL, P. The Cost of a Cloud: Research Problems in Data Center Networks. *ACM SIGCOMM Computer Communication Review* 39, 1 (2008).
- [15] GREENBERG, S., MILLS, E., TSCHUDI, B., RUMSEY, P., AND MYATT, B. Best Practices for Data centers: Lessons Learned from Benchmarking 22 Data Centers. In *ACEEE Summer Study on Energy Efficiency in Buildings* (2006).
- [16] KAISER, J., BEAN, J., HARVEY, T., PATTERSON, M., AND WINIECKI, J. Survey Results: Data Center Economizer Use. *White Paper* (2011).
- [17] KIM, J., RUGGIERO, M., AND ATIENZA, D. Free Cooling-Aware Dynamic Power Management for Green Datacenters. In *Proceedings of the International Conference on High-Performance Computing and Simulation* (2012).
- [18] KOOMEY, J. Growth in Data Center Electricity Use 2005 to 2010, 2011. Analytic Press.
- [19] LIU, Z., CHEN, Y., BASH, C., WIERMAN, A., GMACH, D., WANG, Z., MARWAH, M., AND HYSER, C. Renewable and Cooling Aware Workload Management for Sustainable Data Centers. In *Proceedings of the 12th International Conference on Measurement and Modeling of Computer Systems* (2012).
- [20] MA, A., DOUGLIS, F., LU, G., SAWYER, D., CHANDRA, S., AND HSU, W. RAIDShield: Characterizing, Monitoring, and Proactively Protecting Against Disk Failures. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies* (2015).
- [21] MOORE, J., CHASE, J., RANGANATHAN, P., AND SHARMA, R. Making Scheduling “Cool”: Temperature-Aware Workload Placement in Data Centers. In *Proceedings of the 2005 USENIX Annual Technical Conference* (2005).
- [22] PATEL, C. D., SHARMA, R., BASH, C. E., AND BEITELMAL, A. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers. In *Proceedings of the 8th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems* (2002).
- [23] PATTERSON, M. K. The Effect of Data Center Temperature on Energy Efficiency. In *Proceedings of the 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems* (2008).
- [24] PECK, D. S. Comprehensive Model for Humidity Testing Correlation. In *Proceedings of the 24th Annual Reliability Physics Symposium* (1986).
- [25] PINHEIRO, E., WEBER, W.-D., AND BARROSO, L. A. Failure Trends in a Large Disk Drive Population. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies* (2007).
- [26] RICE, D., PETERSON, P., RIGBY, E. B., PHIPPS, P., CAPPELL, R., AND TREMOUREUX, R. Atmospheric Corrosion of Copper and Silver. *Journal of the Electrochemical Society* 128, 2 (1981).
- [27] SANKAR, S., SHAW, M., VAID, K., AND GURUMURTHI, S. Datacenter Scale Evaluation of the Impact of Temperature on Hard Disk Drive Failures. *ACM Transactions on Storage* 9, 2 (2013).
- [28] SCHROEDER, B., AND GIBSON, G. A. Understanding Disk Failure Rates: What Does an MTTF of 1,000,000 Hours Mean to You? *ACM Transactions on Storage* 3, 3 (2007).
- [29] SHAH, S., AND ELERATH, J. G. Reliability Analysis of Disk Drive Failure Mechanisms. In *Proceedings of the Annual Symposium on Reliability and Maintainability* (2005).
- [30] SRINIVASAN, J., ADVE, S. V., BOSE, P., AND RIVERS, J. A. The Case for Lifetime Reliability-Aware Microprocessors. In *Proceedings of the 41st Annual International Symposium on Computer Architecture* (2004).
- [31] THE GREEN GRID. Data Center Efficiency and IT Equipment Reliability at Wider Operating Temperature and Humidity Ranges, 2012.
- [32] TOLIA, N., WANG, Z., MARWAH, M., BASH, C., RANGANATHAN, P., AND ZHU, X. Delivering Energy Proportionality with Non-Energy-Proportional Systems – Optimizing the Ensemble. In *Proceedings of the HotPower Workshop* (2008).
- [33] TULLMIN, M., AND ROBERGE, P. R. Corrosion of Metallic Materials. *IEEE Transactions on Reliability* 44, 2 (1995).
- [34] UPTIME INSTITUTE. 2014 Data Center Industry Survey Results, 2014.
- [35] VISHWANATH, K. V., AND NAGAPPAN, N. Characterizing Cloud Computing Hardware Reliability. In *Proceedings of the 1st ACM Symposium on Cloud Computing* (2010).