THE IMPACTS OF FILE SYSTEM FRAGMENTATION ON AUTOMOTIVE STORAGE PERFORMANCE

TUXERA Make it work.

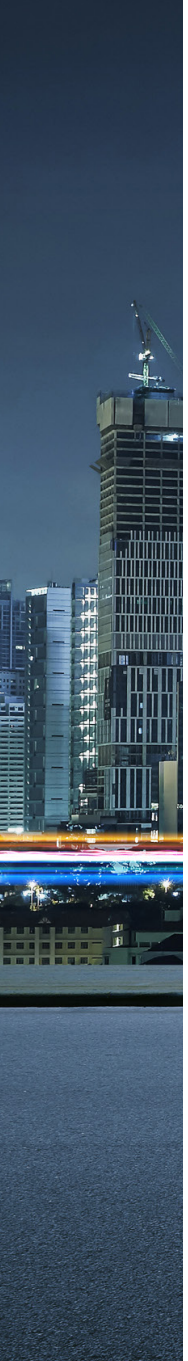

To achieve the highest levels of autonomy, connected cars need increasingly more sensors and cameras, all of which produce massive amounts of data. This data must be processed, analyzed, transferred, and stored using a hybrid solution of edge and cloud storage technologies.

Accordingly, automotive onboard (or edge) data storage capacity needs are also on the rise. Estimates by Western Digital predict that these requirements will soon climb to one terabyte or higher.[1] Advancements in flash memory technology, especially UFS, allow us to record the data pouring into the connected car at lightning-fast speeds.

Out of the box, automotive-grade flash memory hardware offers high inputs/outputs per second (IOPS). But software also acts on top of that hardware, performing various operations that transfer data into and out of the storage. Because of these workloads, achieving the performance listed on paper is a state of nirvana that may never be reached. To get the maximum potential from the storage hardware, you need high-performance storage management software.

COULD YOUR CAR'S FILE SYSTEMS BE DRAGGING YOU DOWN?

At Tuxera, we examine and test how file systems (the commonly known term for storage-management software), affect storage performance. File systems are responsible for allocating space for files on a physical storage medium.[2]

It's important for car makers and suppliers to choose their file system implementations wisely, as file systems impact read and write performance of the storage, the integrity of the data stored, flash endurance or the lifetime of the memory hardware, and data and storage interoperability.

Specific factors that affect file system performance include file size, device partitioning, or the file system implementations themselves. One additional factor Tuxera is currently testing is how fragmentation affects flash performance and lifetime.

WHAT IS FRAGMENTATION?

When a file system is first created on a drive partition, it takes up a small amount of space to build some essential structures.

USE HIGH-PERFORMANCE STORAGE MANAGEMENT SOFTWARE TO GET THE MAXIMUM POTENTIAL FROM YOUR HARDWARE

WHERE FILE SYSTEMS ARE MOST CRITICAL

Automotive systems that require the largest storage capacities are, unsurprisingly, those that generate and/or use heavy amounts of data. Consequently, you'll find file systems in these areas as well. Edge storage requirements can be broadly categorized into three functional domains: infotainment (IVI) and cluster, event recorder (EDR) and drive assistance (ADAS), and telematics or gateway.

- **Telematics/gateway**
- **Event recorder/drive assistance**
- **Infotainment/cluster**

FRAGMENTATION HINDERS STORAGE PERFORMANCE

At this stage, the file system primarily consists of one large adjoining block of free space. As files are created and apps are unpacked and added to the storage, separate files are, theoretically, neatly laid out near each other in sequences. This is a blissful, harmonious beginning—an optimal state of file allocation.

Over the course of time, however, the free space is used up. Files are deleted or need to be rewritten with new data. But the file system may not be able to place the new data near the rest of the file it belongs to, because that space may be occupied by another file's data. The data must then be saved apart from the rest of its file. Thus, a new file "fragment" is allocated elsewhere to the partition. This situation worsens over time as repeated read-manage-erase cycles mean smaller areas of free space. When the file system is nearly full, it typically has a lot of file fragments. In this state, we say the file system is fragmented, or aged (see illustration).

FRAGMENTATION SUPPOSEDLY "NOT A PROBLEM" FOR FLASH TECHNOLOGY

Back in the days when hard disk drives (HDDs) were the prevalent storage technology, the effects of file fragmentation were widely acknowledged. These disks had a spinning platter where the files were stored. A mechanical arm with an electromagnetic head skimmed across the surface of the

A CLOSER LOOK AT A FRAGMENTED FILE SYSTEM

Cells of the same color represent blocks of data that belong to the same file. As the storage fills, files become increasingly fragmented and the fragments are stored in locations on the storage which are not adjacent to each other.

disk to read or write all the data. It's very easy to understand how fragmentation would cause slowdowns in this case, as it takes time for the head to physically move around and find all the scattered fragments of a file.[3]

With flash memory technology, however, there was widespread belief from experts and enthusiasts alike that fragmentation was not an issue.[4] After all, file reading and writing is all electrical in nature and there's no physical arm that must move around finding fragments.

CONTRARY TO POPULAR BELIEF, FRAGMENTATION MATTERS

Fragmentation is a contentious subject when it comes to flash memory performance. For years, developers and tech enthusiasts argued that fragmentation does not significantly affect performance, or if it does, the problem's been solved with modern production file systems.[4] But these are simple assumptions that have been widely accepted for years without valid research to back them up—until recently.

Some groundbreaking studies have shown that these beliefs are wrong. This research suggests that as flash storage hardware gets faster, the software input/output (I/O) stack overhead is an I/O performance bottleneck. Fragmentation is now being reconsidered as one potential cause for sluggish flash storage performance.

IMPACTS OF FRAGMENTED FILE SYSTEMS

A study from 2017 by Conway et al. found that aged file systems cause a 2–5X performance slowdown on mobile flash hardware. In that study, the researchers were surprised by how rapidly aging degrades performance. They conclude, "a user's experience with unaged file systems is likely so fleeting that they do not notice performance degradation. Instead, the performance costs of aging are built into their expectations of file system performance."[4]

Another 2017 study by Han et al. validates this finding. Their research showed that mobile phone performance degraded quickly and drastically as the file system filled up. In their experiments, the Twitter app took 1.6 times longer to launch

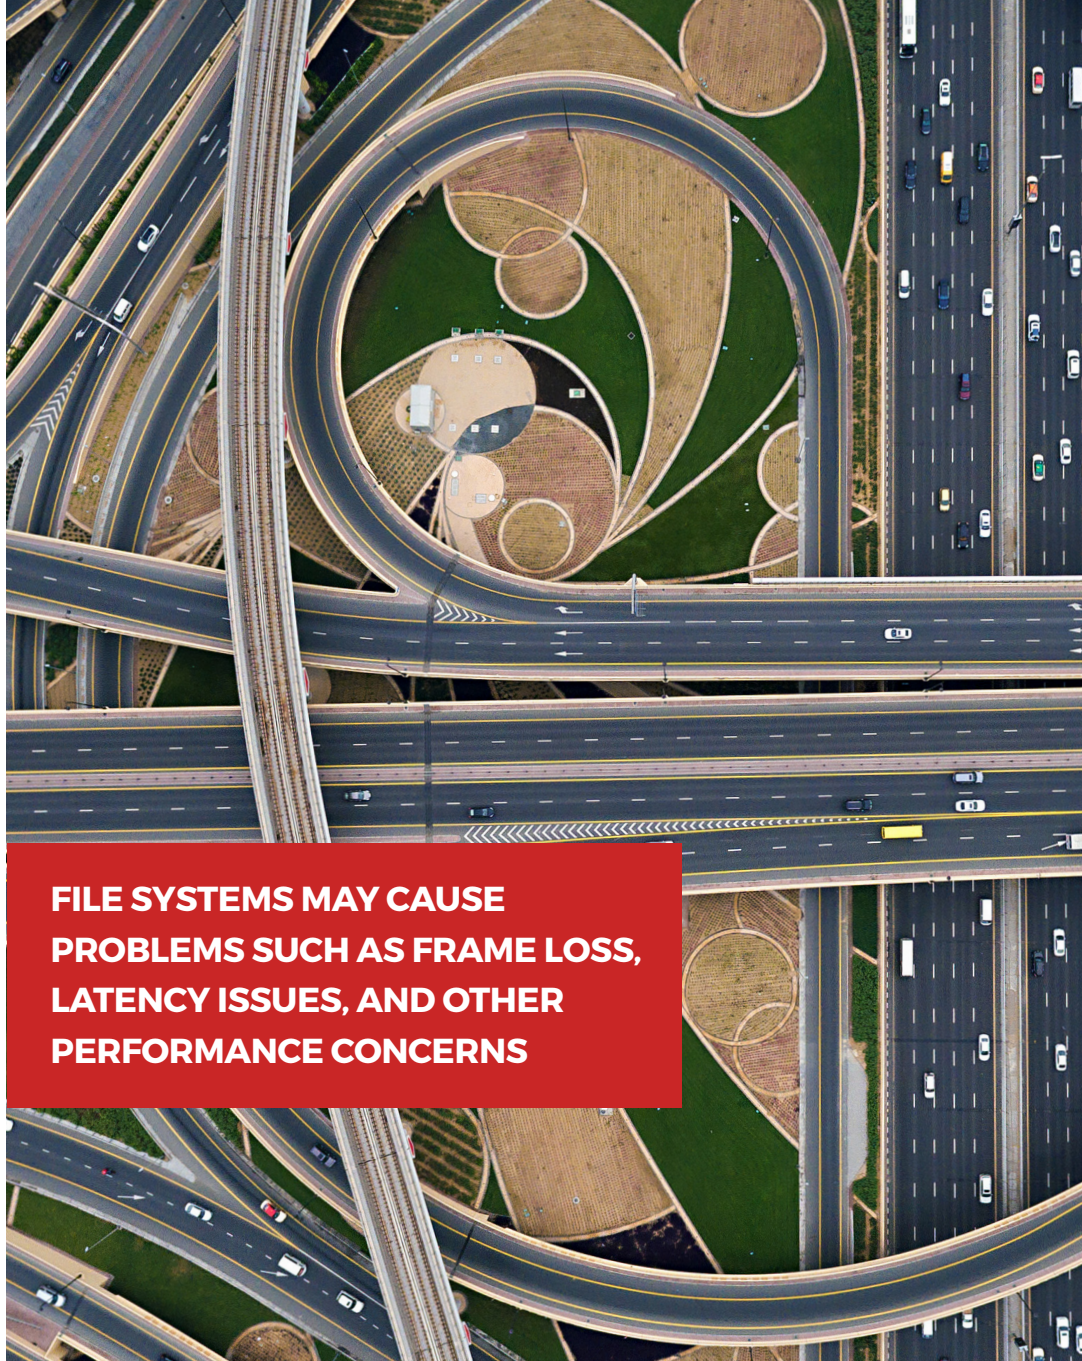

FILE SYSTEMS MAY CAUSE PROBLEMS SUCH AS FRAME LOSS, LATENCY ISSUES, AND OTHER PERFORMANCE CONCERNS

when the file system was 70% full, despite defragmentation on the storage just seven days before. When the file system was 90% full, this launch time delay extended to over twice as long as that of a new file system.[5]

Ji et al. (2016) discovered several interesting things concerning fragmentation. For one, I/O latency is proportional to the degree of fragmentation, and accessing fragmented files results in high I/O frequency. Secondly, as I/O block frequency increased, read performance dropped. Another important finding was that I/O locality matters. Both read and write latency increases when I/Os are highly dispersed.[6]

This means the I/O stack overhead is higher compared to when files are not (or less) fragmented. Thus, fragmentation drags down the overall performance of the storage.

Lastly, Ji et. al demonstrated that Android devices specifically suffer from severe fragmentation in SQLite database files, which in turn slows down application load time.[6] This is no small detail to overlook when an increasing number of car makers are adopting Android as the infotainment system OS of choice.

IMPACTS ON STORAGE PERFORMANCE IN AUTOMOTIVE SYSTEMS

Because smart cars use similar flash storage technologies, we suspect these mobile storage issues also plague automotive storage.

Ext4 is not only one of the most commonly used file systems in mobile phones, but in the automotive industry as well. The file system is a good option for low-data or single-stream automotive use cases. Plus, because it's open-source software, it's free. But when it comes to applications that handle a lot of data and multiple data streams, things get trickier.

Our customers and partners come to us because they suspect the file system could be a root cause of problems such as frame loss, latency issues, and other performance concerns. From our benchmarking, we've observed that over the short term, ext4 performs quite well despite fragmenting files. But over the long term, the degree of fragmentation gets worse as more data is written to the storage, and performance consequently drops.

What's more, automotive applications rarely have a single-threaded data stream. Most consist of at least two, if not four or more threads. For example, an automotive dashcam typically has one or two video streams (plus audio where applicable), as well as GPS data and other system-related files (logs, for example)—all of which are written to the storage at the same time. Such an application would also "clean up" old files when a pre-defined free-space threshold is reached. Under these workloads, our tests show that ext4 becomes heavily fragmented and its performance degrades over time.

TYPICAL MULTI-STREAM AUTOMOTIVE WORKLOAD PERFORMANCE

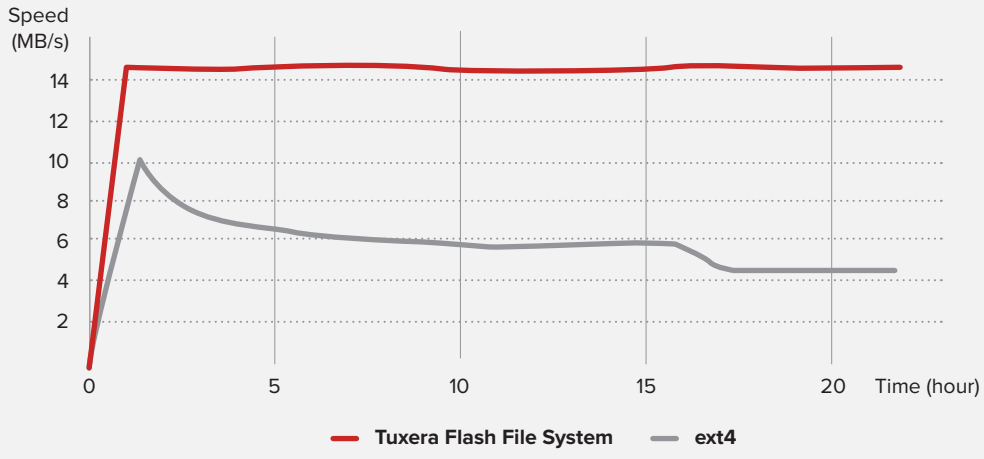

In the short term, ext4 performs quite well. But as the file system stores data over the course of several hours, performance drops. Test setup: Ambarella S2 A2 RH with microSDHC. Actual performance may vary based on the hardware, software, and testing protocols used.

INITIAL PERFORMANCE
EXT4 VS TUXERA FLASH FILE SYSTEM

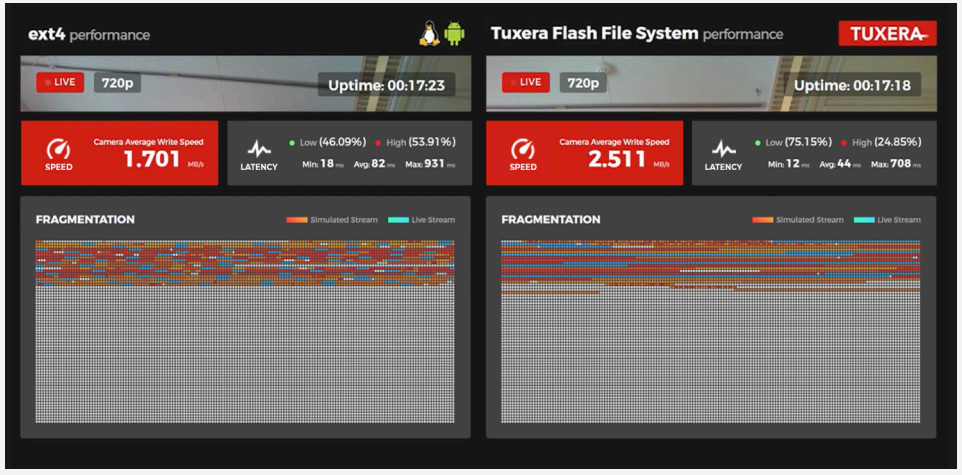

Initially, fragmentation is low and speed is quite high for both file systems.

PERFORMANCE UNDER LONG-TERM WORKLOAD
EXT4 VS TUXERA FLASH FILE SYSTEM

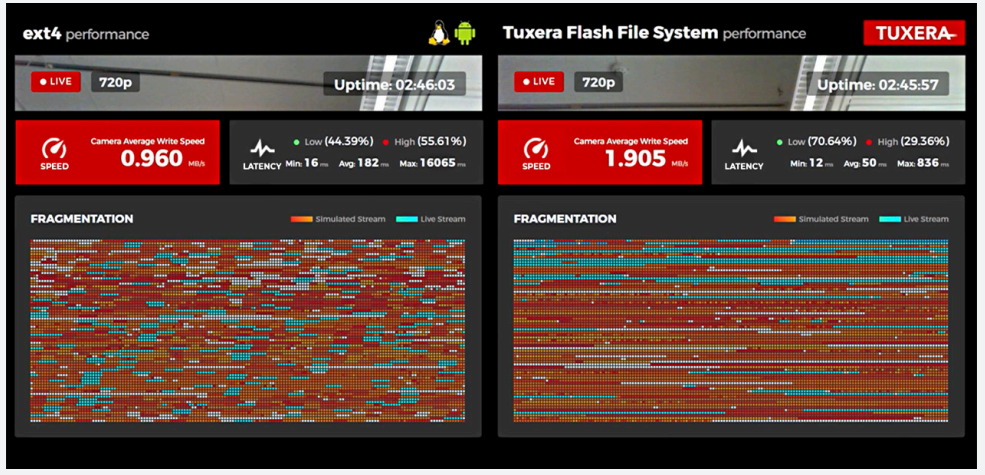

As the storage reaches capacity, you can see ext4 is very fragmented, performance has dropped, and latency has very sharp spikes, plus a high average value.

FRAGMENTATION TESTS RESULTS

You can see the effects of ext4 fragmentation in our time lapse screenshots. Here we show an approximately three-hour test run of ext4 against our proprietary flash-friendly file system implementation, Tuxera Flash File System. Each file system wrote data from a live camera stream, plus four additional streams to simulate sensor data (such as one might find in an ADAS application) to a removable memory card.

The box with cells is a visualization of the data as it's written to the flash memory blocks. Each color represents a unique data stream.

As shown, the data streams handled by ext4 become rapidly fragmented into short segments spread across the storage. At the same time, the red box shows the average write speed and latency. As ext4 becomes more fragmented and the storage fills, speed decreases and latency increases—the latter with some very extreme maximums. Tuxera Flash File System, on the other hand, writes the files in longer, contiguous sequences, and maintains higher speeds with consistently low latency.

See the full time-lapse video of fragmentation and performance.

FRAGMENTATION AFFECTS DRIVER AND PASSENGER EXPERIENCE

The degree of fragmentation and impact on performance varies depending on the use case. In applications with intensive reading, writing, and rewriting of data, (such as cameras for autonomous driving) fragmentation may cause anything from small errors to critical system failure. If the storage is full and heavily fragmented, there will definitely be read/write issues.[3] This can't happen when safety and lives are concerned.

In the case of infotainment, performance loss due to fragmentation boils down to a user experience and customer satisfaction concern. If the IVI system storage is heavily fragmented, this means longer wait times for music and navigation apps to launch. User experience is no small issue when it comes to differentiation, so even minor inconveniences such as app response time cannot be easily dismissed.

FIXING FRAGMENTATION REDUCES FLASH LIFETIME

Early failure of the storage is another potential consequence of heavy

REDUCING FRAGMENTATION SHOULD BE A CORE CHARACTERISTIC OF FILE SYSTEM DESIGN

fragmentation. With the ext4 file system, for example, the only way to guarantee a satisfactory user experience is to periodically defragment the file system.[5] While many defragmentation tools exist, the conventional ext4 utility is e4defrag. Hahn et al. (2017) demonstrated in their research that using e4defrag requires copying significant amounts of data. Thus, weekly defragmentation of the storage could reduce the lifetime by up to 10%.[5] This level of write and erase wear directly affects the recall possibility of a car due to storage failure. As one can see, doing damage control with defragmenting tools is not the optimal solution.

SO WHAT'S THE ANSWER?

The auto industry needs various methods to prevent or alleviate fragmentation to guarantee both user satisfaction and safety, as well as extend flash lifetime.

The trick is to reduce fragmentation in the first place, which relies heavily on the file system design itself. A file system needs to be "smart"—carefully engineered to lay out files as logically as possible, with the lowest amount of fragmentation. If done inadequately, write operations can drastically increase the fragmentation of storage, making every next read or write operation slower.

While preventing fragmentation altogether may be impossible, reducing fragmentation should be a core characteristic of a file system, especially for automotive applications.

Fragmentation is a first-order performance issue[5]—though to what degree in automotive storage? It's a topic we're still exploring. What we do know is that the file system used to interface with the hardware will fragment data, which works against all the advantages of flash memory technology.

MINIMIZE FRAGMENTATION WITH TUXERA'S HIGH-PERFORMANCE SOFTWARE

Car makers and Tier-1 suppliers—let us help you get the maximum performance and lifetime from your entire storage stack.

All our file system implementations are designed to significantly reduce the detrimental effects of fragmentation and improve performance.

- **See Tuxera automotive solutions**
- **Free consultation**

REFERENCES

1 ***What's Driving Automotive Storage?***
Blog post by Martin Booth of published 30 May 2017 in EETimes.

2 ***File system fragmentation*** in Wikipedia.

3 ***The real and complete story—Does Windows defragment your SSD?***
Blog post by Scott Hanselman published 3 December 2014.

4 Conway et al. ***File Systems Fated for Senescence? Nonsense, Says Science!***
15th USENIX Conference on File and Storage Technologies. (2017).

5 Hahn et al. ***Improving File System Performance of Mobile Storage Systems Using a Decoupled Defragmenter.***
2017 USENIX Annual Technical Conference. (2017).

6 Ji et al. ***An Empirical Study of File-System Fragmentation in Mobile Storage Systems.***
8th {USENIX} Workshop on Hot Topics in Storage and File Systems (HotStorage 16). (2016).

OTHER RECOMMENDED READING

Bouganim et al. ***uFLIP: Understanding Flash IO Patterns.***
Conference on Innovative Data Systems Research (CIDR). (2009).

Smith, Keith A. and Seltzer, Margo I. ***File System Aging—Increasing the Relevance of File System Benchmarks.***
Proceedings of the SIGMETRICS 97 ACM SIGMETRICS Conference on Measurement & Modeling of Computer Systems. (1997).

GET IN TOUCH.

HEADQUARTERS
Westendintie 1
02160 Espoo, Finland
General email: **info@tuxera.com**
Product inquiries: **sales@tuxera.com**

CHINA
401, Block A, Vantone Center,
Jia No.6 Chaoyangmenwai Street,
Chaoyang District, Beijing, China
Tel: **+86 138 0109 5473**
Email: **china@tuxera.com**

KOREA
1812 SK V1 Center,
Gasandigital-1-ro 171, Geumcheon-gu,
Seoul 08503, Korea
Tel: +82 2 867 7370
Email: korea@tuxera.com

JAPAN
Level 18 Yebisu Garden Place Tower
4-20-3 Ebisu, Shibuya-ku
Tokyo 150-6018, Japan
Tel: **+81 35 789 5128**
Email: **japan@tuxera.com**

UNITED STATES
440 N.Wolfe Rd.
Sunnyvale, CA 94085, USA
Tel: **+1 408 744 6305**
Email: **usa@tuxera.com**

TAIWAN
10F-1, No.209, Sec 1,
Civic Blvd Datong District
Taipei 103, Taiwan
Tel: **+886 9 2183 2407**
Email: **taiwan@tuxera.com**