Persistent Buffer Management with Optimistic Consistency

Lucas Lersch
lucas.lersch@sap.com
TU Dresden & SAP SE

Wolfgang Lehner
wolfgang.lehner@tu-dresden.de
TU Dresden

Ismail Oukid
ismail.oukid@sap.com
SAP SE

ABSTRACT

Finding the best way to leverage non-volatile memory (NVM) on modern database systems is still an open problem. The answer is far from trivial since the clear boundary between memory and storage present in most systems seems to be incompatible with the intrinsic memory-storage duality of NVM. Rather than treating NVM either solely as memory or solely as storage, in this work we propose how NVM can be simultaneously used as both in the context of modern database systems. We design a persistent buffer pool on NVM, enabling pages to be directly read/written by the CPU (like memory) while recovering corrupted pages after a failure (like storage). The main benefits of our approach are an easy integration in the existing database architectures, reduced costs (by replacing DRAM with NVM), and faster peak-performance recovery.

1 INTRODUCTION

NVM is a persistent media promising higher bandwidth (2×) and lower latency (10×) than modern NAND-flash SSDs. Furthermore, NVM can be attached to the memory bus, thus allowing it to be directly accessed by the CPU through its caches in a much smaller granularity (cache-lines) than regular block devices. Therefore, NVM introduces not only a new layer within the storage hierarchy [3], but it also enables more flexibility regarding data placement.

While a few factors slower, reading data directly from NVM can be done the same way as with DRAM. However, writing data directly to NVM imposes challenges in terms of data consistency. The root of these challenges is the lack of control the application has over data movement between CPU cache and NVM in comparison to the data movement between DRAM and SSD. In other words, the programmer cannot prevent cache lines from being evicted from the CPU cache and written-back to NVM at arbitrary points in time.

Related work address these challenges with solutions that fall in one of three categories (also identified by previous work [16]). First, early proposals leverage NVM as a cheaper alternative to extend DRAM, enabling larger buffer pools [13, 17]. These approaches

focus on reducing write amplification and improving wear leveling on NVM, but they do not enforce any consistency when writing data, and therefore do not leverage persistency. Second, persistent data structures [1, 4, 9, 14, 18] enable direct fine-grained writes to NVM by issuing out-of-place writes and relying on instruction ordering (*SFENCE*) and eagerly flushing cache lines (*CLFLUSH/CLWB*) to make the operation visible (usually by flipping a validity bit). Third, buffer managers were proposed to integrate NVM with the storage hierarchy below DRAM and above SSD [5, 15, 16]. These approaches access data on DRAM, and therefore they have full control of when data is persisted. Optimizations can be made regarding block sizes, but they still impose movement of data between DRAM and NVM.

The first two categories treat NVM more like traditional memory, while the last one treats NVM more like traditional storage. However, NVM is actually both. We consider that a database system should allow NVM to be read and written directly, thus exploring its memory-storage duality to its full potential. Our system accesses NVM like memory, but it guarantees consistency of writes to NVM like storage. We achieve that by integrating NVM in the buffer pool of a database system to either extend DRAM or completely replace it, while still leveraging its persistency in an optimistic way. In other words, we never enforce ordering of writes or eagerly flush cache lines. The main motivation is that corruption occurs when a write operation is partially evicted from the CPU cache to NVM. With the capacity of NVM being significantly higher than CPU caches, the probability of corruption tends to be low and therefore pessimistically enforcing consistency of every write introduces a relatively high overhead by not leveraging the CPU cache.

Recent work [2] has gone into the same direction of allowing data to be accessed directly on NVM. This approach complements our approach by focusing on cost models for optimizing data movement, while we focus on enabling direct writes to NVM to be consistent.

2 BACKGROUND

We give an overview of the techniques that we use to implement a persistent buffer management in a database system.

Buffer Management: We assume a traditional transactional storage manager having a B+Tree as its workhorse data structure. The B+tree is organized such that nodes are represented by pages, which are the unit of data movement and buffering. Optimizations like pointer swizzling [8] and low-overhead replacement policies [10] may apply. The atomicity and durability of writes to pages buffered in DRAM is guaranteed by write-ahead logging (WAL).

Single-page Recovery: Similar to traditional ARIES [12], we assume a page-level physiological logging. This implies that pages are not only a unit of data movement, but also of fault containment and repair [7]. This enables techniques such as write-elision, on-demand instant restart and restore, and single-page repair [6]. These techniques are the base to enable direct writes directly made to NVM to be consistent without eagerly flushing cache lines.

Figure 1: During normal processing reads and writes can access both DRAM and NVM directly (dashed lines). Whole pages can be moved across any devices (solid lines). The trade-offs between DRAM and NVM can be analysed by sliding an abstract persistency bar.

3 SYSTEM DESIGN

The pages of our system are primarily located on SSD and only the warm pages are buffered in NVM. Hot pages might be buffered in a DRAM portion of the buffer pool, as seen in Figure 1.

Normal Processing: During normal transactional processing, a page to be updated will be either on SSD, NVM, or DRAM. In case the page is on SSD, it is loaded to the buffer pool (either to DRAM or NVM, to be decided by a placement policy). In case the page is on DRAM, we have a hit and the page is updated normally. If the page is on NVM, two actions might occur. First, the page might be identified as "heating up" by a placement policy such as 2Q [11], in which case it will be copied to DRAM and updated there. Second, the page might be simply warm, in which case the update is done directly on NVM. In the last case, since atomicity and durability is guaranteed by WAL, issuing *CLFLUSH/CLWB* after updating the page is not necessary as the log serves as the single source of truth. However, when restarting after a system failure, pages that were on NVM might be corrupted because updates were not properly persisted. As a consequence, the current state of a page on NVM is unknown and therefore persistency of NVM cannot be leveraged.

We address this challenge by dividing it in two dimensions: corruption detection and page repair. Each page contains an 8 Byte checksum of the whole page. Whenever the content of a page is modified directly on NVM, the page checksum is re-calculated and updated. At this point, the modifications and the checksum of a page might be persistent or not, since we do not explicitly flush them from the CPU cache to NVM. In case the overhead of calculating the checksum for the whole page after every update is too high, it can be reduced either by updating the checksum only after a certain number of updates (at the cost of higher corruption ratio) or by introducing multiple checksums per page corresponding to fractions of the page (at the cost of higher space consumption).

Corruption Detection After a system failure, the restart process starts with log analysis, which identifies the state (not the content) of pages right before the crash. We assume state-of-the-art page-based on-demand restart [6], therefore a page being requested might still be lingering in the NVM portion of the buffer pool since before the crash. Two steps are required to determine if the page can be used. First, its checksum is calculated and compared to the checksum stored within the page. If the checksums do not match, it is because either the checksum or the updated data were not evicted from the CPU cache to NVM. Second, if the checksums match, the

log sequence number of the last update in the page (*pageLSN*) is compared to the LSN retrieved during log analysis (*expectedLSN*). To summarize, a page might be in one of the following states:

- Corrupted, if checksum does not match
- Behind, if checksum matches and $pageLSN < expectedLSN$
- Current, if checksum matches and $pageLSN = expectedLSN$
- Ahead, if checksum matches and $pageLSN > expectedLSN$

The page is *behind*, if it was updated by a committed transaction (the log records were flushed) but neither the update nor the checksum were evicted from the CPU cache. In this case, the page is in a consistent but outdated state since it violated the write-ahead rule. However, this violation is tolerated on NVM as long as it is guaranteed not to happen on SSD. Analogously, the page is *ahead* if both update and checksum were persisted to NVM but the transactions that made these updates did not commit. Finally, if both checksum and LSNs match, the page is in its most *current* state and ready to be accessed. Except for the *current* state, all other states must be recovered by replaying log records.

Page Repair No assumption can be made about a *corrupted* page, and therefore the only alternative is to retrieve its more recent version from SSD (which acts as a backup) and replay the log records referring to this page to bring it up to its most recent state. A *behind* page is missing committed updates, therefore it can be used as the basis for replaying the log records, not requiring an older version of the page to be fetched from SSD. *Ahead* pages are consistent but contain updates made by uncommitted transactions. Since ARIES-style recovery only allows logical UNDO on a transaction level, the updates of a single page cannot be independently rolled back. Therefore, the same procedure of a *corrupted* page applies for rolling forward. Since it is required that an older version of the page is read from the database, it is convenient that frequently modified pages are flushed to the database often to bound the recovery time in such cases. Fortunately, this can be achieved by an asynchronous cleaning job that iterates through the buffer pool and picks dirty pages to be flushed to the database.

4 CONCLUSION

In this work we propose extending database systems with a persistent buffer pool on NVM. We consider such extension to impose a low implementation effort, since NVM is treated very similarly to DRAM during runtime, while leveraging well-understood recovery algorithms to enforce its consistency. The consequence is that the persistency aspect can be leveraged in an optimistic way, without major changes in the code base to enforce ordering of writes to NVM. Furthermore, cache lines do not have to be eagerly flushed, which enables write operations to be cached in the CPU and amortize the higher latencies of NVM. We also consider that in the short and medium term, NVM will not replace neither DRAM or SSD, but should act in synergy with both. Our design also enables the user to explore different proportions of DRAM and NVM in the buffer pool: more DRAM will lead to higher performance, while more NVM will lead to reduced costs and faster recovery (higher availability). This adds more flexibility for analyzing trade-offs and navigating the performance continuum between modern expensive in-memory databases and more traditional low-cost disk-based systems.

REFERENCES

[1] Joy Arulraj, Justin Levandoski, Umar Farooq Minhas, and Per-Ake Larson. 2018. BzTree: A High-Performance Latch-free Range Index for Non-Volatile Memory. *PVLDB* 11, 5 (2018), 553–565.

[2] Joy Arulraj, Andy Pavlo, and Krishna Teja Malladi. 2019. Multi-Tier Buffer Management and Storage System Design for Non-Volatile Memory. *CoRR* abs/1901.10938 (2019). arXiv:1901.10938 http://arxiv.org/abs/1901.10938

[3] Philippe Bonnet. 2017. What's up with the storage hierarchy?. In *CIDR*.

[4] Shimin Chen and Qin Jin. 2015. Persistent B+-trees in Non-Volatile Main Memory. *PVLDB* 8, 7 (2015), 786–797.

[5] Assaf Eisenman, Darryl Gardner, Islam AbdelRahman, Jens Axboe, Siying Dong, Kim Hazelwood, Chris Petersen, Asaf Cidon, and Sachin Katti. 2018. Reducing DRAM Footprint with NVM in Facebook. In *Proceedings of the Thirteenth EuroSys Conference*. ACM, 42.

[6] Goetz Graefe, Wey Guy, and Caetano Sauer. 2014. Instant Recovery with Write-Ahead Logging: Page Repair, System Restart, and Media Restore. *Synthesis Lectures on Data Management* 6, 5 (2014), 1–85.

[7] Goetz Graefe and Harumi A. Kuno. 2012. Definition, Detection, and Recovery of Single-Page Failures, a Fourth Class of Database Failures. *PVLDB* 5, 7 (2012), 646–655.

[8] Goetz Graefe, Haris Volos, Hideaki Kimura, Harumi Kuno, Joseph Tucek, Mark Lillibridge, and Alistair Veitch. 2014. In-memory Performance for Big Data. *PVLDB* 8, 1 (2014), 37–48.

[9] Se Kwon Lee, K Hyun Lim, Hyunsub Song, Beomseok Nam, and Sam H Noh. 2017. WORT: Write Optimal Radix Tree for Persistent Memory Storage Systems. In *15th USENIX Conference on File and Storage Technologies (FAST 17)*. 257–270.

[10] Viktor Leis, Michael Haubenschild, Alfons Kemper, and Thomas Neumann. 2018. LeanStore: In-Memory Data Management Beyond Main Memory. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. IEEE, 185–196.

[11] Lucas Lersch, Ismail Oukid, Ivan Schreter, and Wolfgang Lehner. 2017. Rethinking DRAM Caching for LSMs in an NVRAM Environment. In *European Conference on Advances in Databases and Information Systems*. Springer, 326–340.

[12] C Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, and Peter Schwarz. 1992. ARIES: A Transaction Recovery Method Supporting Fine-Granularity Locking and Partial Rollbacks Using Write-Ahead Logging. *ACM Transactions on Database Systems (TODS)* 17, 1 (1992), 94–162.

[13] Yi Ou, Lei Chen, Jianliang Xu, and Theo Härder. 2014. Wear-Aware Algorithms for PCM-Based Database Buffer Pools. In *International Conference on Web-Age Information Management*. Springer, 165–176.

[14] Ismail Oukid, Johan Lasperas, Anisoara Nica, Thomas Willhalm, and Wolfgang Lehner. 2016. FPTree: A Hybrid SCM-DRAM Persistent and Concurrent B-tree for Storage Class Memory. In *Proceedings of the 2016 International Conference on Management of Data*. ACM, 371–386.

[15] Steven Pelley, Thomas F Wenisch, Brian T Gold, and Bill Bridge. 2013. Storage Management in the NVRAM Era. *PVLDB* 7, 2 (2013), 121–132.

[16] Alexander van Renen, Viktor Leis, Alfons Kemper, Thomas Neumann, Takushi Hashida, Kazuichi Oe, Yoshiyasu Doi, Lilian Harada, and Mitsuru Sato. 2018. Managing Non-Volatile Memory in Database Systems. In *Proceedings of the 2018 International Conference on Management of Data*. ACM, 1541–1555.

[17] Zhangling Wu, Peiquan Jin, Chengcheng Yang, and Lihua Yue. 2014. APP-LRU: A New Page Replacement Method for PCM/DRAM-Based Hybrid Memory Systems. In *IFIP International Conference on Network and Parallel Computing*. Springer, 84–95.

[18] Jun Yang, Qingsong Wei, Cheng Chen, Chundong Wang, Khai Leong Yong, and Bingsheng He. 2015. NV-Tree: Reducing Consistency Cost for NVM-based Single Level Systems. In *13th USENIX Conference on File and Storage Technologies (FAST 15)*. 167–181.