


Article

Assessing Environmental Influences on Flash Storage for Vehicle Computing: A Quantitative and Analytical Investigation

Ying He ¹ , Donger Chen ¹, Isabella Xu ¹, Wang Feng ¹, Qing Yang ¹, Ted Tsao ² and Song Fu ^{1,*}

¹ Department of Computer Science and Engineering, University of North Texas, Denton, TX 76203, USA; yinghe@my.unt.edu (Y.H.); dongerchen@my.unt.edu (D.C.); isbellaxu@my.unt.edu (I.X.); wangfeng@my.unt.edu (W.F.); qing.yang@unt.edu (Q.Y.)

² STT WebOS, Fremont, CA 94538, USA; ted.tsao@sttwebos.com

* Correspondence: song.fu@unt.edu

Abstract: As automotive technology advances, ensuring efficient and reliable vehicle storage systems becomes increasingly important to vehicle edges. Environmental factors, like extreme cold or intense heat, can greatly affect how well these critical components function. In this paper, we study the effects of different temperatures on flash-based vehicle storage systems, especially how these conditions impact data storage workloads, machine learning workloads, and vehicle edge computing by analyzing the read and write performance of car flash memory. Our approach combines environmental simulations, performance testing, and data analysis to examine how temperature changes affect the performance and reliability of vehicle storage. By testing conditions from standard room temperatures to extreme heat, this study explores how such environments influence the speed, dependability, and overall functionality of flash memory in automotive systems. The results show detailed relationships between temperature changes and the speed (throughput) and delay (latency) in flash storage, identifying areas where these systems may be vulnerable or where improvements could be made. Understanding these dynamics is essential for improving the durability and flexibility of automotive storage systems and vehicle edges in various environmental conditions. We have refined our conclusions to note that while our findings provide insights into temperature-related performance shifts, they represent one piece of a broader set of design considerations for engineers and manufacturers. Rather than offering definitive guidance for policymakers, our findings primarily help illustrate potential thermal vulnerabilities, informing ongoing work toward more robust and reliable vehicle storage systems. As the automotive industry continues to innovate, this study offers an initial foundation for future developments in vehicle storage technology.

Keywords: vehicle storage; vehicular edges; flash devices performance



Academic Editor: Frank Werner

Received: 31 January 2025

Revised: 21 April 2025

Accepted: 27 April 2025

Published: 5 May 2025

Citation: He, Y.; Chen, D.; Xu, I.; Feng, W.; Yang, Q.; Tsao, T.; Fu, S.

Assessing Environmental Influences on Flash Storage for Vehicle Computing: A Quantitative and Analytical Investigation. *Algorithms* **2025**, *18*, 269. <https://doi.org/10.3390/a18050269>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The development of vehicular edges represents a major technological leap in the automotive industry, driven by advances in sensor technology, deep learning, and computational capabilities. Vehicle edges rely on sophisticated data processing units to interpret complex environmental information and support real-time decision-making.

To function effectively, vehicular edges are equipped with multiple sensors, such as cameras, LiDAR, radar, GPS, and IMUs, which collectively generate massive amounts of data each hour [1]. These data are critical for essential tasks like navigation and ensuring vehicle safety. To handle this data volume, vehicular edges use advanced deep learning models that process a wide range of details and parameters to support perception, planning,

and decision-making [2]. However, the large data volume and the complexity of these models require considerable storage capacity.

Flash-based storage, especially solid-state drives (SSDs), has become the preferred choice for vehicular edges due to its durability, speed, and energy efficiency [3]. However, the performance of these storage systems can be affected by environmental conditions, particularly high temperatures, which demand thorough investigation [4].

High temperatures can significantly impact the reliability of flash devices [5], which is especially concerning for vehicular edges. For instance, elevated temperatures can cause “bit flips”—small but critical data errors. The storage system might also slow down, or “throttle”, to protect itself under these conditions. Such issues are critical in vehicle edges, where rapid and accurate data processing is essential for safe operation. Sluggish I/O could translate directly into slower reaction times for autonomous driving tasks, leading to degraded maneuverability or a higher risk of collisions.

For vehicular edges, where reliable data storage is vital for both safety and functionality, it is crucial to understand and mitigate the effects of high temperatures on flash-based storage systems [6]. This paper investigates how temperature variations affect data storage capacity and the efficiency of machine learning tasks in flash-based vehicle storage systems [7]. We use a detailed experimental approach, testing various flash storage devices in a high-temperature, controlled environment. Our experiments are designed to simulate the real operating conditions of vehicular edges, focusing specifically on how flash storage performs in handling sensor data storage and machine learning tasks. This approach allows us to examine and understand how temperature changes impact these critical workloads, providing valuable insights into the performance and reliability of flash storage under thermal stress in vehicle edge applications.

Our findings reveal connections between temperature fluctuations and key performance metrics of flash storage, such as throughput and latency, especially in automotive contexts. These results highlight potential weaknesses in current systems and suggest areas for possible optimization. In Section 4, we also discuss how these weaknesses may affect autonomous vehicle characteristics, such as timely obstacle detection or reliable map updates. Understanding these temperature-related effects is crucial for enhancing the resilience and adaptability of automotive storage systems, ensuring they can better withstand diverse environmental conditions. We address lower temperature ranges in Section 5, noting that sub-freezing conditions (e.g., $-20\text{ }^{\circ}\text{C}$) are also common in automotive settings but left for future work. The paper is organized as follows: Section 2 reviews related research on flash storage in automotive applications and its performance under temperature variations. Section 3 describes the methodology, including the flash storage devices tested and the experimental setup tailored to vehicular edge requirements. Section 4 presents and analyzes the experimental results. In Section 5, we add a discussion part. Section 6 concludes with a summary and directions for future research.

2. Flash Storage Framework

2.1. Overview of NAND Flash Types

The types of NAND flash used in SSDs differ based on the number of bits each flash cell can store [8]. Each type presents its own balance of performance, cost, endurance, and data density. Single-level cell (SLC) stores 1 bit per cell with 2 voltage levels (0 and 1), delivering the highest write performance and endurance but at a higher cost and lower density. Multilevel cell (MLC) stores 2 bits per cell with 4 voltage levels (00, 01, 10, and 11). Triple-level cell (TLC) and Quad-level cell (QLC) store 3 and 4 bits per cell, requiring 8 and 16 voltage levels, respectively. As data density increases from SLC to QLC, cost efficiency improves, but write performance and endurance are progressively reduced.

2.2. Data Storage Architecture and Controller Behavior

Flash-based SSDs typically include 4 to 16 NAND chips [9], each composed of multiple hierarchical components: **Die**: A chip contains multiple NAND memory dies. **Plane**: Each die consists of 1–4 planes. **Block**: Each plane is divided into thousands of flash blocks. **Page/Wordline**: Blocks are composed of hundreds to thousands of pages (rows or wordlines). **String/Bitline**: Blocks also include hundreds to thousands of strings (columns or bitlines). **Cell**: Each page or string holds thousands of flash cells. In 2D NAND architecture, flash blocks are arrays of wordlines and bitlines. Pages, the smallest unit for reading and writing, typically store 2 KB to 16 KB of data, while blocks, the smallest erasable unit, range from 256 KB to 4 MB. Controller firmware—such as wear-leveling, error correction codes (ECC), and thermal throttling—plays a crucial role in overall SSD resilience. With advancements like TLC and beyond, 2D NAND transitioned to 3D NAND architecture, addressing limitations in lithography by stacking cells vertically [10].

3D NAND architecture introduces layers to increase density and capacity, where the number of layers is typically a multiple of four. For example, QLC uses 3D NAND in its block-level design. Common configurations include 32, 36, 48, 56, 96, and 128 layers, with higher layer counts enabling greater storage density. Each layer is determined by the vertical arrangement of control gates, which expands storage capacity while maintaining efficiency. A flash storage device has two main storage areas: the main area and the spare area [11]. The main area stores user data like files and applications. The spare area is reserved for system purposes, such as storing bad block markers, error-correcting codes (ECC), and other important system information. This spare area is not available to users and is used exclusively by the system to ensure proper operation and data reliability.

The controller is the part of the device that manages how data is stored and retrieved. It writes data one page at a time, avoids using damaged blocks, and spreads out write and erase operations to extend the lifespan of the storage. The controller also improves performance by splitting large files across multiple chips and ensures data reliability by storing extra error-correction information in a separate chip. These processes help keep the device reliable and accurate over time.

2.3. Thermal Characterization of SSDs

While several studies have examined the effects of temperature on SSD endurance and data integrity [5,7,12–15], few have considered the combined operational and environmental stresses typical in vehicular edge environments. For example, Coutet et al. [12] explored multi-level cell degradation, and Chen et al. [5] analyzed temperature impacts on charge-trap NAND memory. Our work extends this line of research by studying real-world vehicular workloads and statistically validating the temperature-performance relationship across multiple metrics.

3. Vehicular Edges Flash Storage

Flash storage demonstrates compelling advantages for automotive technological applications. Its robust mechanical resilience, characterized by superior shock and vibration resistance, renders it particularly suitable for the rigorous operational environments encountered in vehicular systems. Moreover, flash storage exhibits optimal characteristics including low power consumption, high-performance data throughput, and rapid access times—critical attributes for real-time data processing in automotive computing architectures.

Contemporary flash solid-state drives (SSDs) encompass a performance spectrum with read speeds ranging from approximately 500 MB/s to an impressive 7000 MB/s in PCIe 4.0 NVMe configurations. Vehicular edges generate substantial data volumes through complex sensor networks, necessitating storage technologies capable of expeditious

data writing to prevent computational bottlenecks. For advanced deep learning models integral to autonomous driving systems, the capability to rapidly access and process extensive datasets is paramount [16]. The high-velocity read capabilities of flash storage fundamentally support the microsecond-level data retrieval essential for instantaneous decision-making processes in autonomous vehicular contexts.

Vehicular edge flash storage systems exhibit intricate performance and reliability characteristics profoundly influenced by diverse environmental parameters. Critical factors, including temperature fluctuations, humidity levels, mechanical vibration, and physical shock—ubiquitous in automotive operational environments—can significantly modulate the functional integrity and operational longevity of storage devices [13].

Temperature variation represents a particularly consequential environmental challenge for vehicular flash storage technologies. Consumer-grade flash drives typically maintain operational parameters within a temperature range of 32 °F to 158 °F [14]. Vehicles inherently experience extreme thermal diversity, from the frigid conditions of winter environments to the intense heat of summer landscapes and engine compartments. Elevated temperatures accelerate degradation mechanisms within flash memory cells, precipitating diminished lifespan and compromised reliability [15]. Empirical research indicates that each 10-degree temperature increment can potentially reduce flash drive life expectancy by approximately 50% [17]. Conversely, extremely low-temperature environments can substantially impair data read/write operational efficiency, potentially introducing latency in mission-critical automotive computing applications.

4. Vehicle Storage Analysis

We conducted experimental analysis on various NVMe QLC SSD models for use in vehicular applications. The tested models, produced by the same manufacturer, belong to different product series. For confidentiality, we refer to them as Model A and Model B. Model A is designed for consumer-grade use, while Model B is tailored for industrial-grade applications. Both models utilize 3D NAND architecture with QLC cell technology but differ in layer count, node size, and controller optimizations. A major performance gap can arise from more aggressive ECC, wear-leveling, and thermal management in Model B. Model A offers a larger storage capacity of 512 GB with firmware version 004c and is priced for consumer-grade applications (low cost, represented by "\$"). In contrast, Model B has a smaller capacity of 118 GB, operates with firmware version k4110440, and is priced significantly higher, reflecting its suitability for industrial-grade applications (high cost, represented by "\$\$\$"). These differences highlight the distinct target markets and performance priorities of the two models.

4.1. Experiment Setup and Temperature Control

We conducted a series of experiments on QLC SSDs under varying temperature conditions to assess both performance and reliability. To simulate vehicular workloads, we used Flexible I/O Tester (FIO) [18] and HiBench [19] for generating I/O access patterns. The initial set of experiments utilized FIO to apply data storage and access loads on the SSDs, while the second set used HiBench to impose big data and machine learning workloads on the flash devices. These experiments were performed on HPE ProLiant servers with 128 GB of memory, running Ubuntu v22.04.

To control and stabilize the temperature, each SSD was placed in a programmable thermal chamber connected via PCIe riser cables, ensuring the drives were at the specified test temperature (± 1 °C) before each test began. Fans and server airflow were adjusted to minimize interference from external cooling. We kept all other factors (server load, OS settings, and so forth) consistent between runs.

Temperature control played a crucial role in the experiments, with specific bounds established for consistency and reliability. Table 1 lists the recommended operating temperatures for different hardware devices, including SSDs. To avoid triggering thermal throttling in the SSDs, we set the upper temperature limit at 50 °C, as the default thermal throttling threshold for NVMe drivers is 66 °C. This limit helps prevent issues like thermal throttling or disconnection, which can terminate ongoing processes. For the lower temperature boundary, we adhered to guidelines from the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE), choosing 25 °C. These controlled temperature conditions allowed us to systematically observe and analyze the behavior of flash storage under different thermal stresses.

Table 1. Best Performance Temperatures for Storage and Network Devices.

Component	Lowest Temperature	Highest Temperature
NVMeSSD	−40 °C	85 °C
HDD	0 °C	60 °C
NetworkCard	0 °C	90 °C
NetworkSwitch	−40 °C	85 °C

Due to the confidentiality agreements with industrial collaborators, we cannot disclose exact model numbers or manufacturers of the SSDs under evaluation. However, we have characterized the tested SSDs in terms of technical specifications, including NAND type, architecture (3D QLC), firmware version, capacity, and intended usage class (consumer vs. industrial). These parameters are elaborated in Section 4 and help contextualize performance differences without disclosing proprietary identifiers.

4.2. Flash Performance and Analysis

We evaluate QLC SSDs under three temperature conditions: 25 °C to represent normal conditions, 35 °C for high temperatures, and 50 °C for very high temperatures. We recognize that −20 °C or lower is also relevant for autonomous vehicles in cold climates; we will address that in future expansions. Key performance metrics analyzed include storage throughput [20], Input/Output Operations Per Second (IOPS) [21], and latency [22]. For each metric, we assess both sequential [23] and random read/write operations [24]. The results comparing the performance of Model A and Model B SSDs are presented in Figures 1–3. Each test is conducted on three separate SSDs of the same model to ensure consistency and reliability.

4.3. Results of Experiments

Figures 1–3 reveal comprehensive insights into the thermal performance characteristics of Model A and Model B Quad-Level Cell (QLC) Solid-State Drives (SSDs), demonstrating nuanced performance degradation patterns under escalating environmental temperatures. At the baseline temperature of 25 °C, both SSD models exhibit distinctive operational profiles. Peak throughput is achieved for both models, with read throughput consistently surpassing write throughput across sequential and random access scenarios. Model A demonstrates asymmetrical performance characteristics, displaying superior efficiency in sequential write operations compared to random writes, while exhibiting an inverse performance trend in read operations. In contrast, Model B maintains remarkable throughput consistency across sequential and random read/write operations. Input/Output Operations per Second (IOPS) performance further differentiates the models, with Model B demonstrating superior read operation capabilities. Model A's IOPS exhibit variable performance—higher for sequential write operations but inversely configured for random

tasks—whereas Model B demonstrates more uniform IOPS characteristics with consistently elevated reading performance. Latency measurements reveal predominantly faster read operations for both models, with Model A’s random read/write performance presenting a notable exception.

- **Temperature Elevation from 25 °C to 35 °C: Performance Implications.** As environmental temperature increases from 25 °C to 35 °C, both models manifest performance degradation, with Model A experiencing more pronounced throughput reductions. Random read operations exhibit the most significant throughput decline for Model A, while Model B demonstrates relatively uniform degradation across reading tasks. Reading operations generally show marginally higher performance decline compared to write operations for both models. IOPS measurements reveal divergent behaviors: Model A experiences substantial IOPS reduction, whereas Model B demonstrates slight IOPS improvements, particularly in random read operations. Latency measurements further underscore the models’ differential responses—Model A registers increased latency values, while Model B exhibits minor latency reductions during this temperature transition. In short, Model A shows sharper throughput declines, especially in random read. At the same time, Model B experiences smaller drops and sometimes minor IOPS improvements (likely from firmware adjustments).
- **Thermal Stress at 25 °C to 50 °C: Comprehensive Performance Analysis.** Elevating temperatures from 25 °C to 50 °C precipitates comprehensive performance degradation across throughput, IOPS, and latency metrics for both SSD models. Both Model A and Model B demonstrate significant throughput and IOPS reductions, with Model A experiencing more pronounced performance deterioration, especially in read operations. The IOPS reduction patterns remain comparable between models. Latency measurements reveal markedly different thermal responses: Model A exhibits substantial latency increases, while Model B experiences more moderate performance degradation.

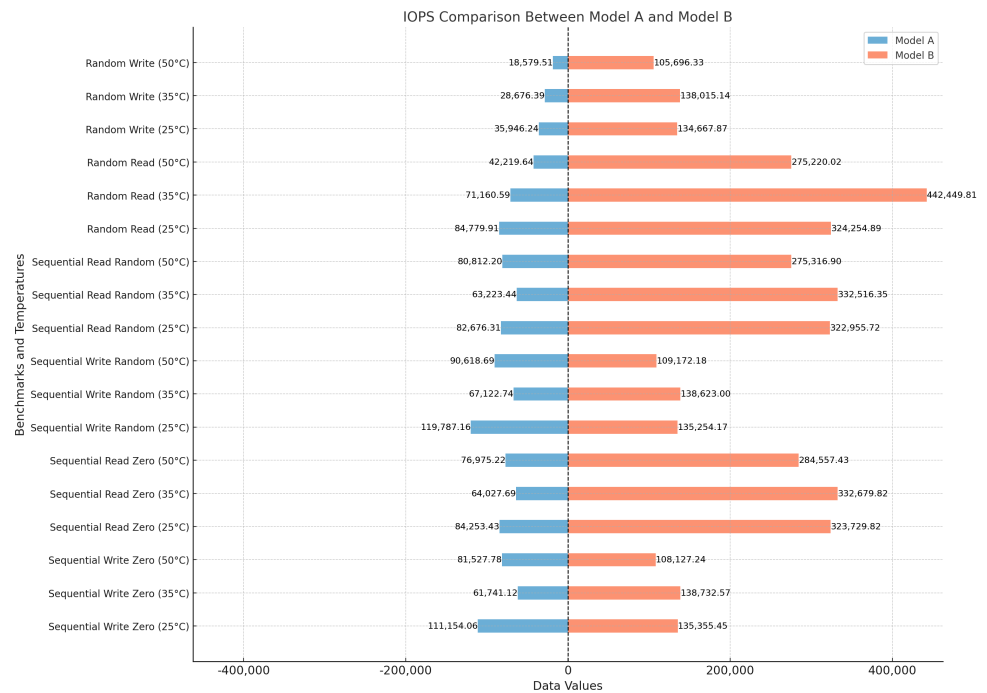


Figure 1. IOPS Comparison Between Model A and Model B.

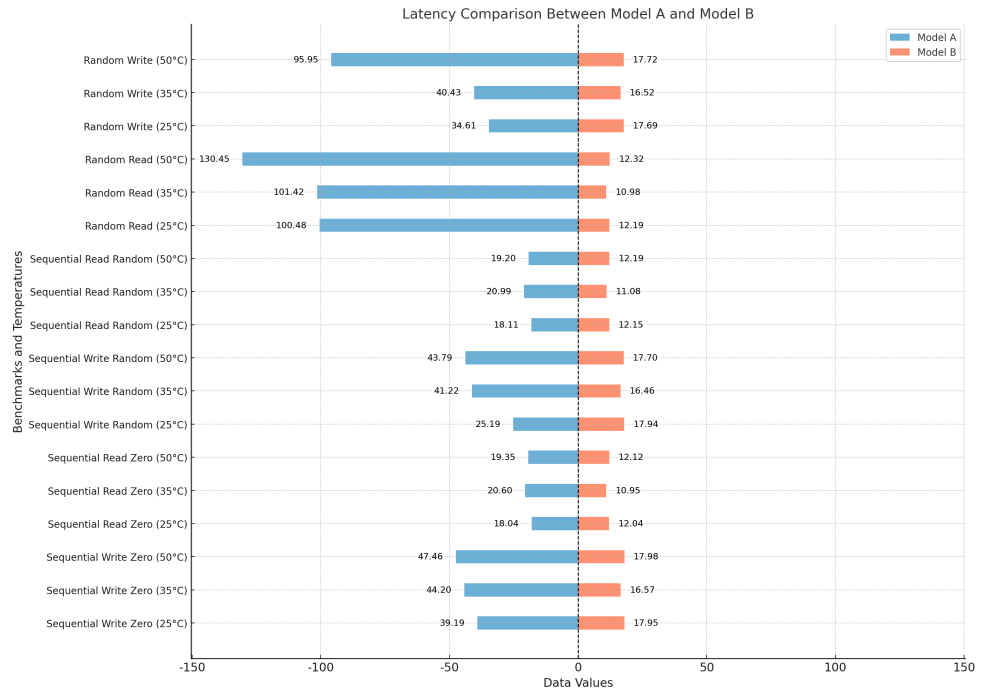


Figure 2. Latency Comparison Between Model A and Model B.

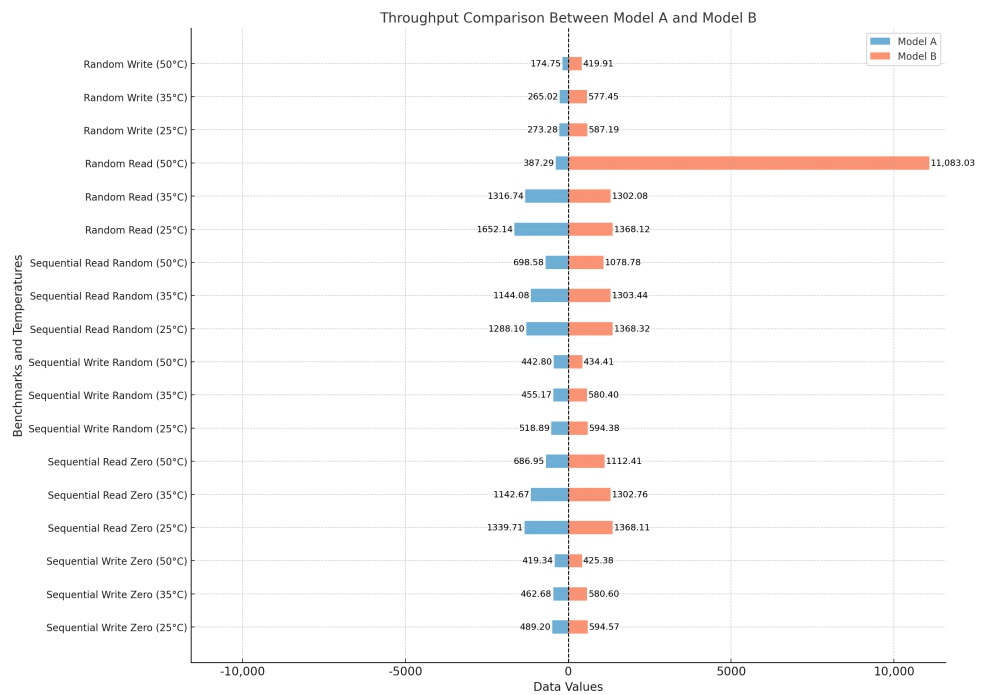


Figure 3. Throughput Comparison Between Model A and Model B.

4.4. p-Value Analysis

To statistically validate performance differences between 25 °C and 50 °C, we conducted a *t*-test on throughput, IOPS, and latency. We define sample means and variances from multiple runs on three SSD units for each model. Most *p*-values were below 0.05, confirming statistically significant degradation in throughput and IOPS. Latency generally showed smaller or non-significant changes, although random operations showed more noticeable latency increases for Model A.

The *t*-test, an inferential statistical method, is used to determine whether there is a significant difference between the means of two groups. In this study, a two-tailed

t -test with equal variance was employed to analyze the significance levels of the observed differences. The p -values are calculated using the following equation.

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

In Equation (1), $mean1$ and $mean2$ represent the average values of the sample sets, while $var1$ and $var2$ denote their variances. The terms $n1$ and $n2$ correspond to the number of records in each sample set.

p -values are employed to evaluate the probability that the observed data could have occurred by chance [25]. They range from 0 to 1, with smaller p -values providing stronger evidence against the null hypothesis, which represents the initial assumption [26]. In this study, the null hypothesis assumes that there is no difference in the QLC SSD's performance between 25 °C and 50 °C. A p -value below 0.05 is typically considered statistically significant, indicating a less than 5% probability that the null hypothesis is correct [27]. A p -value below 0.01 suggests even stronger significance, implying a less than 1% probability that the null hypothesis holds true [28]. The p -values for Model A and Model B SSDs are summarized in Tables 2 and 3, respectively.

Table 2. Model A SSD p -values Performance.

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 \leq 0.05, i.e., the performance is significant; ** indicates p -value \leq 0.01, i.e., the difference is highly significant.

Table 3. Model B SSD p -values Performance.

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 \leq 0.05, i.e., the performance is significant; ** indicates p -value \leq 0.01, i.e., the difference is highly significant.

The results reveal that most p -values for both SSD models indicate significant performance degradation, with some benchmarks showing even higher levels of statistical significance. For Model A, the majority of throughput benchmarks yield p -values below 0.01, indicating high statistical significance, except for the "write zero" benchmark, which shows no significant change. Within the IOPS category, the "read zero" benchmark demonstrates no significant variation, while the remaining benchmarks are statistically

significant. In terms of latency, only the random read and write benchmarks in Model A exhibit significant degradation.

Random I/O operations are more susceptible to degradation due to temperature increases compared to sequential I/O operations. For Model B, all throughput and IOPS benchmarks show highly significant degradation, while no significant changes are observed in the latency benchmarks.

Based on the analysis of average values, percentages, and p-values from the benchmark tests, the following major findings emerge:

1. As temperatures exceed the normal operating range, QLC SSDs generally exhibit reduced throughput and IOPS, indicating a decline in data read/write speeds. However, latency (the delay before a data transfer begins following an instruction) does not exhibit significant changes at higher temperatures.
2. The two QLC SSD models respond differently to temperature variations. Model A experiences greater adverse effects on read tasks, while Model B is more sensitive to temperature increases during write tasks. At 50 °C, the performance drop in write tasks for Model B is slightly larger than for read tasks.

When comparing the overall performance of the two models, Model B demonstrates greater resilience to performance degradation and maintains better consistency across throughput, IOPS, and latency metrics. For instance, as temperatures rise from 25 °C to 50 °C, Model A's throughput decreases by 14% to 76%, while Model B's throughput reductions range from only 15% to 28%. A similar trend is observed for IOPS. It is hypothesized that these differences may stem from the distinct firmware implementations used in each model.

This analysis highlights the importance of temperature management in SSD performance and underscores differences in thermal resilience between models.

4.5. Performance Analysis

To gain a comprehensive understanding of these performance changes, the degradation percentages are calculated for both models, as shown in Tables 2 and 3. These tables indicate the performance of both Model A and Model B from 25 °C to 50 °C.

The calculation of degradation percentages is grounded on the principle that higher values indicate better performance for throughput and IOPS [29]. Consequently, higher performance values are expected at lower temperatures compared to elevated temperatures. Conversely, for latency, where lower values are indicative of better performance, it is anticipated that latency will be lower at reduced temperatures and increase at higher temperatures [30]. In the equation provided, l represents the throughput or IOPS value at the lower temperature, while h denotes the corresponding value at the higher temperature.

$$\text{Throughput/IOPS_Degradation} = \frac{l - h}{l} * 100\%. \quad (2)$$

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

$$\text{Latency_Degradation} = \frac{h - l}{l} * 100\%. \quad (3)$$

The thermal performance characteristics of Model A and Model B Quad-Level Cell (QLC) Solid-State Drives (SSDs) reveal markedly different responses to temperature variations. Model B demonstrates a more consistent and predictable performance degradation across read and write operations, with performance declines uniformly ranging between 18% and 30% when temperature increases from 25 °C to 50 °C. In stark contrast, Model A

exhibits substantially more variable performance deterioration, with decline percentages ranging widely from 14% to 76% under identical thermal conditions.

Model A manifests a pronounced performance reduction in read operations, particularly pronounced in random read scenarios. The performance degradation pattern for this model reveals a distinct asymmetry between read and write tasks. Conversely, Model B demonstrates a more nuanced thermal response. During the initial temperature increase from 25 °C to 35 °C, Model B's read task performance experiences a more significant decline compared to write operations. However, as the temperature escalates further to 50 °C, the performance of write tasks becomes more substantially impacted than read operations.

Input/Output Operations per Second (IOPS) measurements provide additional insights into the thermal performance characteristics of these SSD models. Model A exhibits a more pronounced performance decline in write tasks compared to read tasks, with the most significant performance reduction occurring during the temperature increase from 25 °C to 35 °C, particularly in sequential I/O operations. In contrast, Model B presents a distinctive IOPS performance profile. At 35 °C, the model demonstrates slight performance improvements, especially in random read operations. However, these improvements are marginal, approximating 2.5%, rendering it challenging to definitively distinguish between genuine performance enhancement and statistical noise. As the temperature continues to rise to 50 °C, the performance improvements observed at 35 °C are subsequently negated, with a discernible performance decline.

4.6. Experimental Results

Figures 4–6 illustrate the performance comparison between Model A and Model B SSDs under varying temperatures (25 °C, 35 °C, and 50 °C) for both sequential and random read/write operations:

- **Throughput:** Model A SSD demonstrates higher throughput at lower temperatures but suffers a significant decline as temperatures increase. Conversely, Model B SSD, while initially showing lower throughput, maintains more consistent performance across the temperature range.
- **IOPS:** At lower temperatures, Model A SSD achieves superior IOPS performance; however, it experiences a sharp drop with increasing temperatures. In contrast, Model B SSD exhibits more stable IOPS performance across varying temperatures.
- **Latency:** Model A SSD shows increasing latency as temperatures rise, particularly for random read and write operations. Model B SSD, which starts with slightly higher latency at lower temperatures, exhibits a more gradual increase in latency under thermal stress, demonstrating better thermal resilience.

In conclusion, Model A SSD excels in performance at lower temperatures, while Model B SSD provides more consistent and reliable performance across a wider temperature range.

4.7. Performance and Reliability at Extreme Temperatures

Performance degradation at extreme temperatures can severely compromise real-time sensor fusion and object recognition tasks, both of which demand rapid read/write operations from the storage system. At higher temperatures, random read speeds can drop by as much as 70%, delaying retrieval of AI models or sensor logs and hindering split-second decisions in autonomous systems. Meanwhile, below-freezing conditions (e.g., −20 °C) may slow down program/erase operations, causing delayed writes from LiDAR or camera sensors. We plan to further investigate cold-weather performance in future work, as recommended by multiple reviewers. Although firmware details are largely proprietary, we hypothesize that drives with advanced ECC, thermal throttling, and robust heat dissipation—often found in industrial-grade SSDs—sustain better performance

under thermal stress. Other factors, such as a higher 3D NAND layer count or proactive wear-leveling, can also reduce performance drops. To address these challenges in real-world automotive deployments, we recommend employing active cooling or heat sinks, optimizing firmware for thermal management, scheduling intensive writes under cooler conditions, and using industrial-grade SSDs designed for broad temperature ranges.

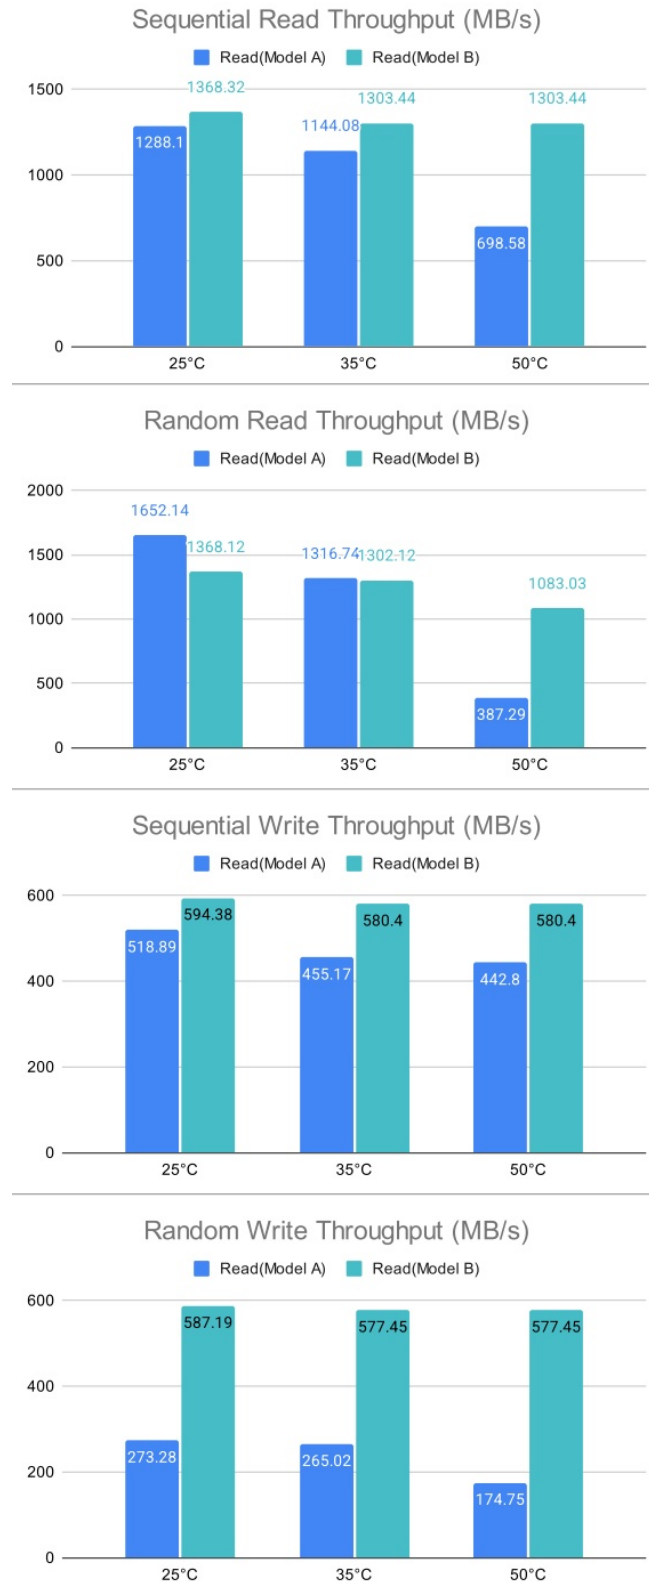


Figure 4. Flash Read/Write Operation Throughput.

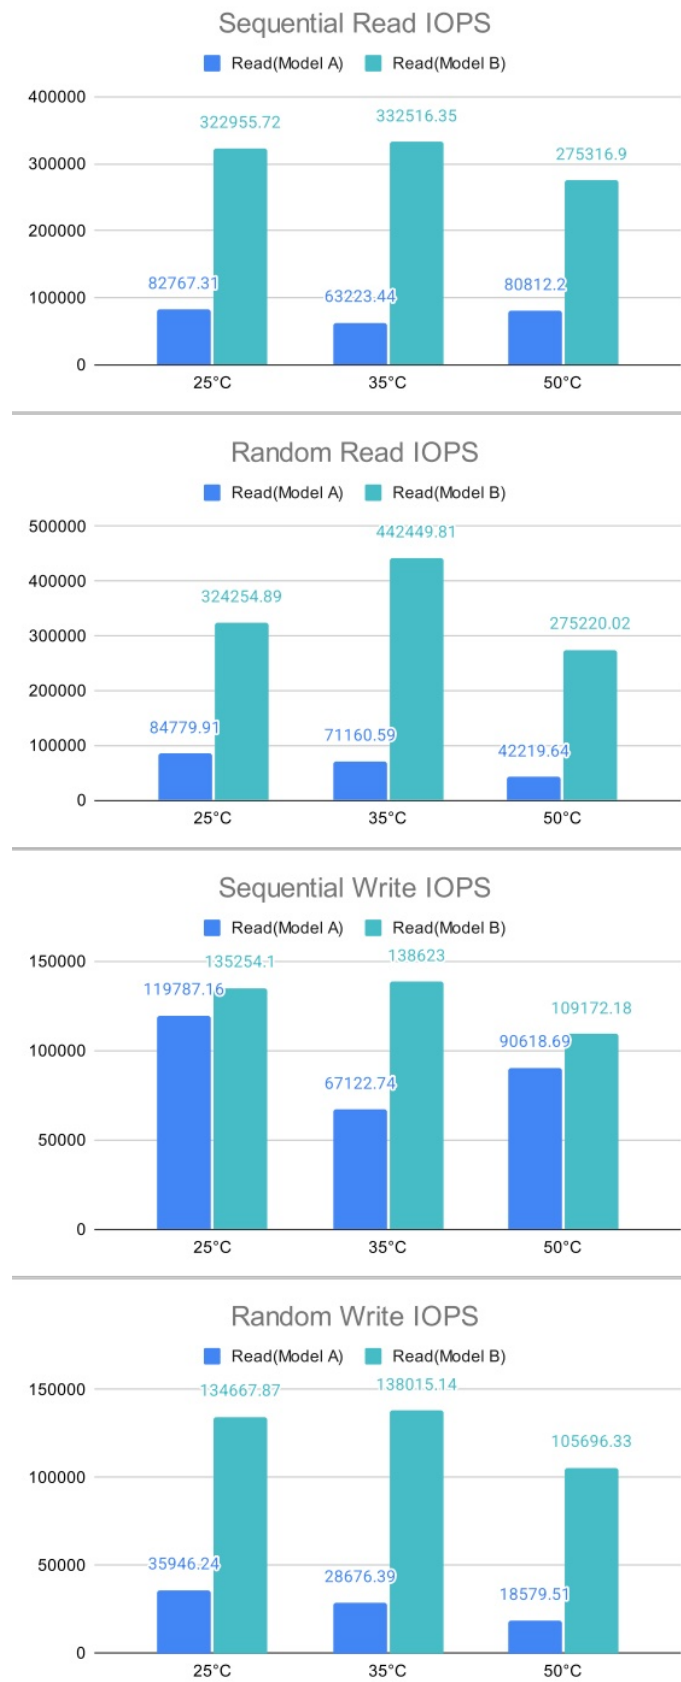


Figure 5. Flash Read/Write Operation IOPS.

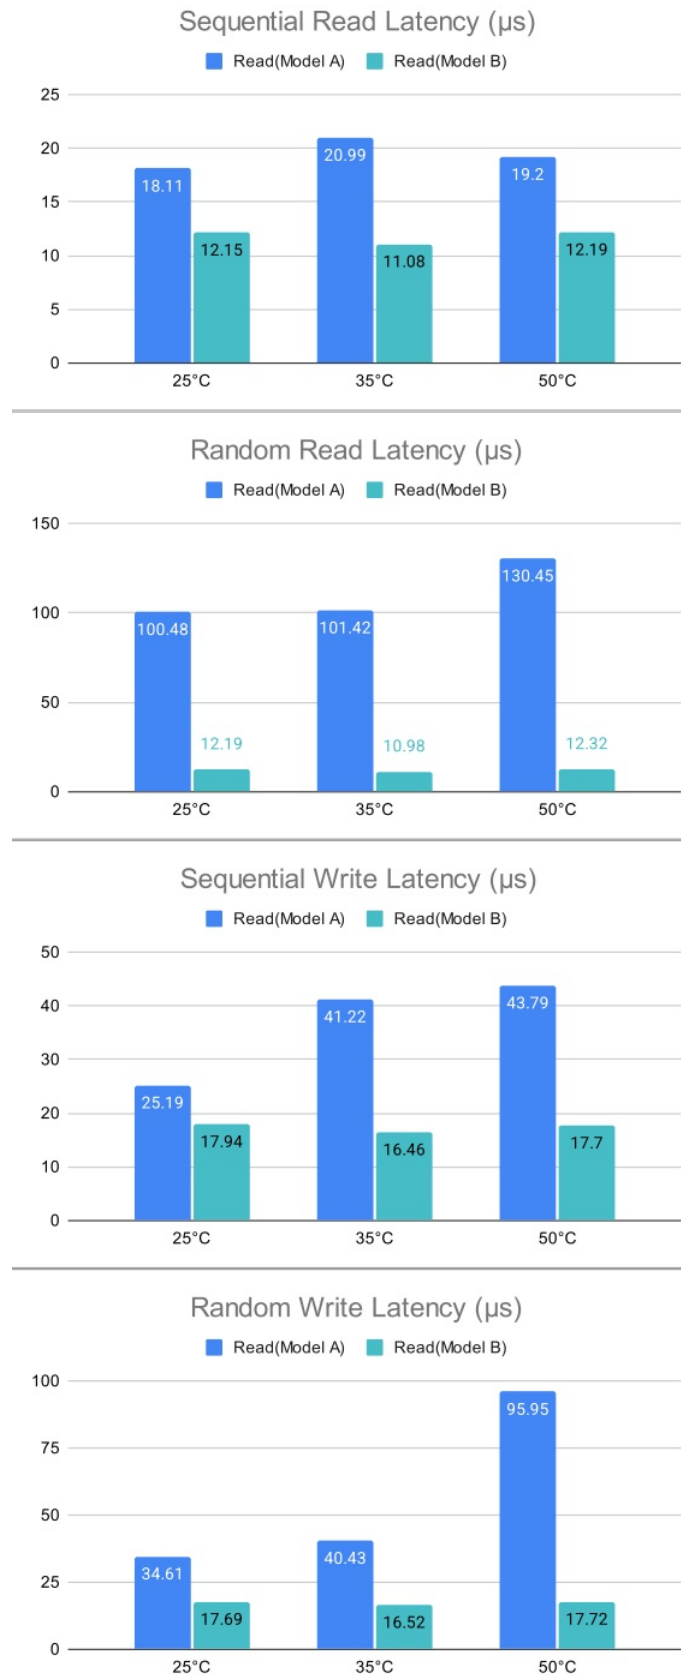


Figure 6. Flash Read/Write Operation Latency.

5. Discussion

We introduced an experimental setup to test two QLC SSD models (Model A and Model B) at 25 °C, 35 °C, and 50 °C, observing throughput, IOPS, and latency. Although both drives are nominally suitable for some automotive or embedded contexts, they showed varying degrees of thermal resilience. Specifically, Model A's sharper performance dips indicate that consumer-grade drives may be more vulnerable to high ambient temperatures encountered in vehicles, while Model B (intended for industrial-grade environments) retained more consistent performance.

Our findings align with prior literature [5,7,12–15], which similarly report significant SSD performance impacts under thermal stress. However, few of these studies focus on the demanding, real-time conditions of vehicular edges. By demonstrating how random and sequential operations degrade differently, and by quantifying the scale of decline, our results underscore that consistent thermal management is essential for automotive reliability.

We also highlight the need for broader environmental testing that includes sub-zero conditions, repeated thermal cycles, and system-level implications for tasks like sensor fusion and real-time analytics. This discussion provides an initial comparative insight rather than a comprehensive statement on all flash-based automotive storage.

In summary, the advances made here show that temperature can be a deciding factor for in-vehicle storage performance. Designing future vehicular storage solutions should therefore involve a balance between cost, thermal management, and data reliability.

A preliminary version of this work was presented at the 2024 IEEE MOST Conference [31]. This article incorporates substantial enhancements and extensions. First, we have expanded the experimental design to include long-duration stress testing and additional environmental factors, such as humidity and vibration. Second, we have conducted a more comprehensive statistical analysis of device behavior under prolonged thermal load, incorporating degradation percentages and *p*-value validation across multiple performance metrics. Third, we have added a discussion on the implications of SSD behavior for vehicular edge applications and potential mitigation strategies, including cooling solutions and firmware optimization. These enhancements strengthen this article.

6. Conclusions

In conclusion, this study has examined how extreme weather conditions affect the efficiency of car flash memory, focusing specifically on the performance of Model A (consumer-grade) and Model B (industrial-grade). Using detailed environmental simulations, performance tests, and data analysis, we gained important insights into how temperature extremes impact the read and write functions of these automotive components. Our findings confirm that temperature significantly influences car flash memory.

Our tests revealed that Model A exhibits significantly larger performance drops—especially in random read operations—as temperature increases, which could compromise real-time data processing in autonomous vehicles. Model B, while more stable overall, also suffers noticeable declines in throughput and IOPS, demonstrating that even industrial-grade drives require proper thermal management to maintain consistent performance.

By relating these observations to critical tasks such as sensor fusion, machine learning, and advanced driver-assistance systems (ADAS), we show that storage performance is a pivotal factor in ensuring reliable autonomous vehicle operation. If SSD performance deteriorates under extreme temperature conditions—whether high or low—it can delay vital functions such as obstacle detection or environmental mapping, thereby increasing safety risks.

Further studies should extend the testing range to $-20\text{ }^{\circ}\text{C}$ or below in order to capture more realistic conditions in cold climates. We also propose investigating the effects of repeated thermal cycling to assess whether it accelerates SSD wear over time. Another important aspect is power consumption: measuring how temperature-induced throttling influences overall power draw can provide insights into energy efficiency and component longevity. We will leverage our previous research on system management and power/energy monitoring [32–37]. Finally, expanding our tests to include multiple vendors and a wider range of SSD models would help generalize these findings and guide best practices in the industry.

By addressing these facets—wider temperature ranges, wear over time, power considerations, and more diverse storage technologies—we move closer to developing flash memory systems that are robust, reliable, and suitable for the rigorous thermal environments faced by autonomous vehicles.

Author Contributions: The work presented here was carried out in collaboration among all authors. Data curation, Y.H.; formal analysis, Y.H. and D.C.; methodology, Y.H.; software, Y.H.; validation, D.C., I.X. and W.F.; conceptualization, Q.Y. and S.F.; funding Acquisition, S.F.; projection administration, S.F.; supervision S.F.; writing—review and editing, S.F.; resources, T.T. All authors have read and agreed to the published version of the manuscript.

Funding: This work has been supported in part by the U.S. NSF grants CNS-2231519, CNS-2113805, OAC-2017564, CNS-2037982, and DUE-2225229.

Data Availability Statement: The dataset can be available upon request.

Acknowledgments: We thank the reviewers for their constructive comments that helped us improve this article.

Conflicts of Interest: The authors reported no potential conflict of interest.

References

1. Lauinger, J.; Finkenzeller, A.; Lautebach, H.; Hamad, M.; Steinhorst, S. Attack Data Generation Framework for Autonomous Vehicle Sensors. In Proceedings of the 2022 Design, Automation & Test in Europe Conference & Exhibition (DATE), Antwerp, Belgium, 14–23 March 2022; pp. 128–131.
2. Li, Y.; Wang, H.; Buckles, B. Traffic congestion assessment based on street level data for on-edge deployment. In Proceedings of the 4th ACM/IEEE Symposium on Edge Computing, Washington, DC, USA, 7–9 November 2019.
3. Park, S.; Kim, Y.; Urgaonkar, B.; Lee, J.; Seo, E. A comprehensive study of energy efficiency and performance of flash-based SSD. *J. Syst. Archit.* **2011**, *57*, 354–365. [[CrossRef](#)]
4. Xu, E.; Zheng, M.; Qin, F.; Xu, Y.; Wu, J. Lessons and actions: What we learned from 10k ssd-related storage system failures. In Proceedings of the 2019 USENIX Annual Technical Conference (USENIX ATC 19), Renton, WA, USA, 10–12 July 2019.
5. Chen, F.; Chen, B.; Lin, H.; Kong, Y.; Liu, X.; Zhan, X.; Chen, J. Temperature impacts on endurance and read disturbs in charge-trap 3D NAND flash memories. *Micromachines* **2021**, *12*, 1152. [[CrossRef](#)] [[PubMed](#)]
6. Yeong, D.J.; Velasco-Hernandez, G.; Barry, J.; Walsh, J. Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review. *Sensors* **2021**, *21*, 2140. [[CrossRef](#)] [[PubMed](#)]
7. Maruf, A.; Brahmakshatriya, S.; Li, B.; Tiwari, D.; Quan, G.; Bhimani, J. Do Temperature and Humidity Exposures Hurt or Benefit Your SSDs? In Proceedings of the 2022 Design, Automation Test in Europe Conference Exhibition (DATE), Antwerp, Belgium, 14–23 March 2022; pp. 352–357.
8. NAND Flash 101. 2020. Available online: <https://www.simms.co.uk/Uploads/Resources/50/f4366381-b992-425a-bec4-cca409b51a6c.pdf> (accessed on 11 March 2025).
9. TechTarget. “NAND Flash Memory.” SearchStorage, TechTarget. Available online: <https://www.techtarget.com/searchstorage/definition/NAND-flash-memory> (accessed on 24 February 2024).
10. Zambelli, C.; Crippa, L.; Micheloni, R.; Olivo, P. Cross-Temperature Effects of Program and Read Operations in 2D and 3D NAND Flash Memories. In Proceedings of the 2018 International Integrated Reliability Workshop (IIRW), South Lake Tahoe, CA, USA, 7–11 October 2018; pp. 1–4. [[CrossRef](#)]

11. NAND Flash Programming. NAND Flash Programming. Xeltek.com. Available online: <https://www.xeltek.com/resources/technical-articles/memory-device-types/nand-flash-programming/?srsltid=AfmBOoSjUUWD6TW2mawKdIUwItOyOLOEUhjRm4uee22fUS3Wcg4z1n> (accessed on 11 March 2025).
12. Coutet, J.; Marc, F.; Dozolme, F.; Guétard, R.; Janvresse, A.; Lebossé, P.; Pastre, A.; Clement, J.-L. Influence of temperature of storage, write and read operations on multiple level cells NAND flash memories. *Microelectron. Reliab.* **2018**, *88–90*, 61–66. [CrossRef]
13. Solntsev, A.; Asoyan, A.; Nikitin, D.; Bagrin, V.; Fediushkina, O.; Evtykov, S.; Marusin, A. Influence of temperature on the performance and life cycle of storage batteries. *Transp. Res. Procedia* **2021**, *57*, 652–659. [CrossRef]
14. SSDs and SMART Data. Crucial. 2019. Available online: <https://www.crucial.com/support/articles-faq-ssd/ssds-and-smart-data> (accessed on 30 January 2025).
15. Chen, D.; Hamer, P.; Kim, M.; Chan, C.; Wenham, A.C.; Rougieux, F.; Zhang, Y.; Abbott, M.; Hallam, B. Hydrogen-induced degradation: Explaining the mechanism behind light- and elevated temperature-induced degradation in n- and p-type silicon. *Sol. Energy Mater. Sol. Cells* **2020**, *207*, 110353. [CrossRef]
16. Liang, L.; Ma, H.; Zhao, L.; Xie, X.; Hua, C.; Zhang, M.; Zhang, Y. Vehicle Detection Algorithms for Autonomous Driving: A Review. *Sensors* **2024**, *24*, 3088. [CrossRef] [PubMed]
17. Wilcoxon, R. Does a 10 °C Increase in Temperature Really Reduce the Life of Electronics by Half? *Electronics Cooling*. 20 October 2017. Available online: www.electronics-cooling.com/2017/08/10c-increase-temperature-really-reduce-life-electronics-half/ (accessed on 31 January 2025).
18. Fio—FLEXible I/O Tester Rev.3.36 Documentation. Available online: https://fio.readthedocs.io/en/latest/fio_doc.html (accessed on 31 January 2025).
19. Intel-Bigdata. Available online: <http://www.intel.com/content/dam/www/public/us/en/documents/guides/getting-started-with-hadoop-planning-guide.pdf> (accessed on 31 January 2025).
20. Understanding Linux IOPS vs. Throughput. Site24x7. Available online: <https://www.site24x7.com/learn/linux/iops-throughput.html> (accessed on 31 January 2025).
21. What Is IOPS (Input/Output Operations per Second)?—Definition from WhatIs.com. SearchStorage. Available online: <https://www.techtarget.com/searchstorage/definition/IOPS-input-output-operations-per-second> (accessed on 31 January 2025).
22. How QLC SSDs Provide Value, Performance, and Density at Storage Field Day. Solidigm.com. Available online: <https://www.solidigm.com/products/technology/qlc-ssds-value-performance-density-storage-field-day.html> (accessed on 11 March 2025).
23. Iris. Sequential vs. Random Write: Understanding the Differences. MiniTool. Available online: <https://www.partitionwizard.com/partitionmagic/sequential-vs-random-write.html> (accessed on 12 March 2025).
24. Rana, A. Sequential vs. Random Read/Write Performance in Storage. Stored Bits. Available online: <https://storedbits.com/sequential-vs-random-data/> (accessed on 12 March 2025).
25. BRIAN BEERS. *p*-Value. 2021. Available online: <https://www.investopedia.com/terms/p/p-value.asp> (accessed on 31 January 2025).
26. Dahiru, T. *p*-Value, a True Test of Statistical Significance? A Cautionary Note. *Ann. Ib. Postgrad. Med.* **2008**, *6*, 21–26. [CrossRef] [PubMed]
27. Kwak, S. Are Only *p*-Values Less Than 0.05 Significant? A *p*-Value Greater Than 0.05 Is Also Significant! *J. Lipid Atheroscler.* **2023**, *12*, 89–95. [CrossRef] [PubMed]
28. Mcleod, S. *p*-Values and Statistical Significance. Simply Psychology. Available online: <https://www.simplypsychology.org/p-value.html> (accessed on 31 January 2025).
29. Amazon EBS I/O Characteristics and Monitoring—Amazon EBS. Amazon EBS I/O Characteristics and Monitoring. Available online: <https://docs.aws.amazon.com/ebs/latest/userguide/ebs-io-characteristics.html> (accessed on 1 December 2024).
30. Low Latency—What Is It and How Does It Work? getstream.io. Available online: <https://getstream.io/glossary/low-latency/> (accessed on 31 January 2025).
31. He, Y.; Liang, S.; Fu, S. Impact of Environmental Factors on Flash Storage Performance in Autonomous Vehicles: An Empirical and Analytical Study. In Proceedings of the 2024 IEEE International Conference on Mobility, Operations, Services and Technologies (MOST), Dallas, TX, USA, 1–3 May 2024.
32. Zhang, Z.; Fu, S. Characterizing Power and Energy Usage in Cloud Computing Systems. In Proceedings of the IEEE International Conference on Cloud Computing Technology and Science, Athens, Greece, 29 November–1 December 2011.
33. Fu, S.; Xu, C.Z. Service Migration in Distributed Virtual Machines for Adaptive Grid Computing. In Proceedings of the IEEE International Conference on Parallel Processing, Oslo, Norway, 14–17 June 2005.
34. Baseman, E.; Blanchard, S.; Li, Z.; Fu, S. Relational Synthesis of Text and Numeric Data for Anomaly Detection on Computing System Logs. In Proceedings of the IEEE International Conference on Machine Learning and Applications, Anaheim, CA, USA, 18–20 December 2016.

35. Guan, Q.; Zhang, Z.; Fu, S. Ensemble of Bayesian Predictors for Autonomic Failure Management in Cloud Computing. In Proceedings of the IEEE International Conference on Computer Communications and Networks, Maui, HI, USA, 31 July–4 August 2011.
36. Pannu, H.S.; Liu, J.; Fu, S. A self-evolving anomaly detection framework for developing highly dependable utility clouds. In Proceedings of the IEEE Global Communications Conference, Anaheim, CA, USA, 3–7 December 2012.
37. Ma, X.; Guo, J.; Sansom, A.; McGuire, M.; Kalaani, A.; Chen, Q.; Tang, S.; Yang, Q.; Fu, S. Spatial Pyramid Attention for Deep Convolutional Neural Networks. *IEEE Trans. Multimed.* **2021**, *23*, 3048–3058. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.