

DOI:10.1145/3318163

Tracing the evolution of the five-minute rule to help identify imminent changes in the design of data management engines.

**BY RAJA APPUSWAMY, GOETZ GRAEFE,
RENATA BOROVIKA-GAJIC, AND ANASTASIA AILAMAKI**

The Five-Minute Rule 30 Years Later and Its Impact on the Storage Hierarchy

THE DESIGN OF data management systems has always been influenced by the storage hardware landscape. In the 1980s, database engines used a two-tier storage hierarchy consisting of dynamic random access memory (DRAM) and hard disk drives (HDD). Given the disparity in cost between HDD and DRAM, it was important to determine when it made economic sense to cache data in DRAM as opposed to leaving it on the HDD.

In 1987, Jim Gray and Gianfranco Putzolu established the five-minute rule that gave a precise answer to this question: “1KB records referenced every

five minutes should be memory resident.”⁹ They arrived at this value by using the then-current price-performance characteristics of DRAM and HDD shown in Table 1 for computing the break-even interval at which the cost of holding 1KB of data in DRAM matches the cost of I/O to fetch it from HDD.

Today, enterprise database engines use a three-tier storage hierarchy as depicted in Figure 1. DRAM or NAND flash solid state device (SSD)-based performance tier is used for hosting data accessed by latency-critical transaction processing and real-time analytics applications. The HDD-based capacity tier hosts data accessed by latency-insensitive batch analytics applications. The archival tier is not used for online query processing, but for storing data that is only accessed rarely during regulatory compliance audits or disaster recovery. This tier is primarily based on tape and is extremely crucial as a long-term data repository for several application domains like physics, banking, security, and law enforcement.

In this article, we revisit the five-minute rule three decades after its inception. We recomputed break-even intervals for each tier of the modern, multi-tiered storage hierarchy and use guidelines provided by the five-minute rule to identify impending changes in the design of data management engines for emerging storage hardware. We summarize our findings here:

► **HDD is tape.** The gap between DRAM and HDD is increasing as the five-minute rule valid for the DRAM-HDD case in 1987 is now a four-hour rule. This implies the HDD-based capacity tier is losing relevance for not just performance sensitive applications, but for all applications with a non-sequential data access pattern.

► **Non-volatile memory is DRAM.** The gap between DRAM and SSD is shrinking. The original five-minute rule is now valid for the DRAM-SSD case, and the break-even interval is less than a minute for newer non-volatile memory (NVM) devices like 3D-



XPoint.²³ This suggests an impending shift from DRAM-based database engines to flash or NVM-based persistent memory engines.

► **Cold storage is hot.** The gap between HDD and tape is also rapidly shrinking for sequential workloads. New cold storage devices that are touted to offer second-long access latency with cost comparable to tape reduce this gap further. This suggests the HDD-based capacity tier will soon lose relevance even for non-performance-critical batch analytics applications that can be scheduled to run directly over newer cold storage devices.

Revisiting the Five-Minute Rule

The five-minute rule explores the trade-off between the cost of DRAM and the cost of disk I/O by providing a formula to predict the break-even interval—the

time window within which data must be reaccessed in order for it to be economically beneficial to be cached in DRAM. The interval is computed as:

$$\frac{\text{PagesPerMBofDRAM}}{(\text{AccessesPerSecondPerDisk})} \times \frac{\text{PricePerDiskDrive}}{\text{PricePerMBofDRAM} (1)} \quad (1)$$

The first ratio in the equation was referred to as the technology ratio, as random I/O access capability of the secondary storage device, and the page size used by the database engine for performing I/O, both directly depend on the hardware technology used for secondary storage. The second ratio, in contrast, is referred to as the economic ratio as pricing is determined by factors other than just hardware technology. Rearranging the formulation by swapping the denominators provides the intuition behind the five-minute rule, as it reduces the equation to price-per-

disk-access-per-second normalized by the price-per-page of DRAM. This term directly compares the cost of performing I/O to fetch a page from disk versus the cost of caching it in DRAM.

Table 1 shows the price, capacity, and performance of DRAM, HDD, and NAND flash-based SSDs across four decades. The values shown for 1987, 1997, and 2007 are those reported by previous revisions of the five-minute rule.^{6,8,9} The values listed for 2018 are performance metrics listed in vendor specifications, and unit price quoted by www.newegg.com as of Mar. 1, 2018, for DRAM, SSD, and HDD components specified in a recent TPC-C report.²⁴

DRAM-HDD. Table 2 presents both the break-even interval for 4KB pages and the page sizes for which the five-minute rule is applicable across four decades. In 1987, the break-even inter-

Table 1. The evolution of DRAM, HDD, and Flash SSD properties.

Metric	DRAM				HDD				SATAFlash SSD	
	1987	1997	2007	2018	1987	1997	2007	2018	2007	2018
Unit price(\$)	5k	15k	48	80	30k	2k	80	49	1k	415
Unit capacity	1MB	1GB	1GB	16GB	180MB	9GB	250GB	2TB	32GB	800GB
\$/MB	5k	14.6	0.05	0.005	83.33	0.22	0.0003	0.00002	0.03	0.0005
Random IOPS	-	-	-	-	5	64	83	200	6.2k	67k (r)/20k (w)
Sequential b/w (MB/s)	-	-	-	-	1	10	300	200	66	500 (r)/460 (w)

val was 400 seconds for 1KB pages. This was rounded down to five minutes, thus, lending the name for the rule. For 4KB pages, the break-even interval was 100 seconds. When the study was repeated in 1997, the break-even interval had increased to nine minutes for 4KB pages, and the five-minute rule was determined to hold only for 8KB pages. Between 1997 and 2007, DRAM and HDD prices dropped further resulting in the economic ratio increasing from 133 (\$2k/\$15) to 1700 (\$80/\$0.05). However, the technology ratio did not drop proportionately due to a lack of improvement in HDD random access latency. As a result, the break-even interval for 4KB pages increased 10x, from nine minutes to 1.5 hours. The five-minute rule was applicable only for 64KB pages in 2007.

Continuing this trend, the break-even interval for DRAM-HDD case today is four hours for 4KB pages. The five-minute rule is valid today for 512KB pages. The break-even interval trend indicates it is more economical to store most data in DRAM instead of the HDD, and the page size trend indicates that

even rare accesses to HDD should be performed in large granularities.

DRAM-SSD. SSDs are being increasingly used as the storage medium of choice in the latency-critical performance tier due to their superior random access capability compared to HDDs. Thus, the five-minute rule can be used to compute a break-even interval for the case where DRAM is used to cache data stored in SSDs. Table 3 shows the interval in 2007, when SSDs were in the initial stages of adoption, and today, based on metrics listed in Table 1.

We see the interval has dropped from 15 minutes to five minutes for 4KB pages. Thus, the five-minute rule is valid for SSDs today. This is in stark contrast with the DRAM-HDD case, where the interval increased 2.7x from 1.5 hours to four hours. In both DRAM-HDD and DRAM-SSD cases, the drop in DRAM cost/MB dominated the economic ratio. However, unlike the 2.5x improvement in random I/Os-per-second (IOPS) with HDDs, SSDs have managed to achieve an impressive 11x improvement (67k/6.2k). Thus, the increase in economic ratio was overshadowed

by the decrease in technology ratio with SSDs, resulting in the interval shrinking.


SSD-HDD. As SSDs can also be used as a cache for HDD, the same formula can also be used to estimate the break-even interval for the SSD-HDD case. From Table 3, we see the break-even interval for this case has increased by a factor of 10x from 2.25 hours in 2007 to 1.5 days in 2018. The SSD-HDD interval is nine times longer than the DRAM-HDD interval of four hours.

Implications. There are two important consequences of these results. First, in 2007, the turnover time in the DRAM-HDD case was six times higher than the DRAM-SSD case (1.5h/15m). In 2018, it is nearly 50x higher (4h/5m). Thus, in systems tuned using economic considerations, one should replace HDD with SSD, as it would not only improve performance, but also reduce the amount of DRAM required for caching data. Second, given the four-hour DRAM-HDD and one day SSD-HDD intervals, it is important to keep all active data in the DRAM or SSD-based performance tier and relegate the HDD-based capacity tier to storing only infrequently accessed data. The growing gap between performance and capacity tiers also implies that SSD vendors should optimize for \$/IOPS, and HDD vendors, in contrast, should optimize for \$/GB. Next, we highlight recent changes in performance and capacity tiers that indicate such targeted optimizations are already underway.

The Performance Tier

NAND flash. NAND flash-based solid-state storage has been steadily inching its way closer to the CPU over the past two decades. When NAND flash was introduced in the early 2000s, solid-state storage was dominated by

Figure 1. Storage tiering for enterprise databases.



DRAM-based SSD products. By the mid 2000s, improvements in performance and reliability of NAND flash resulted in flash-based serial AT attachment (SATA) SSDs gaining popularity in niche application domains. The late 2000s witnessed the emergence of a new breed of peripheral component interconnect express (PCIe) flash SSDs that could deliver two orders of magnitude higher throughput than their SATA counterparts. Since then, a rapid increase in capacity, drop in pricing, and new low-overhead interfaces like non-volatile memory express (NVMe), have all resulted in PCIe flash SSDs displacing their SATA counterparts as server accelerators of choice.

Table 4 (first row) shows the price/performance characteristics of a representative, state-of-the-art PCIe SSD. In comparison to Table 1, we find the PCIe SSD offers five times higher read IOPS and sequential access bandwidth than its SATA counterpart.

NVDIMM. As SSD vendors continue to improve throughput and capacity, the bottleneck in the storage subsystem has shifted from the device itself to the PCIe bus that is used to interface with the SSD. Thus, over the past few years, NAND flash has started transitioning once again from storage devices that are interfaced via the high-latency, bandwidth-limited PCIe bus into non-volatile memory (NVM) devices that are interfaced via the low-latency, high-bandwidth memory bus. These devices, also referred to as non-volatile DIMMs (NVDIMM), use a combination of DRAM and flash storage media packaged together as a dual inline memory module (DIMM).

NVM. Today, NVDIMMs are niche accelerators compared to PCIe SSDs due to a high cost/GB. Unlike these NVDIMM technologies that rely on NAND flash, new NVM technologies that are touted to have better endurance, higher throughput, and lower latency than NAND flash are being actively developed.

Table 4 (second row) shows the characteristics of Intel Optane DC P4800X—a PCIe SSD based on 3D XPoint, a new phase-changed-media-based NVM technology. The cost/GB of 3D XPoint is higher than NAND flash today as the technology is yet to mature. However, preliminary studies

have found that 3D XPoint provides predictable access latencies that are much lower than several state-of-the-art NAND flash devices even under severe load.²³

Break-even interval and implications. When we apply the five-minute rule formula using metrics given in Table 4, we get a break-even interval of one minute for 4KB pages in both the DRAM-NAND Flash PCIe SSD and DRAM-3D XPoint cases. Comparing these results with Table 2, we see that the breakeven interval is 10× shorter when PCIe SSDs or new PM technologies are used as the second tier instead of SATA SSDs. This can be attributed to the drop in technology ratio caused by the improvement in random IOPS.

Implications. Today, in the era of in-memory data management, several database engines are designed based on the assumption that all data is resident in DRAM. However, the dramatic drop in breakeven interval computed by the five-minute rule challenges this trend of DRAM-based in-memory data management due to three reasons. First, recent projections indicate that flash density is expected to increase 40% annually over the next five years.⁵ DRAM, in contrast, is doubling in capacity every three years.¹⁷ As a result,

the cost of NAND flash is likely to drop faster than DRAM. This, in turn, will result in the economic ratio dropping further leading to a reduction in the break-even interval.

Second, modern PCIe SSD is a highly parallel device that can provide very high random I/O throughput by servicing multiple outstanding I/Os concurrently. New non-volatile memory technologies like 3D XPoint promise further improvements in both throughput and access latencies over NAND flash. With interfaces like NVMe, the end-to-end latency of accessing data from PCIe 3D XPoint SSDs is just tens of μs. Thus, further improvements in non-volatile solid-state storage media will result in a drop in technology ratio, thereby reducing the break-even interval further.

Third, SSDs consume substantially lower power than DRAM. The Intel 750 SSD consumes 4W of power when idle and 22W when active. In contrast, 1TB of DRAM in a server would consume 50W when idle and 100W when active.¹ It is also well known that DRAM power consumption increases non-linearly with capacity, as high-density DRAM consumes substantially more power than their low-density counterparts. A recent study that focuses on power

Table 2. The evolution of the page size for which the five-minute rule holds across four decades based on appropriate price, performance, and page size values.

	1987	1997	2007	2018
Break-even (4KB page)	100s	9m	1.5h	4h
Page size (5-minute interval)	1KB	8KB	64KB	512KB

Table 3. The evolution of the break-even interval across four decades based on appropriate price, performance, and page size values.

Tier	1987	1997	2007	2018
DRAM-SSD	—	—	15m	5m
SSD-HDD	—	—	2.25h	1.5d

Table 4. Price/performance metrics for the NAND-based Intel 750 PCIe SSD and 3D-XPoint-based Intel Optane P4800X PCIe SSD.

Device	Capacity	Price(\$)	IOPS(k)	B/w(GB/s)
Intel 750	800GB	589	460	2.5
Intel P4800X	480GB	617	550	2.5

consumption in main memory databases showed that in a server equipped with 6TB of memory, the idle power of DRAM would match that of four active CPUs.¹ Such a difference in power consumption between SSD and DRAM directly translates into higher Operational Expenses (OPEX), and hence, higher Total Cost of Ownership (TCO), for DRAM-based database engines.

Given these three factors, the break-even interval from the five-minute rule seems to suggest an inevitable shift from DRAM-based data management engines to NVM-based persistent-memory engines. In fact, this change is already well under way, as state-of-the-art database engines are being updated to fully exploit the performance benefits of PCIe NVM SSDs.²⁶ Researchers have recently highlighted the fact that data caching systems that trade-off performance for price by reducing the amount of DRAM are gaining market share over in-memory database engines.¹⁸

The Capacity Tier

HDD. Traditionally, HDDs have been the primary storage media used for provisioning the capacity tier. For several years, areal density improvements

enabled HDDs to increase capacity at Kryder’s rate (40% per year), outstripping Moore’s Law. However, over the past few years, HDD vendors have hit walls in scaling areal density with conventional Perpendicular Magnetic Recording (PMR) techniques resulting in annual areal density improvement of only around 16% instead of 40%.¹⁹

HDDs also present another problem when used as the storage medium of choice for building a capacity tier, namely, high idle power consumption. Although enterprises gather vast amounts of data, as one might expect, not all data is accessed frequently. Recent studies estimate that as much as 80% of enterprise data is “cold,” meaning infrequently accessed, and that cold data is the largest growing segment with a 60% Cumulative Annual Growth Rate (CAGR).¹⁰⁻¹² Unlike tape, which consumes no power once unmounted, HDDs consume a substantial amount of power even while idle. Such power consumption translates to a proportional increase in TCO.

Tape. The areal density of tape has been increasing steadily at a rate of 33% per year and roadmaps from the Linear Tape Open consortium (LTO)²⁵

and the Information Storage Industry Consortium (INSIC)¹³ project a continued increase in density for the foreseeable future.

Table 5 shows the price/performance metrics of tape storage both in 1997 and today. The 1997 values are based on the corresponding five-minute rule paper.⁸ The 2018 values are based on a SpectraLogic T50e tape library²² using LTO-7 tape cartridges.

With individual tape capacity increasing 200× since 1997, the total capacity stored in tape libraries has expanded from hundreds of gigabytes to hundreds of petabytes today. Further, a single LTO-7 cartridge is capable of matching, or even outperforming a HDD, with respect to sequential data access bandwidth as shown in Table 6. As modern tape libraries use multiple drives, the cumulative bandwidth achievable using even low-end tape libraries is 1–2GB/s. High-end libraries can deliver well over 40GB/s. These benefits have made tape the preferable media of choice in the archival tier both on-premise and in the cloud, for several applications ranging from natural sciences, like particle physics and astronomy, to movies archives in the entertainment industry.^{15,20} However, random access latency of tape is still 1000× higher than HDD (minutes vs. ms) due to the fact that tape libraries need to mechanically load and wind tape cartridges before data can be accessed.

Break-even interval and implications. Using metrics from Tables 1, 5 to compute the break-even interval for the DRAM–tape case results in an interval of over 300 years for a page size of 4KB! Jim Gray referred to tape drives as the “data motel” where data checks in and never checks out,⁷ and this is certainly true today. Figure 2 shows the variation in break-even interval for both HDD and tape for various page sizes. We see that the interval asymptotically approaches one minute in the DRAM–HDD case and 10 minutes in the DRAM–tape case. The HDD asymptote is reached at a page size of 100MB and the tape asymptote is reached at a size of 100GB. This clearly shows that randomly accessing data on these devices is extremely expensive, and data transfer sizes with these devices should be large to

Table 5. Price/performance characteristics of tape.

	1997	2018
Tape library cost (\$)	10,000	11,000
Number of drives	1	4
Number of slots	14	10
Max capacity per tape	35GB	15TB
Transfer rate per drive (MB/s)	5	750
Access latency	30s	65s

Table 6. Price/performance metrics of DRAM, HDD, and tape.

Metric	DRAM	HDD	Tape
Unit capacity	16GB	2TB	10 × 15TB
Unit cost (\$)	80	50	11,000
Latency	100ns	5ms	65s
Bandwidth	100 GB/s	200 MB/s	4 × 750MB/s
Kaps	9,000,000	200	0.02
Maps	10,000	100	0.02
Scan time	0.16s	3hours	14hours
\$/Kaps	9e-14	5e-09	8e-03
\$/Maps	9e-12	8e-09	8e-03
\$/Tbscan	8e-06	0.003	0.03
\$/TBscan (97)	0.32	4.23	296

amortize the cost of random accesses.

However, the primary use of the capacity tier today is not supporting applications that require high-performance random accesses. Rather, it is to reduce the cost/GB of storing data over which latency-insensitive batch analytics can be performed. Indeed, Gray and Graefe noted that metrics like KB-accesses-per-second (Kaps) are less relevant for HDD and tape as they grow into infinite-capacity resources.⁸ Instead, MB-accesses-per-second (Maps) and time to scan the whole devices are more pertinent to these high-density storage devices. Table 6 shows these new metrics and their values for DRAM, HDD, and tape. In addition to Kaps, Maps, and scan time, the table also shows \$/Kaps, \$/Maps, and \$/TB-scan, where costs are amortized over a three-year time frame as proposed by Gray and Graefe.⁸

Looking at \$/Kaps, we see that DRAM is five orders of magnitude cheaper than HDD, which, in turn, is six orders of magnitude cheaper than tape. This is expected given the huge disparity in random access latencies and is in accordance with the five-minute rule that favors using DRAM for randomly accessed data. Looking at \$/Maps, we see that the difference between DRAM and HDD shrinks to roughly 1,000×. This is due to the fact that HDDs can provide much higher throughput for sequential data accesses over random ones. However, HDD continue to be six orders of magnitude cheaper than tape even for MB-sized random data accesses. This, also, is in accordance with the HDD/tape asymptote shown in Figure 2. Finally, \$/TBscan paints a very different picture. While DRAM remains 300× cheaper than HDD, the difference between HDD and tape shrinks to 10×.

Comparing the \$/TBscan values with those reported in 1997, we can see two interesting trends. First, the disparity between DRAM and HDD is growing over time. In 1997, it was 13× cheaper to use DRAM for a TBscan than HDD. Today, it is 300× cheaper. This implies that even for scan-intensive applications, unsurprisingly, optimizing for performance requires avoiding using HDD as the storage medium. Second, the difference between HDD and tape is following the opposite trend and shrinking over time. In 1997, HDD was

70× cheaper than tape. However, today it is only 10× cheaper. Unlike HDD, sequential data transfer bandwidth of tape is predicted to double for the foreseeable future. Hence, this difference is likely to shrink further. Thus, in the near future, it might not make much of a difference whether data is stored in a tape or HDD with respect to the price paid per TB scan.

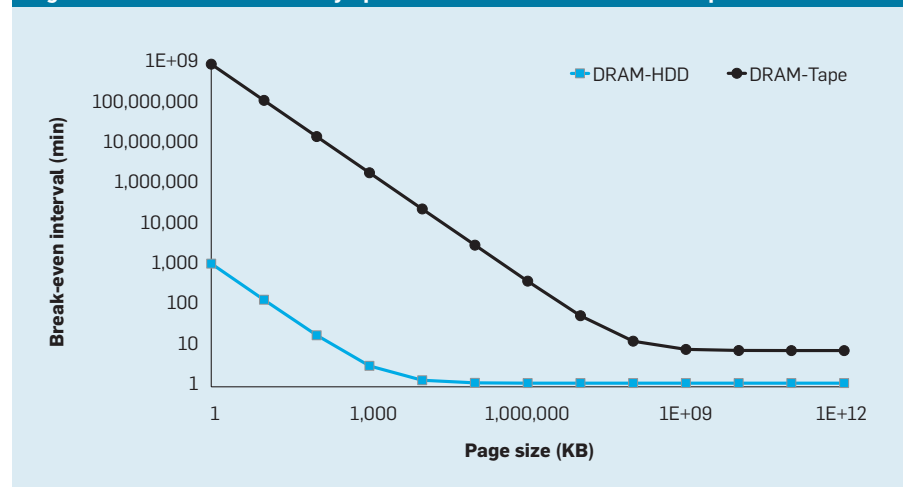
Implications. Today, all data generated by an enterprise has to be stored twice, once in the traditional HDD-based capacity tier for enabling batch analytics, and a second time in the tape-based archival tier for meeting regulatory compliance requirements. The shrinking difference in \$/TBscan between HDD and tape suggests that it might be economically beneficial to merge the capacity and archival tiers into a single *cold storage tier*.³ However, with such a merger, the cold storage tier would no longer be a near-line tier that is used rarely during disaster recovery, but an online tier that is used for running batch analytics applications. Recent hardware and application trends indicate that it might be feasible to build such a cold storage tier.

On the hardware front, storage vendors have recently started building new *cold storage devices* (CSD) for storing cold data. Each CSD is an ensemble of HDDs grouped in a massive array of idle disks (MAID) setup where only a small subset of disks are active at any given time.^{2,4,27} For instance, Pelican CSD provides 5PB of storage using 1,152 SMR disks packed as a 52U rack appliance.² However, only 8% of disks can be spun up simultaneously due

to cooling and power restrictions enforced by hardware. Access to data in any of the spun-up disks can be done with latency and bandwidth comparable to that of the traditional capacity tier. For instance, Pelican, OpenVault Knox, and ArticBlue provide between 1–2GB/s of throughput for reading data from spun-up disks.^{2,21,27} However, accessing data on a spun-down disk takes several seconds, as the disk has to be spun up before data can be retrieved. Thus, CSDs form a perfect middle ground between HDD and tape with respect to both cost/GB and access latency.

On the application front, there is a clear bifurcation in demand between latency-sensitive interactive applications and latency insensitive batch applications. As interactive applications are isolated to the performance tier, the cold storage tier only has to cater to the bandwidth demands of latency-insensitive batch analytics applications. Nearline storage devices like tape libraries and CSD are capable of providing high-throughput access for sequentially accessed data. Thus, researchers have recently started investigating extensions to batch processing frameworks for enabling analytics directly over data stored in tape archives and CSD. For instance, Nakshatra implements prefetching and I/O scheduling extensions to Hadoop so that mapreduce jobs can be scheduled to run directly on tape archives.¹⁴ Skipper is a query-processing framework that uses adaptive query processing techniques in combination with customized caching and I/O scheduling to enable que-

Figure 2. Break-even interval asymptotes for DRAM–HDD and DRAM–tape cases.



ry execution over CSD.³ Skipper even shows that for long-running batch queries, using CSD results in query execution time increasing by only 35% compared to a traditional HDD despite the long disk spin-up latency. With such frameworks, it should be possible for installations to switch from the traditional three-tier hierarchy to a two-tier hierarchy consisting of just a performance tier with DRAM and SSDs, and a cold storage tier with CSDs.

Conclusion and Future Work

Modern database engines use a three-tier storage hierarchy across four primary storage media (DRAM, SSD, HDD, and tape) with widely varying price-performance characteristics. In this article, we revisited the five-minute rule in the context of this modern storage hierarchy and used it to highlight impending changes based on recent trends in the hardware landscape.

In the performance tier, NAND flash is inching its way closer to the CPU resulting in dramatic improvements in both access latency and bandwidth. For state-of-the-art PCIe SSDs, the break-even interval predicted by the five-minute rule is one minute for 4KB pages. Going forward, further improvements in NAND flash and the introduction of new NVM technologies will likely result in this interval dropping further. As the data reuse window shrinks, it will soon be economically more valuable to store most, if not all, data on solid-state storage devices instead of DRAM. This will invariably necessitate revisiting several techniques pioneered by traditional HDD-based database engines, but eschewed by in-memory engines, like buffer caching, on-disk storage layout, and index persistence, to name a few, for these new low-latency, high-bandwidth storage devices.

Traditionally, HDDs have been used for implementing the capacity tier. However, our analysis showed that the difference between HDD and tape is shrinking when \$/TBScan is used as the metric. Given the latency-insensitive nature of batch analytics workloads, it is economically beneficial to merge the HDD-based capacity tier and the tape-based archival tier into a single cold storage tier as demonstrated by recent research.³ However, several open questions still need to

answered in order for the cold storage tier to be feasible in practice.

Over the past few years, several other systems have been built to reduce the cost of storing cold data using alternative storage media. For instance, DT-Store¹⁶ uses LTFS tape archive for reducing the TCO of online multimedia streaming services by storing cold data in tape drives. ROS²⁸ is a PB-sized, rack-scale cold storage library built using thousands of optical discs packed in a single 42U Rack. Today, it is unclear as to how these alternative storage options fare with respect to HDD-based CSD as the storage media of choice for storing cold data. Furthermore, in order for the Cold Storage Tier to be realized in practice, an ideal cold storage media needs to support batch analytics workloads. CSD, tape, and optical media are all primarily used today for archival storage where data is rarely read. Further research is required to understand the reliability implications of using these storage devices under batch analytics workloads.

Finally, with widespread adoption of cloud computing, the modern enterprise storage hierarchy not only spans several storage devices, but also different geographic locations from direct-attached low-latency devices, through network-attached storage servers, to cloud-hosted storage services. The price-performance characteristics of these storage configurations vary dramatically depending not only on the storage media used, but also on other factors like the total capacity of data stored, the frequency and granularity of I/O operations used to access the data, the read-write ratio, the duration of data storage, and the cloud service provider used, to name a few. Given the multitude of factors, determining the break-even interval for cloud storage is a complicated problem that we did not consider in this work. Thus, another interesting avenue of future work is extending the five-minute rule to such a distributed cloud storage setting. □

References

1. Appuswamy, R., Olma, M., and Ailamaki, A. Scaling the memory power wall with dram-aware data management. In *Proceedings of DaMoN*, 2015.
2. Balakrishnan, S. et al. Pelican: A building block for exascale cold data storage. In *Proceedings of OSDI*, 2014.
3. Borovica-Gajic, R., Appuswamy, R., and Ailamaki, A. Cheap data analytics using cold storage devices. In *Proceedings of VLDB 9*, 12 (2016).

4. Colarelli, D. and Grunwald, D. Massive arrays of idle disks for storage archives. In *Proceedings of 2002 Conference on Supercomputing*.
5. Coughlin, T. Flash memory areal densities exceed those of hard drives; <http://bit.ly/2NbDh5T>.
6. Graefe, G. The five-minute rule 20 years later (and how flash memory changes the rules). *Commun. ACM* 52, 7 (July 2009).
7. Gray, J. The five-minute rule; research.microsoft.com/en-us/um/people/gray/talks/fiveminuterule.ppt.
8. Gray, J. and Graefe, G. The five-minute rule ten years later, and other computer storage rules of thumb. *SIGMOD Rec.* 26, 4 (1997).
9. Gray, J. and Putzolu, F. The 5-minute rule for trading memory for disc accesses and the 10-byte rule for trading memory for CPU time. In *Proceedings of SIGMOD*, 1987.
10. Horizon Information Strategies Report. Tiered storage takes center stage, IDC. Technology assessment: Cold storage is hot again – finding the frost point; <http://www.storiant.com/resources/Cold-Storage-Is-Hot-Again.pdf>.
11. Intel. Cold Storage in the Cloud: Trends, Challenges, and Solutions, 2013; <https://intel.ly/2ZG74F6>.
12. I.S.I. Consortium. International magnetic tape storage roadmap; <http://www.insic.org/news/2015roadmap/15index.html>
13. Kathpal, A. and Yasa, G.A.N. Nakshatra: Towards running batch analytics on an archive. In *Proceedings of MASCOTS*, 2014.
14. Lantz, M. Why the future of data storage is (still) magnetic tape; <http://bit.ly/2XChrMO>
15. Lee, J., Ahn, J., Park, C., and Kim, J. Dtstorage: Dynamic tape-based storage for cost-effective and highly-available streaming service. In *Proceedings of CCGRID*, 2016.
16. Lim, K., Chang, J., Mudge, T., Ranganathan, P., Reinhardt, S.K., and Wenisch, T.F. Disaggregated memory for expansion and sharing in blade servers. In *Proceedings of ISCA*, 2009.
17. Lomet, D. Cost/performance in modern data stores: How data caching systems succeed. In *Proceedings of DaMoN*, 2018.
18. Moore, F. Storage outlook 2016; <http://bit.ly/2KBLgao>.
19. Perlmutter, M. The lost picture show: Hollywood archivists cannot outpace obsolescence, 2017; <http://bit.ly/2KDaqWd>.
20. Spectra. Arcticblue deep storage disk. Product, <https://www.spectrallogic.com/products/arcticblue/>.
21. SpectraLogic. Spectralogic t50e; <http://bit.ly/2Ych8pl>.
22. StorageReview. Intel optane memory review. http://www.storagereview.com/intel_optane_memory_review.
23. TPC-C. Dell-microsoft sql server tpc-c executive summary, 2014; http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=114112501.
24. Ultrium. LTP ultrium roadmap;
25. <http://www.lto-ultrium.com/lto-ultrium-roadmap/>.
26. Umamageswaran, K. and Goindi, G. Exadata: Delivering memory performance with shared flash; <http://bit.ly/2LhBVEa>.
27. Yan, M. Open compute project: Cold storage hardware v0.5, 2013; <http://bit.ly/2X6H2Ot>.
28. Yan, W., Yao, J., Cao, Q., Xie, C., and Jiang, H. Ros: A rack-based optical storage system with inline accessibility for long-term data preservation. In *Proceedings of EUROSYS*, 2017.

Raja Appuswamy (raja.appuswamy@eurecom.fr) is an assistant professor in the Data Science Department at EURECOM, Biot, Provence-Alpes-Côte d’Azur, France.

Goetz Graefe (goetzg@google.com), Google, Inc., Madison, WI, USA.

Renata Borovica-Gajic (renata.borovica@unimelb.edu.au) is an assistant professor in the School of Computing and Information Systems at the University of Melbourne, Australia.

Anastasia Ailamaki (anastasia.ailamaki@epfl.ch) is a professor at EPFL, Lausanne, Switzerland, and director of its Data-Intensive Applications and Systems (DIAS) lab.