Optimizing Space Amplification in RocksDB

Siying Dong[1], Mark Callaghan[1], Leonidas Galanis[1],
Dhruba Borthakur[1], Tony Savor[1] and Michael Stumm[2]

[1]Facebook, 1 Hacker Way, Menlo Park, CA USA 94025
{siying.d, mcallaghan, lgalanis, dhruba, tsavor}@fb.com

[2]Dept. Electrical and Computer Engineering, University of Toronto, Canada M8X 2A6
stumm@eecg.toronto.edu

ABSTRACT

RocksDB is an embedded, high-performance, persistent key-value storage engine developed at Facebook. Much of our current focus in developing and configuring RocksDB is to prioritize resource efficiency over more standard performance metrics, such as response time latency and throughput, as long as they remain acceptable. In particular, we optimize space efficiency as long as read/write latencies are able to meet target service-level requirements for the intended workloads, because storage space is most often the bottleneck when using Flash SSDs under typical production workloads at Facebook. RocksDB uses *Log-structured Merge Trees* (LSM) to obtain significant space efficiency and better write throughput while still being good at read performance.

We describe how we apply a number of approaches to reduce storage usage in RocksDB. We discuss how we are able to trade off storage efficiency and CPU overhead, as well as read and write amplification. Based on results of experimental evaluations of MySQL with RocksDB as the embedded storage engine (using TPC-C and LinkBench benchmarks) and based on measurements taken from production databases, we show that RocksDB uses less than half the storage that *InnoDB* uses, yet performs well and in many cases even better than the B-tree-based Innodb storage engine. To the best of our knowledge, this is the first time an LSM-based storage engine has shown competitive performance when running OLTP workloads at large scale.

1. INTRODUCTION

At Facebook (FB), we have made resource efficiency the primary objective in our storage systems strategy: performance needs to be sufficient to meet the needs of our services, but efficiency should be as good as possible to allow for scale.

FB has one of the largest *MySQL* installations in the world, storing many 10s of petabytes of online data. The underlying storage engine for the FB MySQL instances is increasingly being switched over from *InnoDB* to MyRocks, which in turn is based on RocksDB. The switchover is primarily motivated by the fact that MyRocks uses half the storage InnoDB needs and has higher average transaction throughput, yet has only marginally worse read latencies.

RocksDB is an embedded, high-performance, persistent key-value storage system [1] that was developed by FB after forking the code from Google's LevelDB [2]. It was open-sourced in 2013. *MyRocks* is a variant of RocksDB but configured and tuned to be MySQL's storage engine.

RocksDB is used in many applications beyond just MySQL, both within and outside of FB. Within FB, RocksDB is used as a storage engine for *Laser*, a high query throughput, low latency key-value storage service [3], *ZippyDB*, a distributed key-value store with Paxos-style replication [3], *Dragon* to store indices of the Social Graph [4], and *Stylus*, a stream processing engine [3], to name a few. Outside of FB, both *MongoDB* [5] and *Sherpa*, Yahoo's largest distributed data store [6], use RocksDB as one of their storage engines. Further, RocksDB is used by *LinkedIn* for storing user activity and by *Netflix* to cache application data, to list a few examples.

With MyRocks, each MySQL table row is stored as a RocksDB key-value pair: the primary keys are encoded in the RocksDB key and all other row data is encoded in the value. Secondary keys, which are not necessarily unique, are stored as separate key-value pairs, where the RocksDB key encodes the secondary key appended with the corresponding target primary key, and the value is left empty; thus secondary index lookups are translated into RocksDB range queries. All RocksDB keys are prefixed by a 4-byte table-ID or index-ID so that multiple tables or indexes can coexist in one RocksDB key space. Finally, a global sequence ID, incremented with each write operation, is stored with each key-value pair to support snapshots. Snapshots are used to implement multiversion concurrency control, which in turn enables us to implement ACID transactions within RocksDB.

Our primary goal with RocksDB at FB is to make the most efficient use of hardware resources while ensuring all important service level requirements can be met, including target transaction latencies. This focus on efficiency instead of performance is perhaps unique in the community in that database systems are typically compared using performance metrics such as transactions per minute (e.g., tpmC) or response-time latencies. Our focus on efficiency does not imply that we treat performance as unimportant, but rather that once our performance objectives are achieved, we op-

CIDR '17 January 8–11, 2017, Santa Cruz, CA, USA

DOI: 10.1145/1235

timize for efficiency. Our approach is driven in part by the data storage needs at FB (that may differ from that of other organizations):

1. SSDs are increasingly being used to store persistent data and are the primary target for RocksDB;
2. FB primarily relies on shared nothing configurations of commodity hardware in their data centers [7], where data is distributed across a large number of simple nodes, each with 1-2 SSDs;
3. the amount of data that needs to be stored is huge;
4. the read-write ratio is relatively low at roughly 2:1 in many (but not all) cases, given the fact that large memory-based caches are used extensively.

Minimizing space amplification is the key to efficient hardware use because storage space is the bottleneck in environments like the one described above. In a typical production MySQL environment at FB, SSDs service far fewer reads/s and writes/s during peak times under InnoDB than what the hardware is capable of. The throughput level under InnoDB is low, not because of any bottleneck on the SSD or the processing node — e.g., CPU utilization remains below 40% — but because the query rate per node is low. The per node query rate is low, because the amount of data that has to be stored (and be accessible) is so large, it has to be sharded across many nodes to be able to fit, and the more nodes, the fewer queries per node.

If the SSD could store twice as much data, then we would expect storage node efficiency to double, since the SSDs could easily handle the expected doubling of IOPS, and we would need far fewer nodes for the workload. Hence our focus on space amplification. Moreover, minimizing space amplification will make SSDs become an increasingly attractive alternative to spinning disks for colder data storage, as SSD prices continue to sink. In our pursuit to minimize space amplification, we are willing to trade off some extra read or write amplification. Such a tradeoff is necessary because it is not possible to reduce space, read, and write amplification together, as shown by Athanassoulis et al. [8].

In this paper, we describe our techniques for reducing space amplification within RocksDB. Some of the techniques, we believe, are being described for the first time, i.e., dynamic LSM level size adjustment; tiered compression; shared compression dictionary; prefix bloom filter; and different size multipliers at different LSM levels. We discuss how the space amplification techniques affect read and write amplification, and the various tradeoffs involved. Using empirical measurements, we show that RocksDB requires roughly 50% less storage space than InnoDB on average and has a higher transaction throughput, yet increases read latencies only marginally, remaining well within the margins of acceptability. We also discuss tradeoffs between space amplification and CPU overhead, since the CPU may become a bottleneck once space amplification is significantly reduced.

RocksDB is based on Log-Structured Merge-Trees (LSM). The LSM tree was originally designed to minimize random writes to storage as it never modifies data in place, but only appends data to files located in stable storage to exploit the high sequential write speeds of hard drives [9]. As technology changed, LSM trees became attractive because of their low write amplification and low space amplification characteristics. We demonstrate that an LSM-based storage engine can be performance competitive when used on OLTP workloads, which we believe is the first time this is shown.

In the next section, we provide brief background on LSM trees for those unfamiliar with them and then describe our techniques to reduce space amplification in §3. We describe in §4 how to balance space amplification with read amplification and CPU overhead. Finally, in §5 we present the results of experimental evaluations using realistic production workloads (TPC-C and LinkBench) and measurements taken from production instances of the database. We close with concluding remarks.

2. LSM-TREE BACKGROUND

Log Structured Merge Trees (LSM trees) are used in many popular systems today, including BigTable [10], LevelDB, Apache Cassandra [11], and HBase [12]. Here we briefly describe the LSM-tree as implemented and configured in MyRocks at FB by default.

Whenever data is written to the LSM-tree, it is added to an in-memory write buffer called *mem-table*, implemented as a skiplist having O(log n) inserts and searches, and at the same time appended to a Write Ahead Log (WAL) for recovery purposes. If, after a write, the size of the mem-table reaches a predetermined size, then (*i*) the current WAL and mem-table become immutable, and a new WAL and mem-table are allocated for capturing subsequent writes, (*ii*) the contents of the mem-table are flushed out to a "*Sorted Sequence Table*" (SST) data file, and when that is complete, (*iii*) the WAL and mem-table containing the data just flushed are discarded. This general approach has a number of favorable consequences: new writes can be processed concurrently to the flushing of an older mem-table; all I/O is sequential,[1] and except for the WAL, only entire files are ever written.

Each of the SSTs stores data in sorted order, divided into unaligned 16KB blocks (when uncompressed). Each SST also has an index block for binary search with one key per SST block. SSTs are organized into a sequence of levels of increasing size, Level-0 – Level-N, where each level will have multiple SSTs. Level-0 is treated specially in that its SSTs may have overlapping key ranges, while the SSTs of higher levels have distinct non-overlapping key ranges. When the number of files in Level-0 exceeds a threshold (e.g., 4), then the Level-0 SSTs are merged with the Level-1 SSTs that have overlapping key ranges; when completed, all of the merge sort input (L0 and L1) files are deleted and replaced by new (merged) L1 files. For L>0, when the combined size of all SSTs in level-L exceeds a threshold (e.g., $10^{(L-1)}$GB) then one or more level-L SSTs are selected and merged with the overlapping SSTs in level-(L+1), after which the merged level-L and level-(L+1) SSTs are removed.

The merging process, implemented using multiple threads, is called *compaction*, as it removes data marked as deleted and data that has been overwritten (if it is no longer needed).[2] Compaction also has the effect of gradually migrating new updates from Level-0 to the last level, which is why this particular approach is referred to as "leveled" compaction.

[1]There are usually concurrent streams of sequential IO that will cause seeks. However, the seeks will be amortized over LSM-tree's very large writes (many MB rather than KB).

[2]We described *leveled compaction* here, which is different than the compaction methods used by HBase and Cassandra [13]. In this paper, all uses of the term *compaction* refer to leveled compaction.

The process ensures that at any given time, each SST will contain at most one entry for any given key and snapshot. The I/O that occurs during compaction is efficient, as it only involves bulk reads and writes of entire files: if a level-L file being compacted overlaps with only part of a level-(L+1) file, then still the entire level-(L+1) file is used as an input to the compaction and ultimately removed. A compaction may trigger a set of cascading compactions.

A single *manifest file* maintains a list of SSTs at each level, their corresponding key ranges, and some other meta data. It is maintained as a log to which changes to the SST information are appended. Its information is cached in an efficient format in memory to enable quick identification of SSTs that may contain a target key.

A search for a key continues across the levels until the key is found or it is determined that the key is not present in the last level. First, all mem-tables are searched, followed by all Level-0 SSTs and then the SST's at next following levels. At each of these following levels, three binary searches are necessary: one to locate the target SST by searching the Manifest data, one to find the target data block within the SST file by using the index block, and one to search for the key within the data block. Bloom filters (kept in files but cached in memory) are used to eliminate unnecessary SST searches, so mostly only 1 data block needs to be read from disk. Moreover, recently read SST blocks are cached in a block cache maintained by RocksDB and the operating system's page cache, so access to recently fetched data need not result in I/O operations. The MyRocks block cache is typically configured to be 12GB large.

Range queries are more involved and always require a search through all levels since all keys that fall within the range must be located. First the mem-table is searched for keys within the range, then all Level-0 SSTs, followed by all subsequent levels, while disregarding duplicate keys within the range from lower levels. Prefix Bloom filters (§4) can reduce the number of SSTs that need to be searched.

3. SPACE AMPLIFICATION

LSM is typically far more space efficient than a B-tree. Under read/write workloads similar to those at FB, B-tree space utilization will be poor [14] with its pages only 1/2 to 2/3 full (as measured in FB production databases). This fragmentation causes space amplification to be worse than 1.5 in B-tree-based storage engines. Compressed InnoDB has fixed page sizes on disk which further wastes space.

In contrast, LSM does not suffer from fragmentation because it does not require data to be written to SSD page-aligned. LSM space amplification is mostly determined by how much stale data is yet to be garbage-collected. If we assume that the last level is filled to its target size with data and that each level is 10X larger than the previous level, then in the worst case, LSM space amplification will be 1.111..., considering that all of the levels up to the last level combined are only 11.111...% the size of the last level.

RocksDB uses two strategies to reduce space amplification: (i) adapting the level sizes to the size of the data, and (ii) applying a number of compression strategies.

3.1 Dynamic level size adaptation

If a fixed size is specified for each level, then in practice it is unlikely that the size of the data stored at the last level will be 10X the target size of the previous level. In a worse case, the size of the data stored at the last level will only be slightly larger than the target size of the previous level, in which case space amplification would be larger than 2.

However, if we dynamically adjust the size of each level to be 1/10-th the size of the data on the next level, then space amplification will be reduced to less than 1.111....

The level size multiplier is a tunable parameter within LSM. Above, we assumed it is 10. The larger the size multiplier is, the lower the space amplification and the read amplification, but the higher the write amplification. Hence, the choice represents a tradeoff. For most of the FB production RocksDB installations, a size multiplier of 10 is used, although there are a few instances that use 8.

An interesting question is whether the size multiplier at each level should be the same. The original paper on LSM proved that it is optimal to have the same multiplier at each level when optimizing for write amplification [9].[3] It is an open question of whether this also holds true when optimizing for space amplification, especially when considering that different levels may use different compression algorithms resulting in different compression ratios at each level (as described in the next section). We intend to analyze this question in future work.[4]

3.2 Compression

Space amplification can be further reduced by compressing the SST files. We apply a number of strategies simultaneously. LSM provides a number of properties that make applying compression strategies more practical. In particular, SSTs and their data blocks in LSM are immutable.

Key prefix encoding. Prefix encoding is applied on the keys by not writing repeated prefixes of previous keys. We have found this reduces space requirements by 3% – 17% in practice, depending on the data workload.

Sequence ID garbage collection. The sequence ID of a key is removed if it is older than the oldest snapshot needed for multiversion concurrency control. Users can arbitrarily create snapshots to refer to the current database state at a later point in time. Removing snapshot IDs tends to be effective because the 7 byte large sequence ID does not compress well, and because most of the sequence IDs would no longer be needed after the corresponding snapshots that refer to them have been deleted. In practice, this optimization reduces space requirements, depending on the data workload, from between 0.03% (e.g., for a database storing social graph vertexes that will have large values) and 23% (e.g., for a database storing social graph edges that will have empty values) .

Data compression. RocksDB currently supports several compression algorithms, including LZ, Snappy, zlib, and zstandard. Each level can be configured to use any or none of these compression algorithms. Compression is applied on a per-block basis. Depending on the composition of the data, weaker compression algorithms can reduce space requirements down to as low as 40%, and stronger algorithms

[3]The original LSM paper uses a fixed number of levels and varies the multiplier as the database gets larger, but keeping the multiplier the same at each level. LevelDB/RocksDB use a fixed multiplier but varies the number of levels as the database gets larger.

[4]Initial results indicate that adapting the size targets at each level to take into account the compression ratios achieved at each level lead to better results.

down to as low as 25%, of their original sizes on production FB data.

To reduce the frequency of having to uncompress data blocks, the RocksDB block cache stores blocks in uncompressed form. (Note that recently accessed compressed file blocks will be cached by the operating system page cache in compressed form, so compressed SSTs will use less storage space and less cache space, which in turn allows the file system cache to cache more data.)

Dictionary-Based Compression. A data dictionary can be used to further improve compression. Data dictionaries can be particularly important when small data blocks are used, as smaller blocks typically yield lower compression ratios. The dictionary makes it possible for smaller blocks to benefit from more context. Experimentally, we have found that a data dictionary can reduce space requirements by an additional 3%.

LSM trees make it easier to build and maintain dictionaries; they tend to generate large immutable SST files that can be hundreds of megabytes large; a dictionary that is applied to all data blocks can be stored within the file so when the file is deleted, the dictionary is dropped automatically.

4. TRADEOFFS

LSM-trees have many configuration parameters and options that allow an installation to control a number of tradeoffs. Prior work by Athanassoulis et al. established that one can optimize for any two of space-, read-, and write-amplification, but at the cost of the third [8]. Thus, for example, increasing the number of levels (say by decreasing the level multiplier) decreases write amplification, but increases space and read amplification.

As another example, in LSM trees, a larger block size leads to improved compression without degrading write amplification, but negatively affects read amplification (since more data must be read per query). This observation allows us to opt for a larger block size for better compression ratios when dealing with write heavy applications. (In B-Trees, larger blocks degrade both write and read amplification.)

Tradeoffs in many cases involve judgement calls and depend on the expected workload and perceived minimal acceptable quality of service levels. When focusing on efficiency (as we do), it is exceedingly difficult to configure the system to properly balance CPU, disk I/O, and memory utilization, especially because it is strongly dependent on a highly varying workload.

As we show in the next section, our techniques reduce storage space requirements by 50% over InnoDB. This allows us to store twice as much data on each node, which in turn enables significant consolidation of existing hardware. But that also means we double the workload (QPS) per server and must be careful that we have enough CPU, random I/O capacity and RAM to support that.

Tiered compression. Compression generally decreases the amount of storage space required, but increases CPU overheads, since data has to be compressed and decompressed. The stronger the compression, the higher the CPU overhead. In our installations, a strong compression algorithm (like zlib or zstandard) is typically used at the last level even though it incurs higher CPU overhead, because most (close to 90%) of the data is located at that level, yet only a small fraction of reads and writes go to it. In various use cases, applying strong compression to the last level saves an additional 15%–30% in storage space over using lightweight compression only.

Conversely, we do not use any compression at levels 0–2 to allow for lower read latencies at the cost of higher space- and write-amplification, because (*i*) they tend to be accessed more frequently, and (*ii*) they use up only a small proportion of the total storage space. Level-3 up to the lest level use lightweight compression (like LZ4 or Snappy) because its CPU overhead is acceptable, yet it reduces space and write amplification. Reads to data located in the first three levels will more likely be located in (uncompressed) file blocks cached by the operating system because these blocks are frequently accessed. However, reads to data located in levels higher than 2 will have to be uncompressed whether they are located in the operating system file cache or not (unless they are also located in the RocksDB block cache).

Bloom filters. Bloom filters are effective in reducing I/O operations and attendant CPU overheads, but at the cost of somewhat increased memory usage since the filter (typically) requires 10 bits per key. However, as an illustration that some tradeoffs are subtle, we do not use a Bloom filter at the last level. While this will result in more frequent accesses to last-level files, the probability of a read query reaching the last level is relatively small. More importantly, the last-level bloom filter is large (∼10X as large as all lower-level Bloom filters combined) and the space it would consume in the memory-based caches, would prevent the caching of other data that would be being accessed. Measurements indicated that not having a Bloom filter for the last level improved read amplification overall, given our workloads.

Prefix Bloom filters. Bloom filters do not help with range queries. We have developed a prefix Bloom filter that helps with range queries, based on the observation that many range queries are often over a prefix; e.g., the userid part of a (userid,timestamp) key or postid of a (postid,likerid) key. We allow users to define prefix extractors to deterministically extract a prefix part of the key from which we construct a Bloom filter. When querying a range, the user can specify that the query is on a defined prefix. We have found this optimization reduces read amplification (and attendant CPU overheads) by up to 64% on our systems for otherwise costly range queries.

5. EVALUATION

A review of numerous MySQL installations at FB generally reveal that

1. the storage space used by RocksDB is about 50% lower than the space used by InnoDB with compression,
2. the amount of data written to storage by RocksDB is between 10% and 15% of what InnoDB writes out, and
3. the number and volume of reads is 10% – 20% higher in RocksDB that for InnoDB (yet well within the margin of acceptability).

For more meaningful metrics from controlled environments, we present the results of extensive experiments using two benchmarks with MySQL. The first benchmark, LinkBench, is based on traces from production databases that store "social graph" data at FB; it issues a considerable number of range queries [15]. We ran 24 1 hour intervals of LinkBench and measured statistics for the 24th interval to obtain numbers from a steady-state system.[5] The second benchmark is

[5]We also gathered statistics when loading the full LinkBench

the standard TPC-C benchmark.

For both benchmarks, we experimented with two variants: one where the database fit in DRAM so that disk activity was needed only for writes to achieve durability, and one where the database did not fit in memory. We compare the behavior of RocksDB, InnoDB, and TokuDB, configured to use a variety of compression strategies. (TokuDB is another open source, high-performance storage engine for MySQL which at its core uses a fractal tree index tree data structure to reduce space and write amplification [16].)

Figure 1 shows the results from our LinkBench experiments. The hardware and software setup used is described

---

database; the results (not shown due to space) are inline with the steady-state numbers.

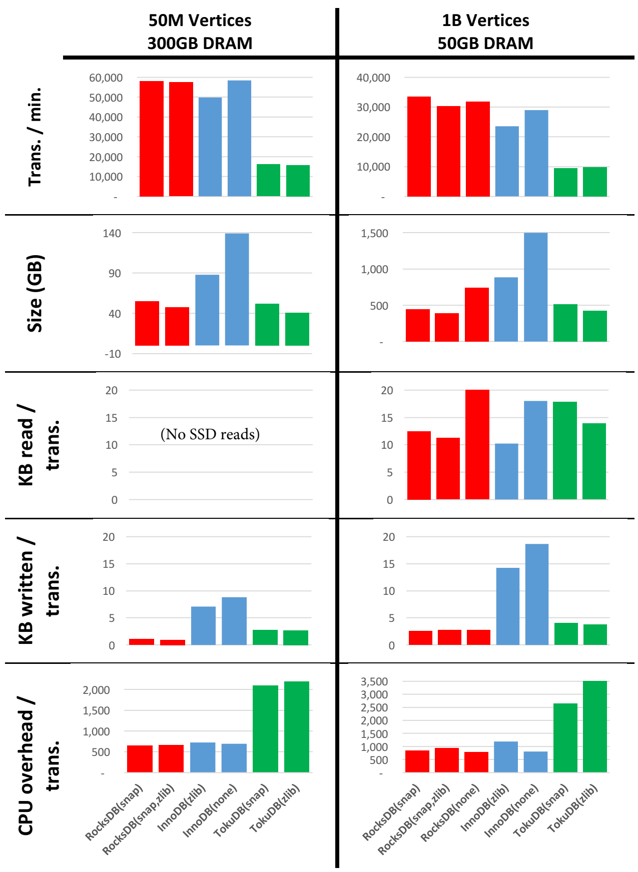

Figure 1: LinkBench benchmark. Statistics gathered from the 24th hour of 24 hour runs with 16 concurrent clients for 3 different storage engines: RocksDB (from FB MySQL 5.6) shown in red, InnoDB (from MySQL 5.7.10) shown in blue, and TokuDB (Percona Server 5.6.26-74.0) shown in green, configured to use the compression scheme(s) listed in brackets. (Sync-on-commit was disabled, binlog/oplog and redo logs were enabled.)

System setup: The hardware consisted of an Intel Xeon E5-2678v3 CPU with 24-cores/48-HW-threads running at 2.50GHz, 256GB of RAM, and roughly 5T of fast NVMe SSD provided via 3 devices configured as SW RAID 0. The operating system was Linux 4.0.9-30

Left hand side graphs: Statistics from LinkBench configured to store 50M vertices, which fits entirely in DRAM.

Right hand side graphs: Statistics from LinkBench configured to store 1B vertices, which does not fit in memory after constraining DRAM memory to 50GB: all but 50GB of RAM was mlocked by a background process so the database software, OS page cache and other monitoring processes had to share those 50GB. The MyRocks block cache was set to 10GB.

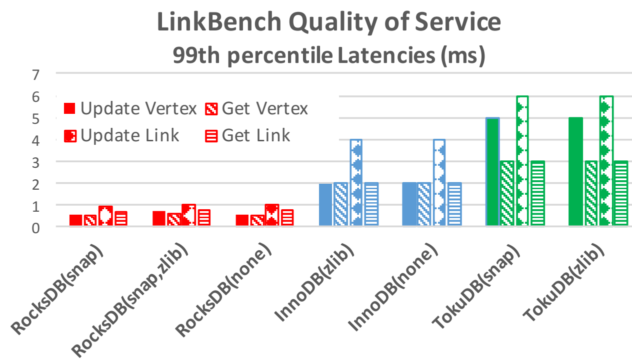

Figure 2: 99th percentile latencies for: Update Vertex, Get Vertex, Update Link, Get Link. The setup of the hardware and storage engines as described in Figure 1. The database with 1B vertices was used with available DRAM constrained to 50GB.

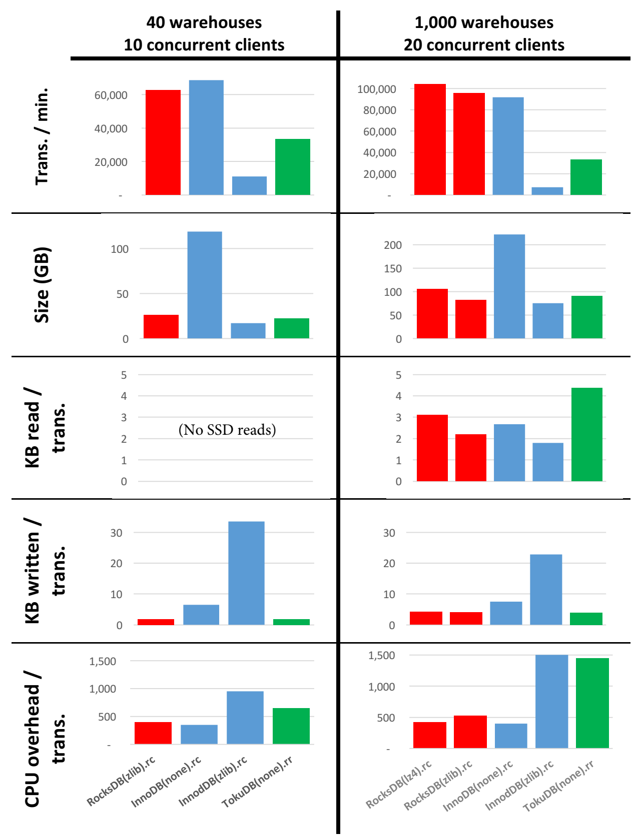

Figure 3: TPC-C. Metrics obtained using the same setup as in Figure 1. The left hand side: configuration of 40 warehouses and 10 concurrent clients; the database fits in memory. The statistics were gathered during the 15th hour interval after 14 hours of operation. The right hand side: configuration of 1,000 warehouses and 20 concurrent clients. The statistics shown were gathered during the 12th hour interval after 11 hours of operation. The transaction isolation levels used is marked as "rc" for READ COMMITTED or"rr" for REPEATABLE READ.

in the figure caption. Some observations for the LinkBench benchmark with a database that does not fit in memory:

- **Space usage:** RocksDB with compression uses less storage space than any of the alternatives considered; without compression, it uses less than half as much storage space as InnoDB without compression.
- **Transaction throughput:** RocksDB exhibits higher throughput than all the alternatives considered: 3%-16% better than InnoDB, and far better than TokuDB. Not visible in the graph: in all cases, CPU is the bottleneck preventing throughput to further increase.
- **CPU overhead:** When stronger compression is used, then RocksDB exhibits less than 20% higher CPU overhead per transaction compared to InnoDB with no compression, but less than 30% as much CPU overhead as TokuDB. RocksDB with strong compression incurs only 80% as much CPU overhead as InnoDB with compression.
- **Write Volume:** The volume of data written per transaction in RocksDB is less than 20% of the volume of data written by InnoDB.[6] RocksDB write volume is significantly lower than TokuDB write volume.
- **Read Volume:** The volume of data read per read transaction in RocksDB is 20% higher than InnoDB when no compression is used, and between 10% and 22% higher when compression is used. RocksDB read volume is significantly less than TokuDB read volume.

Figure 2 depicts the quality of service achieved by the different storage engines. Specifically, it shows the 99-th percentile latencies for read and write requests on both vertices and edges in the LinkBench database. The behavior of RocksDB is an order of magnitude better than the behavior of all the other alternatives considered.

The results of the TPC-C benchmark are shown in Figure 3. The database size statistics are more difficult to interpret here because the TPC-C database grows with the number of transactions. As a result, for example, InnoDB with compression is shown to generate a small-sized database, but this is only the case because this InnoDB variant was able to process far fewer transactions up until the measurements were taken; in fact, InnoDB database size grows much faster in transaction time than RocksDB.

The figure clearly shows that RocksDB is not only competitive on OLTP workloads, but generally has higher transaction throughput while requiring significantly less storage space than the alternatives. RocksDB writes out less data per transaction than all the other configurations tested, yet reads only marginally more and requires only marginally more CPU overhead per transaction.

6. CONCLUDING REMARKS

We described how RocksDB was able to reduce storage space requirements by 50% over what InnoDB would need, while at the same time increasing transaction throughput and significantly decreasing write amplification yet increasing average read latencies by a marginal amount. It did this by leveraging LSM trees and applying a variety of techniques to conserve space.

A number of these techniques are being described for the first time, including: (*i*) dynamic LSM level size adjustment based on on current DB size; (*ii*) tiered compression where different levels of compression are used at different LSM levels; (*iii*) use of a shared compression dictionary; (*iv*) application of bloom filter to key prefixes; and (*v*) use of different size multipliers at different LSM levels. Moreover, we believe this is the first time an LSM-based storage engine has been shown to have competitive performance when running traditional OLTP workloads.

[6]The I/O volume numbers were obtained from *iostat*. The write volume numbers had to be adjusted because *iostat* counts TRIM as bytes written when in fact none are. RocksDB frequently deletes entire files (in contrast to InnoDB) and uses TRIM for that, which iostat reports as if the entire file had been written.

7. REFERENCES

[1] http://rocksdb.org.

[2] http://leveldb.org.

[3] G. J. Chen, J. L. Wiener, S. Iyer, A. Jaiswal, R. Lei, N. Simha, W. Wang, K. Wilfong, T. Williamson, and S. Yilmaz, "Realtime data processing at Facebook," in *Proc. Intl. Conf. on Mgmt. of Data*, 2016, pp. 1087–1098.

[4] A. Sharma, "Blog post: Dragon: A distributed graph query engine," https://code.facebook.com/posts/1737605303120405/dragon-a-distributed-graph-query-engine/, 2016.

[5] https://www.mongodb.com.

[6] https://yahooeng.tumblr.com/post/120730204806/sherpa-scales-new-heights.

[7] http://opencompute.org.

[8] M. Athanassoulis, M. S. Kester, L. M. Maas, R. Stoica, S. Idreos, A. Ailamaki, and M. Callaghan, "Designing access methods: The RUM conjecture," in *Proc. Intl. Conf. on Extending Database Technology (EDBT)*, 2016.

[9] P. O'Neil, E. Cheng, D. Gawlick, and E. O'Neil, "The log-structured merge-tree (LSM-tree)," *Acta Informatica*, vol. 33, no. 4, pp. 351–385, 1996.

[10] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber, "Bigtable: A distributed storage system for structured data," *ACM Trans. on Computer Systems (TOCS)*, vol. 26, no. 2, p. 4, 2008.

[11] A. Lakshman and P. Malik, "Cassandra: a decentralized structured storage system," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 2, pp. 35–40, 2010.

[12] A. S. Aiyer, M. Bautin, G. J. Chen, P. Damania, P. Khemani, K. Muthukkaruppan, K. Ranganathan, N. Spiegelberg, L. Tang, and M. Vaidya, "Storage infrastructure behind facebook messages: Using hbase at scale." *IEEE Data Eng. Bull.*, vol. 35, no. 2, pp. 4–13, 2012.

[13] T. Hobbs, "Blog post: When to use leveled compaction," http://www.datastax.com/dev/blog/when-to-use-leveled-compaction, june 2012.

[14] A. C.-C. Yao, "On random 2—3 trees," *Acta Inf.*, vol. 9, no. 2, pp. 159–170, Jun. 1978.

[15] T. G. Armstrong, V. Ponnekanti, D. Borthakur, and M. Callaghan, "Linkbench: A database benchmark based on the facebook social graph," in *Proc. 2013 ACM SIGMOD Intl. Conf. on Management of Data*, 2013, pp. 1185–1196.

[16] I. Tokutek, "TokuDB: MySQL performance, MariaDB performance," http://www.tokutek.com/products/tokudb-for-mysql/, 2013.