

Impact of Environmental Factors on Flash Storage Performance in Autonomous Vehicles: An Empirical and Analytical Study

Ying He[†], Shuwen Liang[‡], Song Fu[†]

[†]Department of Computer Science and Engineering
University of North Texas, Denton, Texas, USA

[‡]Amazon Web Services

Amazon, Seattle, Washington, USA

{YingHe, ShuwenLiang}@my.unt.edu, Song.Fu@unt.edu

Abstract—In the ever-evolving landscape of automotive technology, the efficiency and reliability of vehicle storage systems is paramount importance. Environmental factors, such as extreme weather conditions characterized by either intense cold or scorching heat, pose significant challenges to the optimal functioning of these critical components [1] [2]. In this paper, we investigate the nuanced effects of varying temperatures on data storage and machine learning workloads on flash based vehicle storage systems: the read and write operations of car flash memory. The work employs a multidimensional approach, incorporating environmental simulations, performance testing, and data analytics to comprehensively analyze the impact of temperature variations on the performance and reliability of vehicle storage. From room temperatures to sweltering heat, the study investigates how these adverse conditions influence the speed, reliability, and overall performance of flash memory in automotive applications. The experimental results reveal intricate relationships between temperature variations and the throughput and latency of flash storage for automotive applications, shedding light on potential vulnerabilities and opportunities for optimization. Understanding these dynamics is crucial for enhancing the resilience and adaptability of automotive storage systems to diverse environmental challenges. This research not only contributes to the broader understanding of the intersection between extreme weather conditions and automotive technology but also provides valuable insights for engineers, manufacturers, and policymakers working towards the development of robust and reliable vehicle storage systems capable of withstanding the rigors of diverse environmental conditions. As the automotive industry continues to push the boundaries of innovation, this study serves as a foundation for future advancements in the realm of vehicular storage technologies.

Index Terms—Vehicle Storage, Flash Devices, Performance, Reliability, Machine Learning Workload.

I. INTRODUCTION

The rise of autonomous vehicles (AVs) marks a significant technological advancement in the automotive industry, driven by rapid developments in sensor technologies, deep learning algorithms, and computational power. These vehicles rely on advanced data processing units, essential for interpreting complex environmental inputs and enabling real-time decision-making.

AVs are equipped with a multitude of sensors, including cameras, LiDAR, radar, GPS, and IMUs, among others. These sophisticated sensors collectively generate terabytes of data every hour [3] [4], a critical component for mission-critical tasks like autonomous navigation and ensuring vehicular safety. To effectively process this deluge of data, AVs utilize advanced deep learning models [5] [6]. These models are designed to learn, perceive, plan, and make decisions. They process an extensive array of details, encompassing numerous parameters and features, to function optimally [7] [8] [9]. The sheer volume of data coupled with the complexity of these deep learning models necessitates a substantial amount of storage space.

Flash-based storage, particularly solid-state drives, have become the preferred choice for vehicular storage due to their durability, speed, and energy efficiency [10]. However, the performance of these storage systems under variable environmental conditions, particularly high temperatures, is a concern that warrants thorough investigation [11].

Elevated temperatures notably affect the performance and reliability of flash devices [12], a fact that is particularly critical for AVs. Increased temperatures can induce bit flips – minor but critical errors in the data stored. Additionally, the storage system might engage in speed throttling as a self-protective measure, slowing down its operation. Such issues are particularly crucial in AVs, which depend on precise and rapid data processing for safe operations. If the storage system malfunctions or underperforms, it can affect essential functions, for example, failing to respond appropriately in critical situations, such as avoiding a collision, or not detecting obstacles on the road. These issues may lead to significant safety hazards or accidents. For AVs, where reliable and consistent data storage is essential for both safety and functionality, understanding and mitigating the impact of high temperatures on flash-based storage is of significant importance.

In this paper, we delve into the detailed effects of temperature variations on the data storage capacity and efficiency of machine learning tasks in flash-based vehicle storage systems. Our study adopts an extensive experimental methodology,

utilizing a variety of flash storage devices in a controlled high-temperature environment. We design the experiments to simulate the real-world operational conditions of AVs, with a particular emphasis on assessing the performance of flash storage in managing sensor data storage and executing machine learning tasks. This approach enables a thorough exploration and understanding of the specific impacts of varying temperatures on these vital workloads, offering valuable insights into the performance and reliability of flash storage under thermal stress in the context of AVs.

The experimental results uncover interrelations between temperature changes and key performance metrics of flash storage, such as throughput and latency, particularly in automotive applications. These findings illuminate potential vulnerabilities in current systems and open avenues for optimization. Gaining insights into these temperature-related dynamics is imperative for improving the resilience and versatility of automotive storage systems, ensuring they are better equipped to handle a variety of environmental challenges.

The paper is organized as follows. Section II provides a literature review, focusing on flash storage in automotive applications and the effects of temperature on its performance. Section III details the methodology, describing the flash storage devices used and the experimental setup tailored to AV workloads. Section IV presents and analyzes the experimental results. The paper concludes with a summary and future research directions in Section V.

II. ARCHITECTURE OF FLASH-BASED STORAGE

Depending on the number of bits stored in each flash cell, there are several types of NAND flash used in SSD. Each type has its distinct performance, cost, endurance, and density trade-offs. Single-level cell (SLC) requires 2 voltage levels (i.e., 0 and 1) to store 1 bit of data, offering the highest write performance and endurance at the cost of price and density. Multilevel cell (MLC) requires 4 voltage levels to represent 2 bits of data (i.e., 00, 01, 10, and 11). Triple-level cell (TLC) and Quadruple-level cell (QLC) require 8 and 16 levels of voltage to store 3 and 4 bits of data, respectively. As the number of bits stored in each cell increases from SLC to QLC, the cost efficiency improves due to higher data density, but this comes at the expense of decreased write performance and endurance.

A. Data Storage

A flash-based storage device typically incorporates between 4 to 16 NAND chips [13]. Each chip is composed of multiple layers, structured as follows. **Die:** A NAND chip may contain several NAND memory dies. **Plane:** Each die consists of 1 to 4 planes. **Block:** Every plane is made up of thousands of flash blocks. **Page/Wordline:** Each block comprises hundreds to thousands of pages (rows, also known as wordlines). **String/Bitline:** Each block includes hundreds to thousands of strings (columns, also known as bitlines). **Cell:** Each page or string contains thousands of flash cells.

In 2D NAND architecture, a flash block is a cluster of wordlines and bitlines. A page is the smallest unit for reading and writing data, whereas a block is the smallest unit that can be erased. A page typically holds 2K, 4K, 8K, or 16KB of data, and the size of a block varies between 256KB and 4MB.

With the transition to TLC and beyond, there has been a shift from 2D to 3D NAND architecture. Each generation of 2D NAND saw a reduction in size and an increase in transistor density, reaching the limits of lithography [14]. 3D flash architecture differs by expanding vertically, increasing density and capacity. QLC, for example, uses 3D NAND in its block-level design. The transition to 3D blocks introduces the concept of layers, with the number of layers being multiples of 4 in the case of QLC. These layers are determined by the count of vertical Control Gates. Common configurations of 3D NAND flash include 32, 36, 48, 56, 96, and 128 layers, with a higher layer count typically resulting in greater storage density.

B. Data Placement

A flash storage device comprises two data storage areas: the main area and the spare area. The main area is primarily used for storing user data such as files and applications. In contrast, the spare area is reserved for critical system information, including bad block markers, Error-Correcting Code (ECC) data, and other metadata. This spare area is typically not accessible to users, being exclusively used by the system for ensuring the correct operations and integrity of the device.

The management of data within a flash device is orchestrated by the controller. The controller is responsible for a range of functions: it directs data to be written to one page at a time, identifies and avoids bad blocks, and implements wear-leveling algorithms to evenly distribute write and erase cycles. Additionally, it manages data transfer through I/O algorithms, efficiently handling tasks like dividing large files across multiple flash chips. To assure data integrity, the controller employs strategies such as writing parities to a separate flash chip, thereby enhancing the device's reliability and error-correction capabilities.

C. Data Operations

NAND flash devices support three basic operations: read, write, and erase.

Write: Data from the file system passes through the flash device connector and then to the controller. The Flash Translation Layer (FTL) executes an address mapping algorithm to determine the physical addresses on the NAND chips. In this process, FTL maps the logical data blocks to NAND pages, which are then written into a block. The controller applies a high positive voltage to the targeted NAND pages and strings. The voltages of the selected cells are altered to represent a logical '0'. Once the cell voltages are updated, the Error Correction Code (ECC) checks the written data before sending a success signal back to the OS. If the drive lacks empty blocks, the controller initiates a garbage collection process to reclaim invalid data blocks before writing new data to flash.

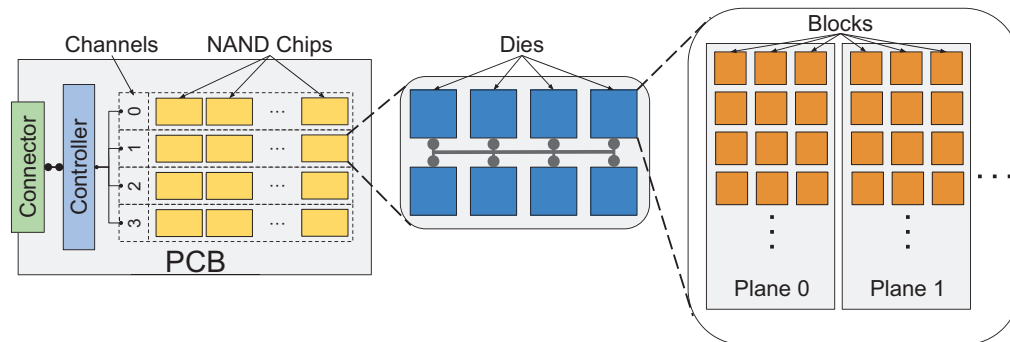


Fig. 1. Major Components of a Flash Storage Device.

Erase: The erase operation in NAND flash is performed only at the block level, whereas write and read can be conducted at the page level. An erase operation involves removing electrons from the storage layer, changing the state of the cell to a logical ‘1’. Typically, a delete request sent from the file system does not immediately remove the data from flash, the controller marks the data as invalid. The garbage collection algorithm, running in the background, determines the appropriate time to apply negative voltage to erase the entire block.

Read: Reading data from flash is similar to the write operation. FTL locates the physical address of the requested data, and the controller applies a read voltage (a medium positive voltage) to NAND pages and strings. Since the medium positive voltage does not change the logical representation of NAND cells, the selected NAND cells respond with the stored logical ‘0’ or ‘1’. The raw data then goes through the NAND decoder. After decoding and verification, the data stream is sent back to the file system.

III. FLASH STORAGE FOR AUTONOMOUS VEHICLES

A. Flash Storage for Vehicular Applications

Flash storage offers several key advantages for automotive applications. Its resistance to mechanical shock and vibration makes it ideal for the demanding conditions encountered in vehicle environments. Furthermore, flash storage boasts low power consumption, high data throughput, and swift access times, which are essential for real-time data processing in vehicles.

Modern flash SSDs typically offer read speeds ranging from about 500 MB/s to as high as 7,000 MB/s for PCIe 4.0 NVMe drives. Autonomous vehicles (AVs) generate a vast amount of data from various sensors. High-speed flash devices can quickly write the data, ensuring that the AV’s computing system is not bottlenecked by data throughput limitations. Additionally, for deep learning models used in AVs, the ability to quickly access and process large datasets is crucial. The high read speeds of flash storage facilitate faster data retrieval, which is vital for real-time decision-making for autonomous driving.

B. Environmental Factors

In vehicular applications, the reliability and performance of flash storage systems are critically influenced by environmental factors. These include temperature variations, humidity, vibration, and shock, prevalent in automotive environments, which can impact the functionality and longevity of flash storage devices in vehicles.

One of the most significant environmental challenges for vehicular flash storage is temperature variation. Flash drives typically operate within a temperature range of 32F to 158F for consumer-grade products. Vehicles are exposed to a wide range of temperatures, from the intense cold of winter environments to the extreme heat of summer conditions or engine compartments. Elevated temperatures can lead to accelerated wear and tear of the flash cells, resulting in reduced lifespan and reliability [15]. Every 10 degree increase in temperature can cut the life expectancy of a flash drive in half [15]. Conversely, extremely low temperatures can affect the speed and efficiency of data read/write operations, potentially leading to slower response times in critical applications.

C. NVMe QLC SSDs for Vehicles

For vehicular applications, quadruple-level cell (QLC) SSDs offer several advantages. 1) They have a higher storage capacity, as QLC SSDs store four bits of data per cell, enabling higher storage densities. This feature makes them ideal for large-capacity drives in vehicles with big data applications. 2) They are more cost-effective. QLC SSDs have a lower cost per gigabyte compared to other SSD types, presenting a budget-friendly option for high-capacity storage requirements. 3) They are more energy-efficient, an essential benefit for vehicles where power consumption is a critical factor.

The disadvantages of QLC SSDs include lower write endurance, slower write speeds, and performance degradation. Specifically, the lifespan of QLC SSDs is generally shorter compared to other types, due to fewer write cycles before cell degradation. Additionally, they tend to have slower write speeds, a notable drawback for tasks requiring frequent data writing, and this slowdown becomes more pronounced as the drive fills up. Moreover, performance can degrade significantly

under heavy workloads or as the storage capacity is maximized.

NVMe, a specialized interface protocol for SSDs, works in conjunction with PCIe to enhance data transfer to and from SSDs. SSDs that utilize the NVMe protocol, known as NVMe SSDs, are capable of achieving 64GB/s bandwidth and a 16GT/s rate with PCIe 4.0, and they do not require an additional power connector. With the introduction of PCIe 5.0, these SSDs are expected to reach speeds double that of PCIe 4.0 [16], which is approximately 20 times faster than those using SATA. Furthermore, NVMe SSDs typically have a smaller physical size compared to SATA-based SSDs. They often lack an outer protective casing, exposing the NAND chip and reducing their overall footprint.

Therefore, QLC SSDs offer a viable solution for high-capacity storage needs where cost is a primary concern and performance is not the main focus. Conversely, NVMe SSDs are ideal for high-performance computing where speed and low latency are crucial. In this study, our focus is on NVMe QLC SSDs for vehicular applications.

IV. ANALYSIS OF PERFORMANCE AND RELIABILITY OF FLASH STORAGE FOR VEHICLES

We have experimentally analyzed different models of NVMe QLC SSDs for vehicular applications. The tested models are from the same manufacturer but belong to different product series. Table I lists their specifications. To protect the manufacturer’s privacy, we refer to them as Model A and Model B. Model A SSDs are designed for consumer-grade applications, while Model B SSDs are intended for industrial-grade applications.

TABLE I
SPECIFICATIONS OF FLASH STORAGE DEVICES EVALUATED IN EXPERIMENTS.

Model	Cell Type	Architecture	Capacity	Firmware	Cost
Model A	NVMe QLC	3D	512GB	004c	\$
Model B	NVMe QLC	3D	118GB	k4110440	\$\$

A. Experiment Setup

We have conducted a series of experiments on QLC SSDs under various temperatures, observing the behavior of flash storage and systematically analyzing both performance and reliability. To simulate vehicular workloads, we utilize the Flexible I/O Tester (FIO) [17] and HiBench [18] to generate I/O accesses on the flash storage. The first set of experiments employ FIO to generate data store and access loads to SSDs. In the second set, we use HiBench to subject the flash devices to big data and machine learning loads. These experiments are conducted on HPE ProLiant servers equipped with 128 GB of memory, running Ubuntu v22.04.

B. Environmental Factors

Table II outlines the recommended operating temperatures for various hardware devices, including flash SSDs. The upper bound temperature refers to the point at which flash SSDs

begin to experience thermal throttling. When the temperature exceeds this threshold, protective measures are triggered. In severe situations, the SSD may disconnect, leading to the termination of running processes. Considering that the default thermal throttling threshold for NVMe drivers is 66°C, we have chosen 50°C as the upper bound temperature in our tests. In alignment with guidelines from the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE), we choose 25°C as the low temperature in our experiments.

C. Flash Performance Under Different Temperatures for Data Read/Write Workloads

We assess QLC SSDs at 25°C, 35°C, and 50°C to mimic normal, high, and very high temperature conditions, respectively. We evaluate a number of critical metrics, including storage throughput, Input/Output Operations Per Second (IOPS), and latency. Within each metric, we examine both sequential and random read/write performance. The comparative results from Model A and Model B flash devices are shown in Tables III and IV. We perform each test on three separate flash SSDs of the identical model.

D. Experimental Results

Tables III and IV show that both Model A and Model B SSDs exhibit similar patterns of performance degradation when the environmental temperature increases from 25°C to 50°C. However, noticeable differences between them are observed.

Specifically, at 25°C, both Model A and Model B QLC SSDs achieve their peak throughput. For both models, the read throughput significantly exceeds the write throughput in both sequential and random contexts. Model A shows higher efficiency in sequential writing compared to random writing, while it exhibits the opposite trend in reading capabilities. Conversely, Model B maintains consistent write and read throughput, whether the operations are sequential or random. In terms of IOPS, Model B outperforms Model A, particularly in read operations. Model A registers higher IOPS in writing than in reading for sequential tasks, but this trend is reversed for random tasks. Model B, however, shows more consistent IOPS values across all operations, with reading operations consistently achieving higher figures. As for latency, reading operations are markedly quicker than writing for both models, with the exception of Model A’s random read/write performance.

As the temperature rises from 25°C to 35°C, both models exhibit a decline in throughput, with Model A experiencing a more noticeable reduction. The most significant decrease in throughput for Model A is observed in random reads. In contrast, Model B shows a similar level of degradation in all reading tasks, both sequential and random. For both models, reading operations tend to degrade slightly more than writing operations with the temperature increase. Notably, Model A’s IOPS values decrease significantly, while Model B’s IOPS demonstrate a slight improvement, especially in random read

TABLE II
RECOMMENDED OPERATING TEMPERATURES FOR STORAGE AND NETWORK DEVICES.

Component	NVMeSSD[19]	HDD[20]	Network Card[21]	NetworkSwitch[22]
Lower-bound Temperature	-40°C	0°C	0°C	-40°C
Upper-bound Temperature	85°C	60°C	90°C	85°C

TABLE III
PERFORMANCE RESULTS OF THROUGHPUT, IOPS, AND LATENCY USING MODEL A SSD.

Benchmarks		Throughput (MB/s)			IOPS			Latency (μ s)		
Temperature		25°C	35°C	50°C	25°C	35°C	50°C	25°C	35°C	50°C
Sequential	write zero	489.20	462.68	419.34	111154.06	61741.12	81527.78	39.19	44.20	47.46
	read zero	1339.71	1142.67	686.95	84253.43	64027.69	76975.22	18.04	20.60	19.35
	write random	518.89	455.17	442.80	119787.16	67122.74	90618.69	25.19	41.22	43.79
	read random	1288.10	1144.08	698.58	82676.31	63223.44	80812.20	18.11	20.99	19.20
Random	random read	1652.14	1316.74	387.29	84779.91	71160.59	42219.64	100.48	101.42	130.45
	random write	273.28	265.02	174.75	35946.24	28676.39	18579.51	34.61	40.43	95.95

TABLE IV
PERFORMANCE RESULTS OF THROUGHPUT, IOPS, AND LATENCY USING MODEL B SSD.

Benchmarks		Throughput (MB/s)			IOPS			Latency (μ s)		
Temperature		25°C	35°C	50°C	25°C	35°C	50°C	25°C	35°C	50°C
Sequential	write zero	594.57	580.60	425.38	135355.45	138732.57	108127.24	17.95	16.57	17.98
	read zero	1368.11	1302.76	1112.41	323729.82	332679.82	284557.43	12.04	10.95	12.12
	write random	594.38	580.40	434.41	135254.17	138623.00	109172.18	17.94	16.46	17.70
	read random	1368.32	1303.44	1078.78	322955.72	332516.35	275316.90	12.15	11.08	12.19
Random	random read	1368.12	1302.08	11083.03	324254.89	442449.81	275220.02	12.19	10.98	12.32
	random write	587.19	577.45	419.91	134667.87	138015.14	105696.33	17.69	16.52	17.72

operations. In terms of latency, Model A shows an increase in latency values as the temperature rises from 25°C to 35°C, whereas Model B displays a minor decrease in latency during the same temperature change.

That is, as the temperature rises from 25°C to 35°C, Model A SSD exhibits a decline in performance across throughput, IOPS, and latency. Conversely, Model B SSD experiences a decrease in I/O throughput but demonstrates slight improvements in both IOPS and latency.

As the temperature increases from 25°C to 50°C, we observe a degradation in all performance metrics - throughput, IOPS, and latency - for both models. Both Model A and Model B show a noticeable decline in throughput and IOPS, with Model A experiencing more severe degradation, particularly in reading operations, a point that will be elaborated upon in later sections. The reduction in IOPS is similar for both models. Regarding latency, Model A exhibits a significant worsening, whereas Model B also undergoes degradation, albeit less markedly than Model A.

E. Performance Analysis

To provide a thorough understanding of these performance shifts, we calculate the degradation percentage for both models, as shown in Tables V and VI. These tables depict how performance metrics change with an increase in temperature starting from 25°C.

Furthermore, the calculation of the degradation percentage is based on the premise that higher values signify improved performance for throughput and IOPS. Thus, we anticipate higher values at lower temperatures compared to those at elevated temperatures. In contrast, for latency, where lower values are desired, we expect that latency will be lower at reduced temperatures than at increased temperatures. In the following equation, l represents the throughput/IOPS value at the lower temperature, and h denotes the throughput/IOPS value at the higher temperature.

$$Throughput/IOPS_Degradation = \frac{l - h}{l} * 100\%. \quad (1)$$

For latency, l denotes latency at a lower temperature, and h represents latency at a higher temperature.

$$Latency_Degradation = \frac{h - l}{l} * 100\%. \quad (2)$$

Both Table V and Table VI illustrate that the I/O bandwidth performance, i.e., the rate of data transfer, deteriorates noticeably as the temperature increases from 25°C to 50°C. The decline in performance is relatively minor when the temperature rises from 25°C to 35°C. However, it becomes significantly more pronounced as the temperature continues to increase.

TABLE V
PERFORMANCE DEGRADATION W.R.T. THROUGHPUT, IOPS AND LATENCY UNDER INCREASED TEMPERATURE USING MODEL A SSD.

Benchmarks		Throughput(MB/S)		IOPS		Latency(μ s)	
Temperature		25°C ->35°C	25°C ->50°C	25°C ->35°C	25°C ->50°C	25°C ->35°C	25°C ->50°C
Sequential	write zero	5.42%	14.28%	44.45%	26.65%	12.78%	21.10%
	read zero	14.71%	48.72%	24.01%	8.64%	14.21%	7.28%
	write random	12.28%	14.66%	43.97%	24.35%	63.64%	73.84%
	read random	11.18%	45.77%	23.61%	2.36%	15.88%	5.95%
Random	random read	20.40%	76.56%	16.06%	50.20%	0.93%	29.82%
	random write	3.02%	36.05%	20.22%	48.31%	16.80%	177.20%

TABLE VI
PERFORMANCE DEGRADATION W.R.T. THROUGHPUT, IOPS AND LATENCY UNDER INCREASED TEMPERATURE USING MODEL B SSD.

Benchmarks		Throughput(MB/S)		IOPS		Latency(μ s)	
Temperature		25°C ->35°C	25°C ->50°C	25°C ->35°C	25°C ->50°C	25°C ->35°C	25°C ->50°C
Sequential	write zero	2.35%	28.46%	-2.50%	20.12%	-7.69%	0.17%
	read zero	4.78%	18.69%	-2.76%	12.10%	-9.05%	0.66%
	write random	2.35%	26.91%	-2.49%	19.28%	-8.25%	-1.34%
	read random	4.74%	21.16%	-2.96%	14.75%	-8.81%	0.33%
Random	random read	4.83%	20.84%	-36.45%	15.12%	-9.93%	1.07%
	random write	1.66%	28.49%	-2.49%	21.51%	-6.61%	0.17%

Model B exhibits a more consistent performance decrease across read and write tasks compared to Model A. For example, with a temperature increase from 25°C to 50°C, all of the performance degradation for Model B are between 18% and 30%. In contrast, the performance declines for Model A vary widely from 14% to 76% under the same temperature conditions. Specifically, Model A demonstrates a greater decrease in performance for read tasks than for write tasks, particularly under random read operations. Model B, on the other hand, shows a different pattern. When the temperature increases from 25°C to 35°C, Model B’s performance in read tasks declines more than in write tasks. However, as the temperature rises from 25°C to 50°C, its write tasks are more adversely affected by the temperature change than the read tasks.

Examining IOPS (a metric indicating the number of read or write operations completed per second), Model A exhibits a greater decrease in performance for write tasks than for read tasks. The most significant decline occurs as the temperature increases from 25°C to 35°C, particularly with sequential I/O operations. In contrast, Model B’s IOPS performance slightly improves at 35°C compared to 25°C, especially with random read operations. However, these improvements are small, around 2.5%, making it challenging to determine whether they are meaningful or merely the result of random fluctuations. As the temperature further rises to 50°C, performance diminishes again.

F. P-value Analysis

Based on the performance results at 25°C, there is a significant performance degradation in both values and percentages when the temperature reaches 50°C. To more confidently validate this finding, we have also employed the T-test to assess

the significance of the results. By analyzing the p-values, we aim to reliably demonstrate the extent of degradation for each benchmark.

The T-test is a type of inferential statistic used to determine if there is a significant difference between the means of two groups. To analyze the significance levels, a two-tailed equal variance T-test is utilized in this experiment. The p-values are calculated using the following equation.

$$P_value = \frac{mean1 - mean2}{\frac{(n1-1)*var1^2 + (n2-1)*var2^2}{n1+n2-2} * \sqrt{\frac{1}{n1} + \frac{1}{n2}}}. \quad (3)$$

In Equation (3), *mean1* and *mean2* represent the average values of sample sets. *Var1* and *var2* are the variances, while *n1* and *n2* are the number of records in sample sets.

P-values are used to assess the likelihood that data can have occurred by chance [23]. They range from 0 to 1, with a smaller p-value indicating stronger evidence against the null hypothesis, which is our initial assumption. In this study, the null hypothesis posits that there is no difference in the QLC SSD’s performance between 25°C and 50°C. Typically, a p-value less than 0.05 is considered statistically significant, implying there is less than a 5% chance that the initial assumption is correct. A p-value less than 0.01 is deemed even more significant, suggesting there is less than a 1% chance that our initial assumption is correct. Table VII and Table VIII present the p-values for Model A SSD and Model B SSD, respectively.

Most p-values for both SSD models point to significant performance degradation, with some indicating more significant results. For Model A, the majority of throughput benchmarks yield p-values less than 0.01, marking them as highly significant, except for the ‘write zero’ benchmark, which exhibits no

significant change. Within the IOPS category, the 'read zero' benchmark shows no significant change, while the rest are significant. In terms of latency, only the random read and write benchmarks in Model A demonstrate significant degradation. Compared to sequential I/O operations, random I/O operations are more susceptible to degradation with temperature increases. For Model B, all benchmarks within the throughput and IOPS categories display extremely significant degradation, whereas the latency category reveals no significant changes.

After examining the average values, percentages, and p-values from the data read/write benchmark tests, we have the following major findings.

- 1) As the temperature exceeds their normal operating range, QLC SSDs tend to exhibit lower throughput and IOPS (indicating how quickly they can read/write data). However, we observed that the latency (the delay before a transfer of data begins following an instruction) of the SSDs does not significantly change with higher temperatures.
- 2) When comparing two different models of QLC SSDs, their reactions to temperature changes vary. Model A is more adversely affected by temperature increases for read tasks, whereas Model B is more sensitive to temperature changes for write tasks. At 50°C, the performance decrease for write by Model B is slightly greater than for read.

Additionally, when evaluating the overall performance of the two models, Model B appears to manage performance drops more effectively and maintains greater consistency in throughput, IOPS, and latency, even as it degrades. For instance, as temperatures increase from 25°C to 50°C, Model A's throughput can decrease by anywhere from 14% to 76%, whereas Model B's decreases range from 15% to 28%. A similar pattern is observed in IOPS. We suspect that the distinct firmware used by each model might be a contributing factor.

G. Experiments Using Machine Learning and Big Data Applications

The FIO benchmark test provides an understanding of how different QLC SSD models respond to rising temperatures. However, real-life scenarios are more complex. With the increasing prevalence of machine learning and big data and the use of QLC SSDs for high-end storage, examining more realistic applications is crucial. The experiments enhance our comprehension of the performance and reliability of flash-based storage under elevated temperatures for vehicle applications. In this part of the study, we test QLC SSDs Model A and Model B in environments of 25°C, 35°C, 45°C, and 50°C. We subject them to various machine learning and big data tasks to observe how their data handling performance varies at these temperatures.

The experiment is conducted in two steps. First, we establish the big data environment and conduct various benchmarks at a temperature of 25°C. We utilize Blktrace to monitor the I/O

activity, enabling us to individually record the I/O traces for Model A and Model B SSDs.

In the subsequent step, we employ FIO to replay these recorded traces 100 times, albeit under different temperature conditions. Throughout these replays, we collect data on throughput and other performance metrics using various monitoring tools. The rationale behind replaying the traces 100 times stems from preliminary tests, which suggested that QLC SSDs do not show immediate or consistent degradation across all benchmarks. It is observed that different benchmarks exhibit degradation at varying intervals. Repeating the process 100 times ensures that all benchmarks not only show a reduction in throughput but also maintain this decline steadily over time, preventing secondary phase of degradation.

H. Machine Learning and Big Data Benchmarks and I/O Trace Analysis

In this set of experiments, the workloads include micro benchmarks, machine learning applications, and search applications. For the micro benchmarks, we test DFSIOE-read, DFSIOE-write, and sort. Additionally, we test the Bayes and KMeans from the machine learning applications. For the search benchmarks, we test Nutch indexing and PageRank. All settings are presented in Table IX.

- **Byes:** Bayes is a machine learning algorithm used for classification. In the HiBench benchmark suite, it is utilized to generate documents with words following a Zipfian distribution [25].
- **Kmeans:** Kmeans is a machine learning algorithm designated for clustering data into groups. In HiBench, the testing data follow Uniform Distribution and Gaussian Distribution [25]. In the experiment, it is configured to produce five clusters. The Kmeans workload operates in phases, processing one cluster at a time until completion.
- **Sort:** This benchmark involves sorting a collection of randomly generated records. It uses simple functions for map and reduce, like those used in the MapReduce framework, which actually carries out the sorting task [24].
- **Dfsioe:** As part of the HiBench suite, the Dfsioe benchmark represents an enhanced version of the DFSIO benchmark. It aims to evaluate the data management capability of HDFS in a cluster. This is achieved by initiating numerous tasks that involve both writing to and reading from the system. The benchmark assesses three critical metrics: the average I/O rate per task, the average data throughput per task, and the total throughput of the cluster [25].
- **Nutchindexing:** This workload examines the indexing component of Nutch, a web crawler software. It simulates web data, with both links and words adhering to the Zipfian distribution, to test the effectiveness of data patterning [25].
- **Pagerank:** Pagerank is an algorithm employed by search engines to rank web pages in search results. The Pagerank benchmark within HiBench evaluates the performance of

TABLE VII
P-VALUES AND PERFORMANCE OF MODEL A SSD

Benchmarks		Throughput(MB/s)	IOPS	Latency(μ s)
Temperature		25°C vs. 50°C	25°C vs. 50°C	25°C vs. 50°C
Sequential	write zero	$1.91E-01^N$	$1.78E-01^*$	$4.49E-01^N$
	read zero	$9.95E-05^{**}$	$8.80E-02^N$	$9.59E-02^N$
	write random	$8.51E-04^{**}$	$1.16E-03^{**}$	$1.32E-01^N$
	read random	$7.26E-04$	$4.76E-02^{**}$	$1.68E-01^N$
Random	random read	$1.70E-08^{**}$	$5.94E-07^{**}$	$1.09E-04^{**}$
	random write	$3.70E-03^{**}$	$1.13E-02^*$	$3.81E-02^*$

N indicates p-value > 0.05 , i.e., the performance is NOT significant;
* indicates p- value ≤ 0.05 , i.e., the performance is significant;
** indicates p-value ≤ 0.01 , i.e., the difference is highly significant.]

TABLE VIII
P-VALUES AND PERFORMANCE OF MODEL B SSD

Benchmarks		Throughput(MB/s)	IOPS	Latency(μ s)
Temperature		25°C vs. 50°C	25°C vs. 50°C	25°C vs. 50°C
Sequential	write zero	$8.85E-05^{**}$	$1.17E-04^{**}$	$8.39E-01^N$
	read zero	$7.41E-04^{**}$	$4.52E-03^{**}$	$4.01E-01^N$
	write random	$3.51E-05^{**}$	$2.81E-04^{**}$	$5.56E-01^N$
	read random	$9.12E-04^{**}$	$2.80E-03^{**}$	$1.68E-01^N$
Random	random read	$9.50E-04^{**}$	$1.94E-03^{**}$	$8.49E-01^N$
	random write	$1.01E-04^{**}$	$2.04E-04^{**}$	$2.10E-01^N$

N indicates p-value > 0.05 , i.e., the performance is NOT significant;
* indicates p- value ≤ 0.05 , i.e., the performance is significant;
** indicates p-value ≤ 0.01 , i.e., the difference is highly significant.]

TABLE IX
MACHINE LEARNING AND BIG DATA BENCHMARKS TESTED IN EXPERIMENTS.

Benchmarks		Dataset	Size
Micro	Wordcount	Huge	32000MB
	Sort	Huge	3200MB
	Dfsioe-read	Huge	256*100MB
	Dfsioe-read	Huge	256*100MB
Machine Learning	Bayes	Large	pages: 100000 classes: 100 ngrams:2 number of clusters: 5 dimensions: 20
	Kmeans	Large	number of samples: 20000000 samples per input file: 4000000 maximum iteration: 5 k: 10
Search	Nutchindexing	Small	pages: 1000000 pages: 5000
	Pagerank	Small	number of iterations: 3 blocks: 0 block width: 16

Pagerank algorithm using web data structured in a Zipfian distribution [25].

Table X presents the read/write ratios for each benchmark across both SSD models. In the following equation, n_r represents the number of read operations and n_w signifies the number of write operations.

$$Ratio(read, write) = \frac{n_r}{n_w}. \quad (4)$$

Model A and Model B SSDs undergo similar workloads as evaluated through the following benchmarks.

- Bayes and Kmeans: In these benchmarks, read requests are sporadic throughout the processes. Note Kmeans exhibits a surge in write requests towards its conclusion.
- Dfsioe: Dfsioe-read and Dfsioe-write function according to their designations. Dfsioe-read focuses on read requests, whereas Dfsioe-write emphasizes write requests. Both of them encompass a mix of read and write activities, not solely one or the other.
- Wordcount: Initially, read and write requests are almost equally distributed. Over time, however, read requests significantly outnumber write requests, rendering Wordcount predominantly read-intensive.
- Sort: This benchmark experiences a notable surge in write

TABLE X
READ/WRITE RATIOS OF TESTED BENCHMARK APPLICATIONS.

Benchmarks		Model A	Model B
Micro	Wordcount	2.05	2.37
	Sort	0.38	0.57
	Dfsioe-read	0.54	0.81
	Dfsioe-write	0.48	0.55
Machine Learning	Bayes	0.17	0.13
	Kmeans	0.62	0.69
Web Search	Nutchindexing	1.12	1.17
	Pagerank(raw)	0/90	0/98

TABLE XI
THROUGHPUT OF MODEL A SSD UNDER VARYING TEMPERATURES.

Temperature		25°C		35°C		45°C		50°C	
Throughput(MB/s)		Read	Write	Read	Write	Read	Write	Read	Write
Micro	Wordcount	1411.26	5.21	1426.44	5.26	1443.36	5.33	1258.60	4.64
	Sort	82.32	777.51	84.57	798.68	82.96	783.61	75.118	710.82
	Dfsioe-read	1325.00	39.89	1333.22	40.15	1346.66	40.57	1211.80	36.48
	Dfsioe-write	0.09	878.88	0.09	879.34	0.09	842.89	0.08	766.32
Machine Learning	Bayes	73.68	710.12	79.15	763.13	78.19	754.11	65.74	633.71
	Kmeans	118.85	758.67	120.49	768.89	117.44	749.81	101.64	649.03
Web Search	Nutchindexing	223.07	534.30	221.57	530.71	218.84	524.31	181.62	435.12
	Pagerank	NA	167.45	NA	167.52	NA	165.48	NA	155.26

TABLE XII
THROUGHPUT OF MODEL B SSD UNDER VARYING TEMPERATURES.

Temperature		25°C		35°C		45°C		50°C	
Throughput(MB/s)		Read	Write	Read	Write	Read	Write	Read	Write
Micro	Wordcount	1251.92	4.64	1248.90	4.63	1242.98	4.60	931.92	3.45
	Sort	127.99	530.67	127.83	529.74	127.40	527.88	85.49	354.63
	Dfsioe-read	1179.52	38.97	1174.62	38.78	1173.83	38.78	898.39	29.67
	Dfsioe-write	0.06	586.81	0.06	588.60	0.06	576.44	0.04	357.49
Machine Learning	Bayes	50.61	562.66	50.55	564.17	51.09	565.51	32.85	364.28
	Kmeans	95.80	537.14	96.17	539.01	96.03	538.17	65.69	368.47
Web Search	Nutchindexing	198.95	490.39	199.70	491.57	199.51	491.51	127.20	313.28
	Pagerank	NA	437.59	NA	441.89	NA	444.79	NA	444.46

requests midway through the process, peaking at 7,000 requests for Model A SSD and 5,000 for Model B SSD, while read requests remain comparatively low. Thus, Sort is characterized as write-intensive.

I. Experimental Results and Performance Degradation Analysis

Table XI and Table XII display the average I/O throughput of Model A and Model B SSDs under different temperature.

When the temperature is at or below 45°C, both Model A and Model B SSDs achieve their highest read throughput running Wordcount. This benchmark, together with the read-intensive Dfsioe-read, displays throughput close to their peak performance. When running Dfsioe-write, both models of SSDs reach their highest write throughput. Other write-focused benchmarks, such as Sort, Bayes, and Kmeans, also exhibit a relatively high write throughput. However, Pagerank, which is also write-intensive, shows a lower I/O throughput compared to the other benchmarks that prioritize write.

Figures 2, 3, and 4 compare the performance of Models A and B SSDs across varying temperatures (25°C, 35°C, and 50°C) for both sequential and random read/write operations.

- Throughput of read/write operations: Model A SSD generally exhibits a higher throughput at lower temperatures but experiences a significant drop as the temperature increases. In contrast, Model B SSD, while starting with a lower throughput, maintains more consistent performance across different temperatures.
- IOPS of read/write operations: At lower temperatures, Model A SSD shows superior IOPS performance but suffers a notable decrease as the temperature rises. Model B SSD demonstrates more stable IOPS figures across increasing temperatures.
- Latency of read/write operations: Model A SSD experiences increased latency with rising temperatures, particularly for random read and write operations. Model B SSD, starting with slightly higher latency at lower temperatures, shows a smaller increase in latency at higher temperatures, indicating better resilience under thermal stress.

In summary, Model A SSD performs better at lower temperatures, whereas Model B SSD offers more consistent performance across a broader temperature range.

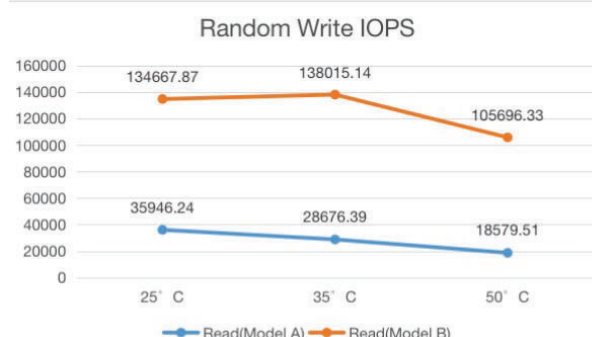
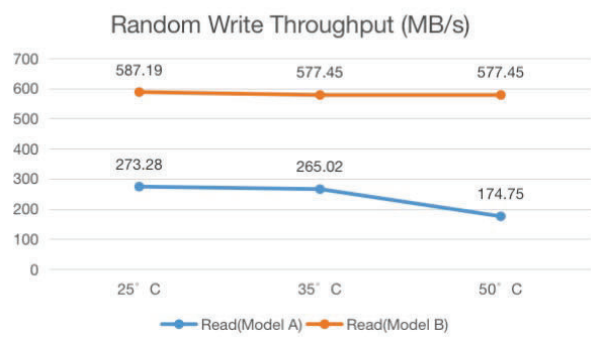
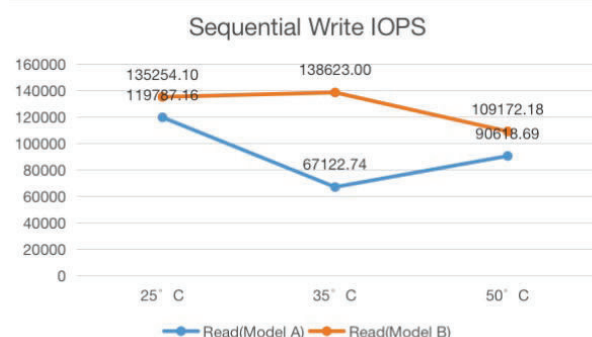
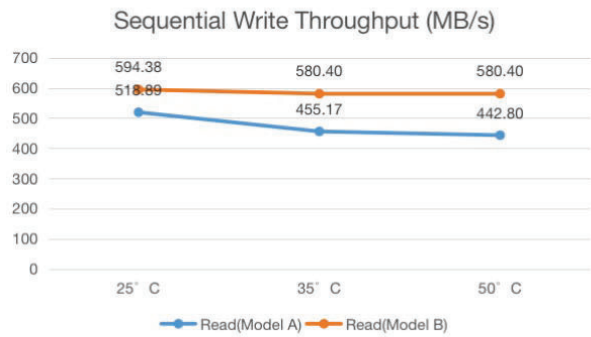

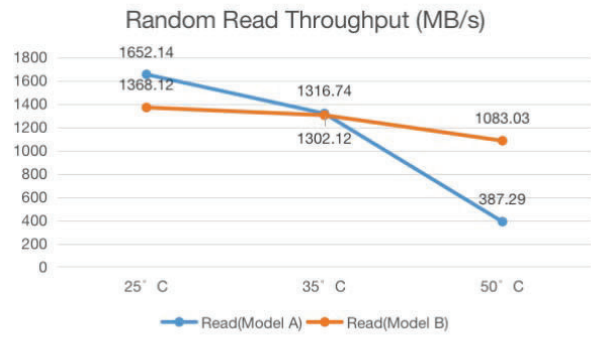
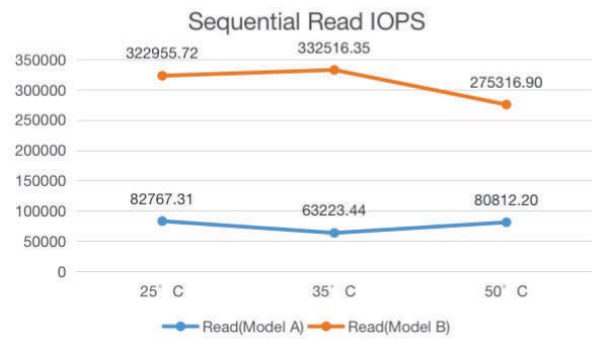
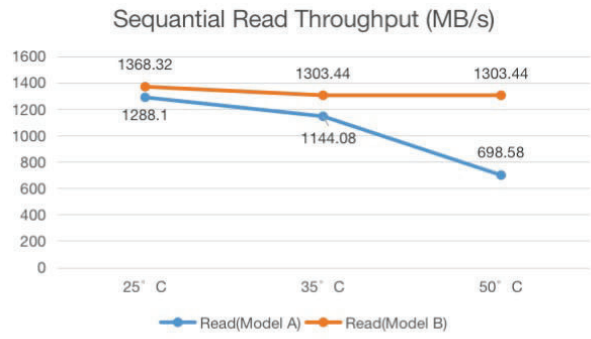


Fig. 2. Throughput for Flash Read/Write Operations.

Fig. 3. IOPS for Flash Read/Write Operations.

V. CONCLUSIONS

This paper explores the intricate relationship between extreme weather conditions and the efficiency of automotive flash memory, with a focus on the distinctive characteristics of different SSDs. Through extensive environments, performance evaluations, and data analysis, we gain insights into how temperature extremes affect the read and write operations of these essential automotive components.

Our work highlights the significant impact of temperature on automotive flash memory. Different models of SSDs, each embodying different aspects of automotive technology, show varied sensitivities to extreme temperatures. Under lower temperatures, we observe a slowdown in the responsiveness of certain flash memory components, affecting both read and write speeds. On the other hand, intense heat is associated with decreased reliability and higher error rates.

The side-by-side comparison of different models of SSDs has yielded a thorough understanding of the vulnerabilities and strengths present in different automotive memory systems. This information is crucial for designers and manufacturers to improve the resilience and adaptability of automotive flash memory in various climates, ensuring consistent performance regardless of external temperature fluctuations.

As the automotive industry moves toward connected and autonomous vehicles, the findings from this study are invaluable in guiding the design and development of future automotive memory technologies. The approach taken in this investigation has offered a detailed perspective, enriching our overall understanding of the challenges and opportunities that extreme weather conditions present.

The effect of temperature on automotive flash memory is clear, and our study lays the groundwork for further research and innovation in this vital area of automotive technology. By recognizing and mitigating the impacts of temperature extremes, we are setting the stage for the development of more durable, reliable, and adaptable automotive flash memory systems capable of withstanding the varied climates encountered on roads.

ACKNOWLEDGMENT

This work has been supported in part by the U.S. NSF grants CNS-2231519, CNS-2113805, OAC-2017564, CNS-2037982, and DUE-2225229. We thank the reviewers for their constructive comments, which helped us improve this paper.

REFERENCES

- [1] Wu, Dan, et al. "Experimental investigation of threshold voltage temperature effect during cross-temperature write-read operations in 3-D NAND flash." *IEEE Journal of the Electron Devices Society* 9 (2020): 22-26.
- [2] Johnson, R. Wayne, et al. "The changing automotive environment: high-temperature electronics." *IEEE transactions on electronics packaging manufacturing* 27.3 (2004): 164-176

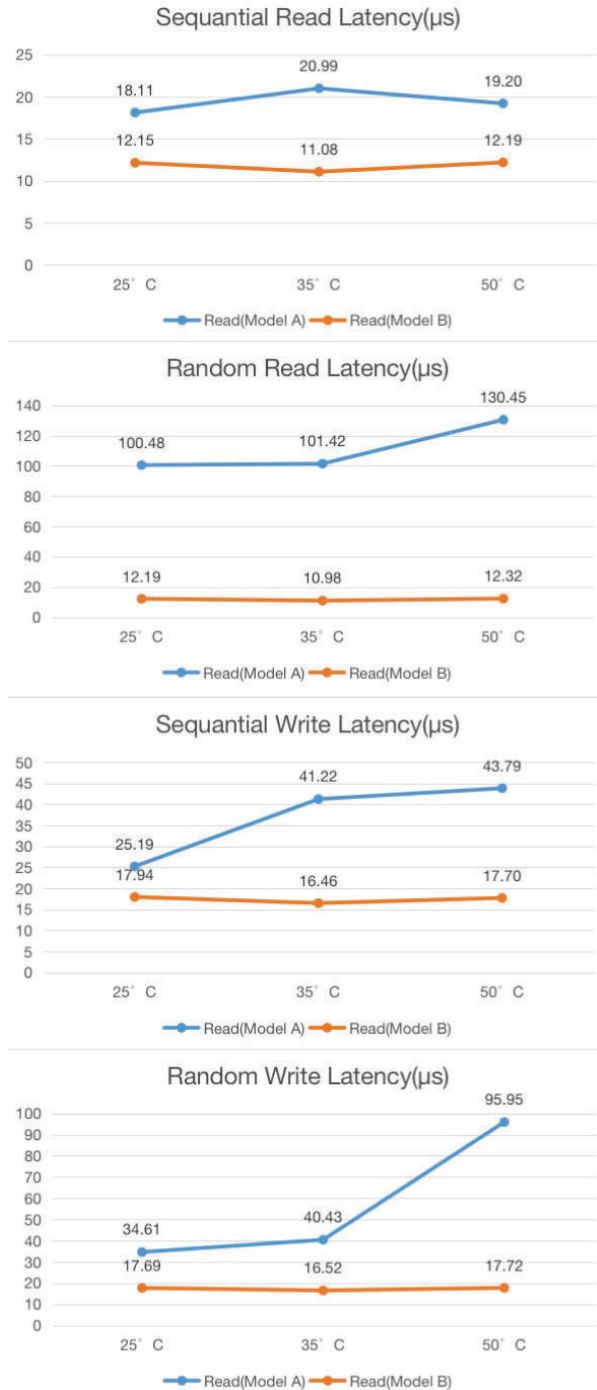


Fig. 4. Latency for Flash Read/Write Operations.

- [3] D. J. Yeong, G. Velasco-Hernandez, J. Barry, and J. Walsh, "Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review," *Sensors*, vol. 21, no. 6, p. 2140, Mar. 2021, doi: <https://doi.org/10.3390/s21062140>.
- [4] Vargas, Jorge, et al. "An overview of autonomous vehicles sensors and their vulnerability to weather conditions." *Sensors* 21.16 (2021): 5397.
- [5] Miglani, Arzoo, and Neeraj Kumar. "Deep learning models for traffic flow prediction in autonomous vehicles: A review, solutions, and challenges." *Vehicular Communications* 20 (2019): 100184.
- [6] Jebamikyous, Hrag-Harout, and Rasha Kashef. "Autonomous vehicles perception (avp) using deep learning: Modeling, assessment, and challenges." *IEEE Access* 10 (2022): 10523-10535.
- [7] Mozaffari, Sajjad, et al. "Deep learning-based vehicle behavior prediction for autonomous driving applications: A review." *IEEE Transactions on Intelligent Transportation Systems* 23.1 (2020): 33-47.
- [8] Bachute, Mrinal R., and Javed M. Subhedar. "Autonomous driving architectures: insights of machine learning and deep learning algorithms." *Machine Learning with Applications* 6 (2021): 100164.
- [9] Kuutti, Sampo, et al. "A survey of deep learning applications to autonomous vehicle control." *IEEE Transactions on Intelligent Transportation Systems* 22.2 (2020): 712-733.
- [10] Park, Seonyeong, et al. "A comprehensive study of energy efficiency and performance of flash-based SSD." *Journal of Systems Architecture* 57.4 (2011): 354-365.
- [11] Xu, Erci, et al. "Lessons and actions: What we learned from 10k ssd-related storage system failures." 2019 USENIX Annual Technical Conference (USENIX ATC 19), 2019.
- [12] Chen, Fei, et al. "Temperature impacts on endurance and read disturbs in charge-trap 3D NAND flash memories." *Micromachines* 12.10 (2021): 1152.
- [13] TechTarget. "NAND Flash Memory." SearchStorage, TechTarget, n.d., <https://www.techtarget.com/searchstorage/definition/NAND-flash-memory>. Accessed 24 Feb. 2024.
- [14] C. Zambelli, L. Crippa, R. Micheloni and P. Olivo, "Cross-Temperature Effects of Program and Read Operations in 2D and 3D NAND Flash Memories," 2018 International Integrated Reliability Workshop (IIRW), South Lake Tahoe, CA, USA, 2018, pp. 1-4, doi: 10.1109/IIRW.2018.8727102.
- [15] Wilcoxon, Ross. "Does a 10°C Increase in Temperature Really Reduce the Life of Electronics by Half?" *Electronics Cooling*, 20 Oct. 2017, www.electronics-cooling.com/2017/08/10c-increase-temperature-really-reduce-life-electronics-half/.
- [16] "PCIe 5.0 vs PCIe 4.0 - Everything to Know — Simms International." www.simms.co.uk, www.simms.co.uk/tech-talk/pci-50-everything-you-need-to-know/. Accessed 13 Jan. 2024.
- [17] Fio - FLEXible I/O Tester Rev.3.36 Documentation, <http://fio.readthedocs.io/en/latest/fio-doc.html>.
- [18] Intel-bigdata, <http://www.intel.com/content/dam/www/public/us/en/documents/guides/getting-started-with-hadoop-planning-guide.pdf>.
- [19] AKCP, How temperature affects it storage, <https://www.akcp.com/blog/how-temperature-affects-it-data-storage/>, 2021.
- [20] M.2 nume, <https://www.atpinc.com/products/industrial-ssds-nvme-m.2-wide-temperature>, 2020.
- [21] Softing, Pci controlnet:high-temperature network card for pci bus computers.
- [22] Why ethernet switches can take the heat (or cold), <https://iebmedia.com/technology/why-ethernet-switches-can-take-the-heat-or-cold/>, 2014.
- [23] BRIAN BEERS, P-value, <https://www.investopedia.com/terms/p/p-value.asp>, 2021.
- [24] L. S. S. Reddy B. Thirumala Rao, Scheduling data intensive workloads through virtualization on mapreduce based clouds, *International Journal of Distributed and Parallel Systems (IJDPS)*, vol. 3, 2012, pp.99-110.
- [25] Intel-bigdata, <https://www.intel.com/content/dam/www/public/us/en/documents/guides/getting-started-with-hadoop-planning-guide.pdf>.