

Skeena: Efficient and Consistent Cross-Engine Transactions

Jianqiu Zhang
jianqiu@sfu.ca
Simon Fraser University

Tianzheng Wang
tzwang@sfu.ca
Simon Fraser University

Kaisong Huang
kha85@sfu.ca
Simon Fraser University

King Lv
lvjinqian@huawei.com
Huawei Cloud Database Innovation Lab

ABSTRACT

Database systems are becoming increasingly multi-engine. In particular, a main-memory database engine may coexist with a traditional storage-centric engine in a system to support various applications. It is desirable to allow applications to access data in both engines using *cross-engine* transactions. But existing systems are either only designed for single-engine accesses, or impose many restrictions by limiting cross-engine transactions to certain isolation levels and table operations. The result is inadequate cross-engine support in terms of correctness, performance and programmability.

This paper describes Skeena, a holistic approach to cross-engine transactions. We propose a lightweight snapshot tracking structure and an atomic commit protocol to efficiently ensure correctness and support various isolation levels. Evaluation results show that Skeena maintains high performance for single-engine transactions and enables cross-engine transactions which can improve throughput by up to 30× by judiciously placing tables in different engines.

CCS CONCEPTS

• **Information systems** → **Database transaction processing**; **Main memory engines**; **DBMS engine architectures**.

KEYWORDS

Cross-engine transactions; snapshot isolation; OLTP; MySQL

ACM Reference Format:

Jianqiu Zhang, Kaisong Huang, Tianzheng Wang, and King Lv. 2022. Skeena: Efficient and Consistent Cross-Engine Transactions. In *Proceedings of the 2022 International Conference on Management of Data (SIGMOD '22)*, June 12–17, 2022, Philadelphia, PA, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3514221.3526171>

1 INTRODUCTION

Traditional database engines are storage-centric: they assume data is storage-resident and optimize for storage accesses. Modern database servers often feature large DRAM that fits the working set or entire databases, enabling memory-optimized database engines [5, 21, 28, 35–37, 41, 43, 61, 63] that perform drastically better with lightweight concurrency control, indexing and durability designs.

Now suppose you are a database systems architect, and inspired by recent advances, built a new memory-optimized engine. But




Figure 1: Multi-engine database system. Data accesses are routed to the corresponding storage engines.

soon you found it was difficult to attract users: some do not need such fast speed; some say “I want it only for some tables or part of my application.” A common solution is to integrate the new engine into an existing system that initially uses a traditional engine. The result is a *multi-engine* database system (Figure 1). The application can judiciously use tables in both engines. Although engines share certain services (e.g., SQL parser), each engine is autonomous with its own indexes, concurrency control, etc. Some systems [20, 47, 55] already take this approach for easier migration and compatibility.

1.1 Cross-Engine: a Poorly-Supported Necessity

As an experienced architect—perhaps even before users did—you realized it was necessary to support cross-engine transactions. For example, a financial application may use a memory table for fast trading and keep other data in the traditional engine for low cost; yet the user may need to access both engines for recent and historical trading data in one ACID transaction [19]. The application may use a unified SQL interface to access all engines, but since each engine is implemented as an autonomous “package,” the system has to use each engine’s own transaction abstractions; we refer to them as *sub-transactions*. A *transaction* consists of at least one sub-transaction. In Figure 1, *S* is single-engine with *S₁*, while *T* is cross-engine with *T₁* (memory-optimized) and *T₂* (storage-centric).

Cross-engine transactions can be very useful, but existing support is inadequate in terms of correctness, performance and programmability. First, although simply starting and committing sub-transactions suffice to support single-engine transactions, doing so does not ensure correct cross-engine execution. A transaction over two engines that both use snapshot isolation (SI) [6] can still see inconsistent data and run under a lower-than-SI isolation level. Even if both engines ensure serializability, the overall execution is not necessarily serializable. Simply committing sub-transactions also risks atomicity if a sub-transaction fails to commit.

SIGMOD '22, June 12–17, 2022, Philadelphia, PA, USA

© 2022 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 2022 International Conference on Management of Data (SIGMOD '22)*, June 12–17, 2022, Philadelphia, PA, USA, <https://doi.org/10.1145/3514221.3526171>.

Second, prior designs are not aimed at achieving high performance in modern multi-engine systems, which are *fast-slow* where a (much faster) memory-optimized engine and a (much slower) storage-centric engine coexist in a single node. So it is vital for the cross-engine solution to impose low (if any) overhead, especially on the faster engine. Prior solutions [8, 11, 24, 31, 57, 59, 60] ignored this hidden requirement by assuming a cluster of similar systems.

Finally, past solutions are often at odds with (1) keeping engine autonomy for maintainability as engines are typically developed by different teams (but still of the same vendor), and (2) easing application development. They often require non-trivial application changes and limit functionality, by forcing users to pre-declare whether a transaction is cross-engine or to use certain isolation levels [20]; both can be complex and affect performance.

1.2 Skeena

We present Skeena, a holistic approach to efficient and consistent cross-engine transactions in the context of multi-versioned, fast-slow systems. We make three key observations to guide Skeena’s design. First, as noted by prior work [8], inconsistent snapshots can be avoided by carefully selecting a snapshot in each engine. This requires efficiently tracking snapshots that can be safely used by later transactions. Second, in addition to using correct snapshots and enforcing sub-transactions commit in the same order across engines, for serializability it suffices to require each engine use commit ordering, i.e., forbid schedules where commit and dependency orders mismatch [1, 57]. Many concurrency control protocols exhibit this property, including the widely-used 2PL and optimistic concurrency control (OCC) [21, 38, 43, 51, 63]. Finally, engines are developed and/or well understood by the same vendor, potentially allowing non-intrusive changes to engines for more optimizations.

Based on these observations, we design Skeena to consist of (1) a cross-engine snapshot registry (CSR) for correct and efficient snapshot selection and (2) an extended pipelined commit protocol for atomicity and durability without (expensive) traditional 2PC. Skeena can be easily plugged into an existing system.

Conceptually, CSR maintains mappings between commit timestamps (therefore snapshots) in one engine and those in another. A transaction may start by accessing any engine using the latest snapshot s . Upon accessing another engine E , it queries CSR using s to select a snapshot in E using which would avoid incorrect executions. Further, with CSR one only needs to set each engine to use a serializable protocol that exhibits commit ordering to guarantee serializability. Later, we discuss the detailed algorithms to realize this idea and techniques that make CSR lightweight and easy to maintain. In fast-slow systems, CSR incurs negligible overhead as the storage accesses in the traditional engine present a bigger bottleneck, and single-engine transactions do not access CSR at all.

Leveraging the fact that engines can communicate via fast shared memory (e.g., in the same address space), Skeena extends the widely-used group/pipelined commit protocols [34, 64, 68] to ensure atomicity and durability. Upon commit, the worker thread detaches the transaction and places it on a commit queue, before continuing to work on the next request. Meanwhile, a background committer thread monitors the queue and durable log sequence numbers in both engines to dequeue transactions whose sub-transactions’

log records have been fully persisted. This way, Skeena ensures cross-engine transactions are not committed (i.e., with results made visible to the application) until all of its sub-transactions are committed, while avoiding expensive 2PC.

We adopted Skeena in MySQL to enable cross-engine transactions across its default InnoDB and ERMLA [36], an open-source main-memory OLTP engine. This required 83 LoC out of over 200k LoC of the entire codebase. Evaluation on a 40-core server shows that Skeena retains the memory-optimized engine’s high performance, and incurs very low additional overhead for cross-engine transactions. By judiciously placing tables in both engines, Skeena can help improve the throughput of realistic workloads by up to 30× compared to using traditional engines.

Note that our goal is *not* to build faster database engines, nor to invent new concurrency control protocols for cross-engine transactions; both are well studied by prior work. Instead, we aim to (1) enable cross-engine transactions without excessive overhead and (2) explore practical designs for modern fast-slow systems.

1.3 Contributions

This paper makes five contributions. ❶ We analyze the correctness requirements of cross-engine transactions under various isolation levels, ranging from read committed to serializable. ❷ We distill a set of desirable properties and design principles to be followed by multi-engine systems. ❸ We propose Skeena, a holistic approach to consistent cross-engine transactions by leveraging the properties of the fast-slow multi-engine architecture. ❹ We show Skeena’s feasibility and explore practical design issues by integrating an open-source memory-optimized engine (ERMLA) into MySQL alongside its storage-centric engine (InnoDB). ❺ Through comprehensive experiments, we explore the potential and distill useful recommendations of using cross-engine transactions to improve performance and reduce storage costs under realistic workloads. Skeena is open-sourced at <https://github.com/sfu-dis/skeena>.

2 BACKGROUND

In this section, we give the necessary background for cross-engine transactions and motivate our work.

2.1 Modern Fast-Slow Multi-Engine Systems

We have described the idea of multi-engine systems in Section 1. Several production systems already adopted the fast-slow architecture: SQL Server supports memory-optimized tables managed by its Hekaton main-memory engine [21, 46]; PostgreSQL supports additional engines through foreign data wrapper [55], which is used by Huawei GaussDB to integrate a main-memory engine [5].

Multi-engine systems bear similarities to distributed and federated database systems [8, 9, 11, 18, 24, 31, 42, 56, 60], but are unique in several ways. As Table 1 summarizes, a multi-engine system integrates engines developed and/or understood by the same vendor; in contrast, federated systems consist of opaque systems developed by different vendors. Distributed systems typically involve a set of nodes that run the same engine carefully designed to support distributed transactions, exhibiting low autonomy. Fast-slow systems integrate different engines that vary in performance, so an

Table 1: Multi-engine vs. distributed and federated systems.

	Multi-Engine	Federated	Distributed
Engine Internals	Transparent	<i>Opaque</i>	Transparent
Engine Types	Heterogeneous	Heterogeneous	<i>Homogeneous</i>
Autonomy	<i>Almost full</i>	Full	Low
Scalability	<i>Up and/or out</i>	Out	Out


inefficient cross-engine solution may penalize single-engine transactions, defeating the purpose of adopting a fast engine; mitigating such overhead is the major goal of our work. Note that multi-engine systems often allow slightly trading autonomy for performance and compatibility, e.g., by managing schemas in all engines centrally. However, federated systems allow little room for doing so, as each system is usually a proprietary package. Multi-engine systems can scale up and out, whereas the other two types of systems mainly focus on scaling out. We focus on single-node fast-slow systems and leave scaling out as future work. Finally, both multi-engine and federated systems may present applications with a unified interface (e.g., SQL). But this does not automatically guarantee correctness for cross-engine transactions. Unlike federated systems which already address this issue [11, 31, 60], modern fast-slow systems either completely lack the support for cross-engine transactions (e.g., MySQL) or come with many restrictions (e.g., SQL Server); we elaborate in Section 2.4 after introducing more necessary background next.

2.2 Database Model and Assumptions

Now we lay out the preliminaries for analyzing cross-engine transactions in fast-slow systems.

Multi-Versioning. Many popular systems are multi-versioned, including storage-centric (e.g., MySQL InnoDB, PostgreSQL and SQL Server) and memory-optimized (e.g., Hekaton [21], ERMA [36] and Cicada [43]) engines. Given the wide adoption, we focus on multi-versioned systems. Following prior work [1, 2, 8, 13, 65], we model databases as collections of records, each of which is a totally-ordered sequence of *versions*. Updating a record appends a new version to the record’s sequence. Inserts and deletes are special cases of updates that append a valid and special “invalid” version, respectively. Obsolete versions (as a result of deletes/updates) are physically removed only after no transaction will need them, using reference counting or epoch-based memory management [10, 36].

Reading a record requires locating a proper version; this is dictated by the concurrency control protocol. We base on a common design [21, 36, 41, 67] where the engine maintains a global, monotonically increasing counter that can be atomically read and incremented. Note that in a multi-engine system, engines maintain their own timestamp counters; for now, we assume single-engine transactions and expand to cross-engine cases later. Each transaction is associated with a begin timestamp and a commit timestamp, both drawn from the counter. Upon commit, the transaction obtains its commit timestamp, which determines its commit order by atomically incrementing the counter. Each version is associated with the commit timestamp of the transaction that created it. Transactions

**Figure 2:** Inconsistent snapshots. (a) S uses an older (newer) snapshot in E_1 (E_2). (b) U sees T_1 ’s results, but does not see T_2 ’s.

access data using a *snapshot* (aka *read view*), which is a timestamp that represents the database’s state at some point in (logical) time.

Isolation Levels. For read committed, we always read the latest committed version. SI allows the transaction to read the latest version created before its begin timestamp obtained upon transaction start or the first data access. A transaction can update a record if it can see the latest committed version. Serializability can be achieved by locking [15, 16, 22] and certifiers [13, 38, 65]. We aim to enforce these isolation levels in the presence of cross-engine transactions.

Cross-Engine ACID Properties. Analogous to maintaining ACID properties in a single engine, a multi-engine system must maintain these for both single- and cross-engine transactions:

- **Atomicity:** All the sub-transactions should eventually reach the same commit or abort conclusion, i.e., either all or none of the sub-transactions commit in their corresponding engines.
- **Consistency:** All transactions (single- or cross-engine) should transform the database from one consistent state to another, enforcing constraints within and across engines.
- **Isolation:** Changes in any engine made by a cross-engine transaction must not be visible until the cross-engine transaction commits, i.e., all sub-transactions have committed.
- **Durability:** Changes made by cross-engine transactions should be persisted while guaranteeing atomicity.

Enforcing cross-engine ACID requires careful coordination of sub-transactions to avoid anomalies, as we describe next.

2.3 Cross-Engine Anomalies

The relative ordering of sub-transaction begin/commit events directly determines correctness, as certain ordering may lead to anomalies and violate ACID requirements, as described next.

Issue 1: Inconsistent Snapshots. There are two cases where a transaction may be given an inconsistent snapshot. In Figure 2(a), S started in E_1 with snapshot 1000, and T started in E_2 with snapshot 100. Suppose another transaction in E_1 committed by incrementing E_1 ’s timestamp counter to 3000. Then, T accesses E_1 , which assigns T_1 its latest snapshot 3000, and S_2 obtains snapshot 200 in E_2 . Compared to S , T sees a newer version of the database in E_1 , but an older version in E_2 . This would require S and T to start before each other, which is impossible in a correct SI schedule [1]. This corresponds to the “cross” phenomenon in distributed SI (DSI) [8].

Another anomaly (isolation failure) may allow transactions to see partial results. In Figure 2(b), T first commits T_1 with timestamp 4000. Until T_2 is fully committed, T is still in-progress, so none of its changes should be visible to other transactions. Meanwhile, U starts

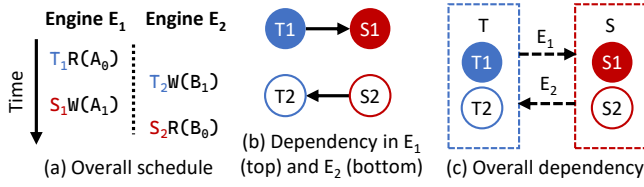


Figure 3: Non-serializable execution of cross-engine transactions. (a) Each engine executes a serializable schedule (b) without cyclic dependencies. (c) Overall cyclic dependency between T and S .

in E_2 with timestamp 250, and continues to open U_1 : since U_1 started after T_1 committed, by definition it should see T_1 's changes. Thus, U sees partial results: T 's results are visible in U 's snapshot in E_1 , but not E_2 . This anomaly corresponds to the serial-concurrent phenomenon in DSI [8]. Compared to inconsistent (skewed) snapshots which concern the order in which sub-transactions are opened, isolation failure arises when sub-transaction begin and commit actions are interleaved and inflict different write-read dependency orders in different engines.

Issue 2: Serializability. Even if both engines guarantee full serializability, the overall execution may not be serializable. As Figure 3(a) shows, S and T are concurrently executing in two engines that offer serializability. Each engine runs a serializable schedule, with an anti-dependency shown in Figure 3(b). However, as shown in Figure 3(c), the overall execution exhibits write skew with cyclic dependencies ($T \rightarrow S \rightarrow T$), indicating non-serializable execution.

Issue 3: Atomicity and Durability. A cross-engine transaction should commit either all or none of its sub-transactions. Distributed systems usually solve this problem with 2PC, but newer engines may not support it [7, 71]. 2PC's coordination overhead can also be heavyweight for shared memory, slowing down the (faster) main-memory engine. As we describe later, additional checks are needed in addition to a traditional 2PC prepare-commit protocol. Thus, 2PC may not be the best choice for single-node multi-engine systems.

2.4 State-of-the-Art and Motivation

Prior work can avoid the anomalies [8, 9, 11, 18, 24, 31, 42, 56, 57, 60], but they targeted distributed and federated systems without considering single-node fast-slow systems. For example, certain solutions for DSI [8] guarantee correct snapshots with a central coordinator node and global IDs, requiring non-trivial engine changes. 2PC as we have described may also not suit fast-slow systems.

The mismatch between prior approaches and fast-slow systems has led to missing or limited cross-engine support in real systems, motivating our work. For example, MySQL supports various engines under a unified interface. Users are free to issue multi-engine transactions, but correctness is undefined as MySQL does not coordinate snapshot and commit ordering *across* engines. All the anomalies in Section 2.3 could occur. Compared to MySQL, SQL Server supports cross-engine transactions with various restrictions [20]. For example, if both the traditional engine and Hekaton use SI, cross-engine transactions are not allowed at all, yet SI is among the most popular isolation levels in Hekaton [20]. These significantly limit the potential of cross-engine transactions.

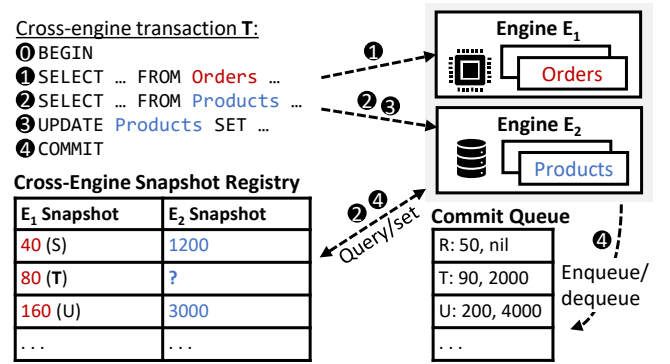


Figure 4: Skeena overview. ①—① Transactions access data without explicitly declaring whether they are cross-engine. ② Upon accessing an additional engine, the transaction ③ consults CSR to obtain a proper snapshot. ④ Cross-engine transactions use CSR for commit check and if passed, goes through the pipelined commit protocol.

3 DESIGN PRINCIPLES

We distill a set of desired properties and design principles that a cross-engine mechanism like Skeena should follow:

- **Low Overhead.** The mechanism should introduce as low overhead as possible. It should try not to penalize single-engine transactions, especially those in the faster engine.
- **Engine Autonomy.** Engines should be kept as-is, or only be minimally modified to work with the cross-engine mechanism or optimize for performance.
- **Full Functionality.** The mechanism should support various isolation levels for both single- and cross-engine transactions, unless it is limited by individual engine capabilities.
- **Transparent Adoption.** The application should not be required to make logic changes. Rather, it should only need to declare the “home” engine of each table in the schema.

4 SKEENA DESIGN

Skeena targets fast-slow systems with a main-memory and a storage-centric engine. We first give an overview of Skeena, and then discuss its design in detail.

4.1 Overview

Skeena ensures correct snapshot selection and atomic commit. As Figure 4 shows, Skeena consists of (1) the cross-engine snapshot registry (CSR) that tracks valid snapshots and (2) a pipelined commit protocol for atomically committing cross-engine transactions. Now we describe the high-level transaction workflow under Skeena.

Initialization. Transactions (single- or cross-engine) can keep using the database system's unified APIs (e.g., SQL), without providing additional hints (e.g., whether a transaction will be cross-engine). Figure 4 shows an example SQL program, which is written in the same way as without Skeena. Skeena does not force transactions to run under specific isolation levels. However, the system may allow users to specify an isolation level (e.g., SET TRANSACTION ISOLATION LEVEL in MySQL [48]). As part of the integration effort, Skeena can detect and enforce such settings across all engines.

Algorithm 1 Snapshot selection for cross-engine transactions.

```
1 def select_snapshot(e1_snap, engine &e2):  
  # Find existing snapshots that could be used  
2 candidates[] = CSR.forward_scan_1st(e1_snap)  
  if candidates is empty:  
3    # No existing mapping, obtain the latest from e2  
    e2_snap = e2.timestamp_counter  
4  else:  
    # Use the latest snapshot mapped to s <= e1_snap  
5    e2_snap = max(candidates)  
6    CSR.map(e1_snap, e2_snap)  
7  return e2_snap
```

Data Accesses. The system routes requests to the target engine which uses a sub-transaction to access data. Skeena requires no change to the existing routing mechanism. Upon start or accessing the first record, the sub-transaction obtains a snapshot. Depending on whether the transaction is single- or cross-engine, the system may directly give the latest snapshot in the underlying engine, or use CSR to obtain a snapshot that would not cause anomalies (steps ②–④ Query/set). If such a snapshot does not exist, the transaction will be aborted; we quantify the impact in Section 6.

Finalization. To commit, a cross-engine transaction consults CSR to verify that committing it would not lead to inconsistent snapshots for future transactions; single-engine transactions commit without using Skeena. Transactions that have passed CSR verification are marked as pre-committed to go through the pipelined commit protocol (step ④ Dequeue). If the verification fails, we abort the transaction by rolling back all of its sub-transactions.

In the rest of this section, we describe how Skeena facilitates snapshot selection, atomicity and durability, beginning with CSR.

4.2 Cross-Engine Snapshot Registry

The key to avoiding inconsistent snapshots is to ensure the sub-transactions of different cross-engine transactions follow the same start order in each engine [8]. That is, if T 's sub-transaction T_1 uses an older snapshot than S_1 does in engine E_1 , then T_2 should also use an older snapshot compared to that of S_2 in E_2 . For example, in Figure 4, T first started as a single-engine transaction accessing Orders in E_1 , using snapshot 80. When T starts to access Products in E_2 , T needs to use a snapshot (s) in E_2 such that s is between the snapshots of its “neighbors” in E_1 , i.e., S_1 and T_1 . Thus, T may use any valid E_2 snapshot between 1200 and 3000 (inclusive), although using 3000 would allow it to see fresher data.

To facilitate such a snapshot selection process, CSR tracks valid snapshots (i.e., commit timestamps of past cross-engine transactions) that can be safely used by future cross-engine transactions. Conceptually, CSR is a table of many-to-many mappings, where each “row” (CSR entry) is a pair of snapshots (i.e., commit timestamps), one from each engine as depicted by Figure 4. When a transaction crosses to an additional engine, it uses the current engine’s snapshot as the key to query CSR for a snapshot in the other engine. As Algorithm 1 shows, to access a new engine $e2$, the worker thread issues a non-inclusive forward scan over CSR using the snapshot in the current engine $e1$ as the key ($e1_snap$) to obtain a set of

Algorithm 2 Commit check for cross-engine transactions.

```
1 def cross_engine_commit_check(sub_t1&, sub_t2&):  
  # Obtain lower and upper bounds for sub-transaction t2  
2 low = -inf  
  candidates[] = CSR.reverse_scan_1st(sub_t1.commit_ts)  
3  if candidates[] is not empty:  
    low = max(candidates)  
4  
  high = +inf  
5  candidates[] = CSR.forward_scan_1st(sub_t1.commit_ts)  
  if candidates[] is not empty:  
6    high = min(candidates)  
7  
8  # Check if committing t2 would cause future anomalies  
9  if low > sub_t2.commit_ts or high < sub_t2.commit_ts:  
10    return false  
11  else:  
12    # Check passed, setup mapping and return  
13    CSR.map(sub_t1.commit_ts, sub_t2.commit_ts)  
14    return true
```

candidate snapshots. The scan returns once a first key greater than $e1_snap$ is met or no such key is found. If the scan returns an empty set, then no past transaction has set up any mapping, or the current transaction is using the latest $e1$ snapshot. Then we use the latest $e2$ snapshot (lines 4–6). However, if any candidate is found, we must take an $e2$ snapshot that is already mapped to $e1_snap$ to avoid anomalies (lines 7–9). Finally, we ensure the mapping is recorded at line 10. Under SI, the algorithm is executed only once per transaction when it becomes cross-engine. Subsequent accesses continue to use the previously acquired snapshots.

In addition to acquiring snapshots, committing a cross-engine transaction implicitly limits the ranges of snapshots a (future) cross-engine transaction may use: recall that the commit timestamp of a previous transaction T in fact is the snapshot of a future transaction that reads the results generated by T . Thus, CSR also tracks commit timestamps of cross-engine transactions. Algorithm 2 describes the process at a high level. Here, we assume the sub-transaction commit timestamps are given (as the `commit_ts` member in each sub-transaction); we revisit this assumption later in more detail. The idea is to ensure that committing a cross-engine transaction—i.e., adding a new mapping entry to CSR—would not add skewed snapshots to CSR. Thus, upon commit, we issue a reverse scan and a forward scan over CSR using a sub-transaction’s (sub_t1) commit timestamp to obtain the lower and higher bounds for the other commit timestamp (lines 4–11 in Algorithm 2). If sub_t2 ’s commit timestamp falls between the higher and lower bounds, we can safely commit this cross-engine transaction and setup a new mapping in CSR (line 18). Otherwise, the transaction must be aborted. Note that the mapping step in Algorithm 1 is still necessary: (1) single-engine commits are not covered by CSR to avoid unnecessary overheads, and (2) a cross-engine transaction may access data generated by single-engine transactions and form new cross-engine snapshots.

Since a transaction may access engines in any order (e.g., from the storage-centric engine and crosses over to the memory-optimized

engine, and vice versa), CSR needs to support queries from *either* engine. CSR may be implemented using a relational table in one of the supported engines with full-table scan or two range indexes, each of which is built on a “column” of the CSR table. However, this can create dependency on a particular engine and incur much table and index maintenance overhead. A practical design must also support concurrency. We address these issues next.

4.3 Lightweight Multi-Index CSR

We take advantage of the unique properties of fast-slow systems to devise a lightweight CSR that mitigates the above issues.

Anchor Engine. Compared to storage-centric engines, it is typically much cheaper to obtain snapshots in main-memory engines. This is often as simple as manipulating an 8-byte counter in a lock-free manner without using a mutex [69]. For example, ERMIA [36] only needs to read the counter; Hekaton [21] increments the counter using atomic fetch-and-add (FAA) [32] to keep the process efficient.¹ On the contrary, obtaining a snapshot in a storage-centric system can be much more complex. For example, MySQL InnoDB needs to take multiple mutexes to compute watermark values (see Section 5).

Leveraging the existence of a fast and a slow engine, Skeena designates an *anchor* engine and always follows the snapshot order in the anchor engine. The anchor engine should be the one where it is cheaper to acquire a snapshot (usually the memory-optimized one). Then a transaction always starts by acquiring the latest snapshot from the anchor engine, and uses it to query CSR when it extends to the other engine. This allows us to maintain one-to-many mappings (instead of many-to-many mappings), which simplifies CSR to become a range index that uses the anchor engine’s snapshots as “keys” and lists of snapshots in the other engine as “values.” We currently use Masstree [44], a high-performance in-memory index, but any concurrent data structure that supports range queries would suffice. A side effect is transactions that only access the “slower” engine become cross-engine and need to use Algorithms 1–2. As Section 6 shows, the overhead is negligible compared to data accesses which may involve the storage stack while CSR is fully in-memory.

Using the main-memory engine as the anchor is an optimization, not a requirement: in theory any engine can be the anchor. In case a heavyweight engine has to be the anchor, cross-engine transactions may incur higher overhead for creating snapshots (thus lower overall performance). This will in turn reduce the pressure on CSR which is less frequently accessed and maintains fewer snapshots.

Multi-Index. Since CSR tracks cross-engine snapshots and commit histories, its size can grow quickly, slowing down query speed over time; entries that are no longer needed should also be cleaned up. We solve these problems by partitioning the CSR by snapshot ranges, reminiscent of multi-rooted B-trees [53]. The result is a multi-index design shown in Figure 5. Each partition is an index and covers a unique range of snapshots so that a transaction only uses a single index. In Figure 5, the first two indexes cover mappings in the ranges of [30, 400] and [401, 500], respectively. Each partition has a fixed capacity (number of keys), and a new index is created when the current open index is full. Therefore, there is always one and only one open index that can accept new mappings;




Figure 5: Multi-index CSR. Each index covers a range of anchor snapshots, and is recycled in its entirety when no longer needed.

other indexes are read-only but can continue to serve existing transactions for snapshot selection. However, since inactive indexes are read-only, if a transaction needs to setup a new mapping in an inactive index during snapshot selection or commit check, it must be aborted; in practice, such aborts are rare as we evaluate in Section 6.

Snapshot Acquisition and Commit Check. With multiple indexes and an anchor engine, a transaction acquires snapshots by (1) obtaining a snapshot S from the anchor engine, (2) locating the index I that covers S , and (3) using S to query I and if needed, create a new entry in I following Algorithm 1 with $e1_snap = S$ and CSR at line 3 being I . Note that steps 2 and 3 are only executed if the transaction accesses the non-anchor engine. For example, if the main-memory engine is the anchor and the transaction only accesses an in-memory table, steps 2 and 3 are never executed. For step 2, we track all the indexes in a list/array (e.g., C++ `std::vector`). Each entry records the minimum snapshot of the partition and a pointer to the index. Since we keep only one open index, entries in the list are sorted by snapshot ranges. We search for I by traversing the list backwards and stopping at the first entry whose smallest snapshot is smaller than or equal to the given snapshot. In step 3, a new entry is inserted if and only if I is open; otherwise the transaction is aborted. Commit check follows a similar logic and can proceed only if I is open. Likewise, CSR in lines 4–9 of Algorithm 2 refers to I . Multiple threads may execute the above three steps concurrently, for which we describe our solution next.

4.4 CSR Concurrency and Maintenance

Now we discuss how Skeena handles concurrent accesses and manages/recycles indexes in multi-index CSR.

Concurrency. Although latches can be a potential bottleneck in multicore systems [33], a latch-based solution in Skeena can be efficient by leveraging the fast-slow property: compared to executing transactions in the slower engine, using latches and high-performance indexes present negligible overhead and little impact on overall performance for cross-engine transactions. Each index is protected by a mutex, and we protect the array of all indexes using a reader-writer lock for mutual exclusion between threads that only query an index without modifying the list using the reader mode and those that may add or remove an index using the writer mode. A transaction starts by latching the list of indexes in shared (reader) mode to locate the target index I . Then the thread latches I for exclusive access to run Algorithm 1. If the thread needs to create a new index, e.g., if I is full or the list is empty, it (1) releases the list latch, (2) requires the latch in writer mode to allow inserting to the list, and (3) checks if such an index has been inserted by another thread between steps 1 and 2, and if so, retries the entire process after releasing the list latch; otherwise we proceed by (4)

¹Such designs are common in multi-versioned engines [26, 43, 67, 69]. When the engine is integrated into a full system, e.g., MySQL, FAA’s overhead is negligible.

appending the new index with a new mapping to the list. Commit check follows the same logic so we do not repeat the details.

Index Maintenance. Using multiple indexes simplifies garbage collection (GC) as we can delete an entire index once its mappings are no longer needed, instead of issuing many key delete operations to an index. To recycle, we first iterate over all the active transactions to find the oldest anchor-engine snapshot (`min_snap`). Then we exclusively latch the list of indexes and scan through it to remove stale indexes that cover ranges below `min_snap`. In case long-running transactions prevent `min_snap` from growing, one may further find opportunities to remove unused indexes covering newer but still unused ones, reminiscent of GCing long version chains in multi-versioned systems [10]. For example, assume a transaction still uses the left-most index in Figure 5, yet no transactions need the middle index (401–500), which can then be first recycled; the oldest index is recycled later when it is no longer needed. Our experiments do not indicate this approach to be necessary and recycling is fast as it is purely in-memory. Recycling is triggered between CSR accesses based on a user-defined threshold (e.g., once per 5000 accesses); it could also be delegated to background threads.

4.5 Commit Protocol

Once all accesses are finished, Skeena checks whether both sub-transactions can commit using engine-level commit timestamps that represent the sub-transactions’ commit ordering. Thus, Skeena needs to obtain the sub-transaction’s commit timestamp from each engine. This is usually easy for memory-optimized engines which break the commit process into pre- and post-commit [21, 36, 39]. During pre-commit, the engine assigns a commit timestamp and uses it to determine whether the transaction can commit without violating correctness criteria (e.g., serializability). If so, post-commit will finish the commit by marking new records as visible, finalizing log records, etc. Otherwise the transaction is aborted. Some (mainly storage-centric) engines may not explicitly expose such pre- and post-commit interfaces, but engines in a multi-engine system are maintained by the same vendor. This justifies simple changes in engines to expose the pre- and post-commit interfaces, which is straightforward in practice by breaking a monolithic “commit” function into a pre- and post-commit function (Section 5). Single-engine transactions directly execute the two steps without commit check.

With the pre- and post-commit interfaces, Skeena commits a cross-engine transaction in three steps. (1) Pre-commit both sub-transactions to obtain commit timestamps. (2) Use the timestamp obtained from the anchor engine to conduct the commit check. (3) If the check passes, post-commit both sub-transactions. From a high-level, Skeena’s commit protocol resembles 2PC: step 1 may correspond to 2PC’s prepare phase that collects commit decisions from each engine; step 3 may correspond to 2PC’s commit phase. However, Skeena differs from 2PC by requiring an additional check (step 2) even after all the engines have pre-committed the transaction. So an “all-yes” result from the 2PC-equivalent prepare phase does not necessarily mean a cross-engine transaction can commit.

Before both sub-transactions are post-committed, changes by either should be kept invisible, yet from the perspective of an engine, a post-committed (sub-)transaction is fully committed with its results visible. Skeena must ensure partial results are not visible until

all sub-transactions are post-committed. We observe that a simple yet effective solution is to extend the pipelined commit protocol [34] which was initially proposed to hide log flush latency. It decouples transactions waiting for log flushes and worker threads to keep I/O off the critical path. Upon commit, instead of directly issuing a log flush, the thread detaches the transaction and appends it to a global commit queue (or a partitioned queue to avoid introducing a central bottleneck). Results by these transactions are immediately visible internally but are not returned to applications until their log records have been persisted. A daemon tracks transactions awaiting log durability on the commit queue, and dequeues transactions whose log records have been persisted. Some systems use this approach to improve throughput without sacrificing correctness [33, 34, 64, 66].

Based on this idea, Skeena (1) pushes both sub-transactions onto the commit queue upon post-commit and (2) has the commit daemon monitor both engines’ log flushes to dequeue transactions. If an engine already implements commit pipelining, Skeena can directly extend it. Note that single-engine and read-only transactions must also use commit pipelining [30] as they may read cross-engine transactions’ results; we quantify its impact on latency in Section 6.

4.6 Durability and Recovery

In a multi-engine system, each engine implements its own approach to durability and crash recovery. Sub-transactions still follow their corresponding engines’ approach to persist data and log records. Checkpoints can be taken as usual independently by each engine. To ensure atomicity of cross-engine transactions, Skeena can record the pre- and post-commit of cross-engine transactions, by maintaining a standalone log or piggybacking on individual engines. The latter can be easier to implement: upon pre-commit we append a `commit-begin` record, and after post-commit finishes, the engine appends a `commit-end` record. During recovery, each engine executes its recovery mechanism and rolls back changes done by cross-engine transactions whose sub-transactions are not fully committed. Alternatively, the recovery procedure may inspect each engine’s log and truncate at the first “hole” where only one sub-transaction of a cross-engine transaction is committed. This is safe because transactions that depend on partially committed cross-engine transactions will wait on the commit queue and their results were never made visible to applications.

4.7 Serializability

As noted by prior work [57], disallowing anti-dependencies (i.e., using commit order as dependency order) in all engines is sufficient for cross-engine serializability. This translates into choosing a concurrency control protocol for each engine where a sub-transaction can only commit if its read records are not concurrently modified by a newer transaction. A wide range of engines [5, 21, 26, 37, 39, 41, 43, 63] exhibit this property based on 2PL (by blocking readers and writers) and OCC (by verification at commit time). Some protocols [13, 14, 27, 65] can tolerate certain safe anti-dependencies, but would require implementing verification in Skeena. This needs engines to expose dependency information, tightly coupling Skeena with engine design and sacrificing engine autonomy. Thus, we take the former approach that imposes no engine-level changes.

4.8 Correctness

Skeena’s theoretical foundation comes from DSI [8] and commit ordering (CO) [57] which respectively ensure consistent snapshots and serializability. Different from DSI and CO, Skeena targets single-node instead of distributed systems. In essence, Skeena implements DSI and CO for shared memory. Both DSI and CO enforce sufficient conditions with their correctness formally proved. Thus, we argue for Skeena’s correctness by showing Skeena enforces the same conditions as DSI and CO. We first lay out the necessary notations used by DSI [8] and our adaptation for shared memory:

- c_x : Commit of transaction x ;
- c_x^i : Commit of transaction x on node/engine i ;
- b_x : Begin of transaction x ;
- b_x^i : Begin of transaction x on node/engine i ;
- $SN(x, y)$: All nodes/engines accessed by x and y ;
- $x < y$: Transaction x is serialized before y ;
- $op_1 < op_2$: Begin/commit operation op_1 is ordered before op_2 .
For begin operations, $op_1 \leq op_2$ is allowed as begin timestamps can be acquired by reading a counter.

Then DSI defines the following sufficient conditions for correctness:

THEOREM 4.1. *Suppose each node i enforces correct local SI. If all the local schedules satisfy the following rules, we can construct a correct DSI schedule, i.e., the execution is correct DSI [8]:*

- $\exists c_x^i < c_y^j \rightarrow c_x < c_y$ (Rule 1)
- $\exists b_x^i < c_y^j \rightarrow b_x < c_y$ (Rule 2)
- $\exists c_x^i < b_y^j \rightarrow c_x < b_y$ (Rule 3)
- $\exists b_x^i \leq b_y^j \rightarrow b_x \leq b_y$ (Rule 4)
- $c_x < c_y \rightarrow \forall j \in SN(x, y) : c_x^j < c_y^j$ (Rule 5)
- $b_x < c_y \rightarrow \forall j \in SN(x, y) : b_x^j < c_y^j$ (Rule 6)
- $c_x < b_y \rightarrow \forall j \in SN(x, y) : c_x^j < b_y^j$ (Rule 7)
- $b_x < b_y \rightarrow \forall j \in SN(x, y) : b_x^j < b_y^j$ (Rule 8)

Analogous to DSI, in a shared-memory environment, Rules 1–4 construct the partial order of begin/commit events of *cross-engine* transactions (“distributed transactions” in DSI). Then, Rules 5–8 enforce the same partial order across all the *engines* (“nodes” in DSI). We omit the detailed proof (available elsewhere [8]); our goal is to show the conditions enforced by Skeena satisfy Rules 1–8.

THEOREM 4.2. *Skeena enforces the same partial order of cross-engine transactions on all engines as defined by Theorem 4.1 in a shared-memory environment.*

PROOF. In Algorithm 1, given some b_x^i , Skeena either uses the latest timestamp (line 6) or the latest snapshot from engine j whose corresponding snapshot in i is no newer than b_x^i . Therefore, Algorithm 1 enforces Rules 4 and 8 of Theorem 4.1. Similarly, Algorithm 2 forbids $< c_x^i, c_y^j >$ to be inserted into CSR if c_y^j does not maintain the same partial order on engine i (lines 14–15), satisfying Rules 1 and 5. Note that begin timestamps are in fact previously committed transactions’ commit timestamps, enforcing Rules 3, 4, 6 and 7. \square

For serializability, Skeena requires the use of concurrency control protocols to follow the requirements of CO [57], without introducing any additional algorithms. Hence, Skeena can support serializability using commit ordering.

Table 2: Requirements to achieve different isolation levels.

Isolation Level	Requirements
Read Committed	<i>Engines:</i> Isolation level \geq Read Committed <i>Skeena:</i> Refresh snapshot per record access
Snapshot Isolation	<i>Engines:</i> Isolation level \geq Snapshot Isolation <i>Skeena:</i> New snapshot upon start/first access
Serializable	<i>Engines:</i> Serializable, no anti-dependencies <i>Skeena:</i> New snapshot upon start/first access

4.9 Discussions

In essence, Skeena is a coordinator that enforces correct snapshots and atomic commit in fast-slow systems. Table 2 lists the requirements to achieve different overall isolation levels. To achieve an overall isolation level of read committed (RC), in most systems this means to acquire a new snapshot per record access. Guaranteeing SI or serializable isolation levels usually requires obtaining a snapshot upon transaction start or the first record access. Since Skeena does not implement extra concurrency control logic to avoid tight coupling with engines, for all isolation levels (e.g., SI), each engine needs to run at least at it (e.g., SI) or higher to ensure sub-transactions are correctly scheduled. Thus, the overall isolation level guaranteed by Skeena is at most the lower level being used across all engines. For example, if two engines respectively use RC and SI, then Skeena can only guarantee RC overall.

Our focus has been on dual-engine systems with interpreted queries. To support more engines, a straightforward way is to extend the multi-index CSR to become a hierarchy of indexes, each is the anchor of the next lower-level engine. This would allow Skeena to enforce ordering between sub-transactions. The downside is more complex CSR maintenance which may require more efficient CSR designs. If more changes are tolerable, one may introduce additional global ordering to simplify sub-transaction ordering. However, this can potentially couple the design with engine internals. Most memory-optimized engines compile queries to machine code [21, 45, 49]. Skeena is orthogonal to whether queries are compiled or interpreted, although we focus on the latter as accesses in conventional engines can cancel out compilation’s benefits.

Finally, Skeena can be applied to systems that (1) support multiple engines and (2) follow the database model in Section 2.2. Both are widely available in practice. For example, MySQL and SQL Server already support multiple engines. PostgreSQL can support additional engines using foreign data wrapper. Many systems, including MySQL, PostgreSQL and SQL Server, employ multi-versioning. We discuss in depth how Skeena can be used by MySQL later; for space limitation we do not expand on other systems. Moreover, Skeena does not require significant engine-level changes. The most notable (yet simple) change (mainly for conventional engines) is exposing commit ordering via a pre-commit interface. Skeena only expects the commit/abort decision of sub-transactions, without dictating engine internals, such as whether cascading abort is possible or how writes and versions are organized. The remaining effort is mainly put into integrating Skeena with existing multi-engine support. As Section 5 describes, these changes are not intrusive or complex.

5 SKEENA IN PRACTICE

We explore the effort needed to adopt Skeena in real systems, by enabling cross-engine transactions in open-source MySQL² between its default storage-centric InnoDB and ERMA [36]. MySQL defines a set of core interfaces (e.g., search, update and commit) for engines to implement [47]. This allowed us to integrate ERMA easily with < 2000 LoC.³ InnoDB and ERMA share MySQL’s SQL layer and thread pool.⁴ The application specifies each table’s home engine in its schema, which is managed by existing MySQL features.

To use CSR, it is necessary to understand each engine’s database model. ERMA closely follows our database model. To obtain a snapshot, the thread reads the counter without latching, which is much cheaper than InnoDB based on latching. So we use ERMA as the anchor engine. InnoDB uses transaction IDs (TIDs) to determine record visibility and ordering. We describe how we reconcile the differences between real implementation and our database model. Each read-write transaction is uniquely identified by a TID drawn from a central counter. Each record is stamped with the TID of the transaction that last updated it. Updates are handled in-place. Old versions are generated on-demand using undo logs. The freshness (or the amount of undo log to apply) is determined by the transaction’s read view (snapshot) which is acquired upon the first data access. A read view consists of low/high watermarks (TIDs) and an active transactions list captured at the read view’s creation time. The transaction is not allowed to see versions created by transactions with TIDs above the high watermark, but can see the results of transactions with TIDs below the low watermark. Versions created by transactions with TIDs between the two watermarks are invisible if they are active. As a result, read views are not directly comparable, deviating from our database model.

Our solution is to use the high watermark in CSR.⁵ Specifically, sub-transactions in InnoDB first acquire the latest read view using the original approach. We then adjust its high watermark using CSR and leave the active transactions list unchanged. In case the new high watermark is even lower than the low watermark, we adjust both to be the same as the high watermark. The sub-transaction can then use the adjusted read view as usual to test record visibility.

Upon commit, InnoDB assigns a `serialisation_no` drawn from the TID counter to denote commit ordering, which we use for Skeena’s commit check.⁶ We broke the monolithic commit function into pre/post-commit functions (Section 4.5). The pre-commit function only acquires a `serialisation_no`, leaving the remaining logic to post-commit. For atomic commit, we piggyback on ERMA’s commit pipelining. We extend the commit entry design in ERMA to include commit LSNs in both engines, along with a MySQL callback for notifying the client of concluded transactions.⁷

In total, we modified 83 LoC in InnoDB for it to use Skeena to choose read views and commit sub-transactions. CSR is implemented as a separate module of ~600LoC.⁸ For ERMA, we only modified its commit pipelining code to consider both engines.

²Based on MySQL 8.0 at <https://github.com/mysql/mysql-server>.

³Details in `ha_erma.cc` in our code repository (<https://github.com/sfu-dis/skeena>).

⁴Adopted from https://github.com/percona/percona-server/blob/8.0/sql/threadpool.*.

⁵Lines 2233–2254 of `trx0trx.cc` in our code repository.

⁶Lines 1378–1480 of `trx0trx.cc` in our code repository.

⁷Lines 180–222 of `sm-log-alloc.cpp` in our code repository.

⁸Details in `gtt.{cc,h}` in our code repository.

6 SKEENA IN ACTION

We empirically evaluate Skeena under microbenchmarks and realistic workloads. Through experiments, we show that:

- Skeena retains the performance benefits brought by memory-optimized engines in fast-slow systems;
- Skeena only incurs a very small amount of overhead for cross-engine transactions;
- By judiciously placing tables in different engines, Skeena can effectively improve performance for realistic workloads.

6.1 Experimental Setup

We run experiments on a dual-socket server equipped with two 20-core Intel Xeon Gold 6242R CPUs (80 hyperthreads in total), 384GB of main memory and a 400GB Micron SSD with peak bandwidth of 760MB/s. Each CPU has 35.75MB of cache and is clocked at 3.1GHz. All experiments are conducted in MySQL 8.0 with InnoDB and ERMA. We use SysBench [52] to issue benchmarks. To reduce networking overhead, we pin MySQL server and the client (SysBench) to two different CPU sockets, and use a Unix Domain Socket between the server and client [70]. We use `jmalloc` [25] to avoid memory management becoming a major bottleneck. We use SI (repeatable read in InnoDB) to run all experiments as they are widely used in practice, and reinitialize the database for each run which then starts with a warm buffer pool. We report the average throughput and latency of three 60-second runs.

ERMA is memory-optimized so all records are in heap memory. For InnoDB, we test both the memory- and storage-resident cases: the memory-resident variant (InnoDB-M) uses a large enough buffer pool to avoid accessing storage; the storage-resident variant (InnoDB) uses a small buffer pool that would mandate accessing the storage stack. To stress test Skeena, we store persistent data (such as data files and logs) in `tmpfs`, so that I/O is as fast as memory, making it easier to expose Skeena’s overhead. To understand the performance under more realistic workloads, we also run experiments using a real SSD; `tmpfs` is used unless otherwise specified.

6.2 Benchmarks

We use YCSB-like [17] microbenchmarks and TPC-C [62] (based on Percona’s implementation [54]) to test Skeena and explore the effect of cross-engine transactions.

Microbenchmarks. We devise three microbenchmarks based on access patterns: read-only, read-write and write-only. Unless otherwise specified, each transaction accesses ten records randomly chosen from a set of tables following a uniform distribution: (1) for read-write transactions, eight out of the ten accesses are point reads and two are updates, and (2) for each engine, we create 250 tables, each of which contains a certain number of records depending on whether the experiment is memory- or storage-resident for InnoDB. Each record is 232-byte, consisting of two INTEGERS and one VARCHAR. For memory-resident experiments, each table contains 25000 records, bringing the total data size of 250 tables to ~1.35GB; the buffer pool size in InnoDB is set to 32GB. For storage-resident experiments, we set each table to contain 250000 records, and the total data size is ~13.5GB; we set the buffer pool to be 2GB. Under both settings, ERMA is populated with the same amount of data as InnoDB (i.e., 500 tables across two engines).

Table 3: Throughput (TPS) of single-engine microbenchmarks (80 connections) and TPC-C (50 connections). Skeena (-S) incurs negligible overhead and retains ERMIA’s high performance.

Scheme	Read-only	Read-write	Write-only	TPC-C
ERMIA	1,427,071	1,252,146	1,091,606	7,550
ERMIA-S	1,430,137	1,253,368	1,095,056	7,546
InnoDB-M	1,326,710	930,249	710,697	626
InnoDB-MS	1,310,809	915,406	711,425	612
InnoDB	456,672	420,328	194,446	277
InnoDB-S	453,781	420,474	194,412	261

TPC-C. We use TPC-C for the dual-purpose of (1) testing Skeena under non-trivial transactions, and (2) exploring the potential benefits of cross-engine transactions in realistic scenarios. We run both memory- and storage-resident experiments: the former sets the scale factor to be the number of connections and the latter uses 200 warehouses. For memory-resident experiments, each connection works on a different home warehouse, but the 1% of New-Order and 15% of Payment transactions may respectively access a remote warehouse; we set InnoDB buffer pool to be 32GB which is large enough to hold all the data (~14GB). With 200 warehouses for storage-resident experiments, the total data size is ~55GB, for which we set InnoDB to use a buffer pool of 5GB and set each thread (connection) to always pick a random warehouse as its home warehouse to ensure the footprint covers the entire database. Finally, we gradually move tables from InnoDB to ERMIA, making the affected transactions cross-engine. This allows us to explore the effectiveness of cross-engine transactions and distill several useful suggestions on how to optimize performance in fast-slow systems; we discuss more detailed setups later.

6.3 Single-Engine Performance

An important goal of Skeena is to ensure single-engine transactions (especially those in the faster engine, ERMIA) pay little additional cost. We evaluate this aspect by turning Skeena on and off under six ERMIA- and InnoDB-only variants. To stress test Skeena, we use the memory-resident InnoDB (InnoDB-M) and the storage-resident InnoDB with tmpfs (InnoDB-S). Table 3 summarizes the results; variants with Skeena turned on carry an S suffix. In all cases, Skeena incurs negligible overhead with the slightly more complex logic in commit pipelining. Note that “single-engine” transactions in InnoDB/InnoDB-M are in fact cross-engine, as they must follow the start order in the anchor engine (ERMIA) even if they do not access any records in ERMIA. This means CSR will only maintain a single mapping (using ERMIA’s initial snapshot) which incurs a constant but very small amount of overhead (up to 5.6%); garbage collection is also never needed with a single mapping. Compared to InnoDB, InnoDB-M performs up to over ~3× better thanks to its large buffer pool. InnoDB-M and InnoDB-MS perform similarly to ERMIA under the read-only microbenchmark, but fall behind as we add more writes to the workload, signifying the potential benefits a memory-optimized engine could bring (more later). ERMIA-S performs as well as ERMIA since CSR is never used. These results verify that Skeena retains the advantage of memory-optimized engines.

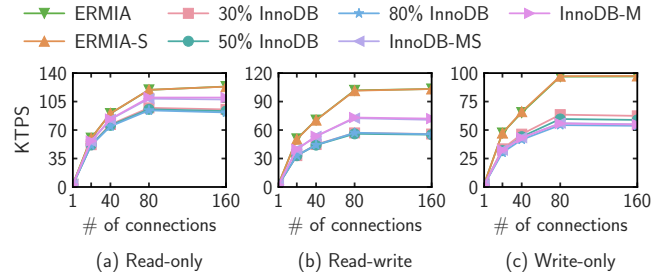


Figure 6: Throughput under memory-resident microbenchmarks. CSR cost can be comparable to that of reading records, causing InnoDB-M to outperform cross-engine cases.

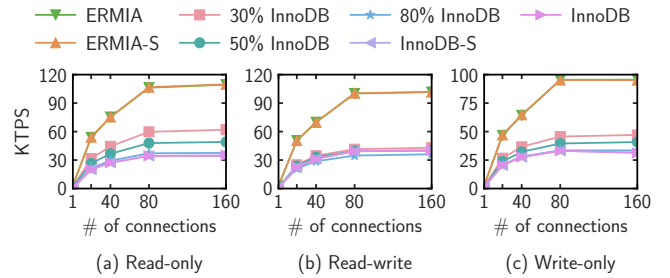


Figure 7: Throughput under storage-resident microbenchmarks. Performance improves with more accesses in ERMIA.

6.4 Cross-Engine Performance

Now we explore the behavior of cross-engine transactions using microbenchmarks. For each transaction, we vary the percentage of InnoDB and ERMIA accesses out of ten accesses. For example, with 30% InnoDB, three accesses per transaction are done in InnoDB, the remaining seven accesses go to ERMIA. We use the same -M and -S notations from Section 6.3 for single-engine transactions whose results are shown to calibrate expectations. For cross-engine transactions we mark the percentage of InnoDB accesses and note whether the experiment is memory- or storage-resident as needed.

InnoDB is more heavyweight, so more accesses in it should lower performance, e.g., transactions with 30% InnoDB accesses should perform better than those that only access InnoDB. However, Figures 6(a)–(b) show the opposite: InnoDB-M outperforms the cross-engine 30–80% InnoDB. The reason is two-fold. First, ERMIA writes a commit log record for read-only transactions. So with more ERMIA accesses, CSR becomes larger and slower to access. This is non-negligible for read-intensive workloads, which are very lightweight in ERMIA. Second, under InnoDB-MS, CSR is very small and only maintains one mapping as we mentioned earlier. However, under 30–80% InnoDB, more ERMIA accesses lead to more mappings in CSR, which then becomes more expensive to query. The memory-resident write-only workload follows the expectation in Figure 6(c), although the difference is not significant due to InnoDB’s low raw performance. As the workload becomes storage-resident, Skeena’s overhead becomes negligible, with more ERMIA accesses leading to higher performance: in Figure 7, 30% InnoDB is up to 75%/40% faster than InnoDB for read-only/write-only workloads.

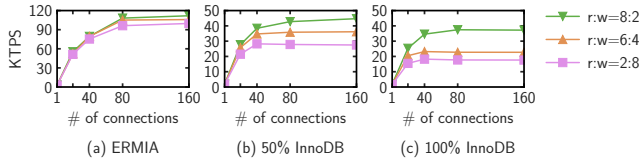


Figure 8: Throughput under storage-resident microbenchmarks with different read/write ratios.

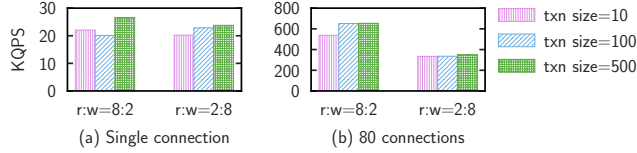


Figure 9: Throughput under storage-resident microbenchmarks with 50% InnoDB accesses and different transaction sizes.

6.5 Impact of Transaction Size and Mix

We quantify the impact of transaction sizes along with the impact of read/write mixes. We first show the throughput of the storage-resident workload with 50% InnoDB accesses under varying transaction sizes and read/write ratios. In Figure 9, longer transactions can lower performance, but do not delay CSR index recycling as all the threads start and commit transactions at roughly the same pace. We also ran experiments with a mix of long and short transactions, where a fixed 0–20% of connections only run long transactions with 500 queries. We observe the number of indexes increases by ~50 per second (with 1000 entries per index). Once the number of indexes reaches a configurable threshold (1000 in our experiments), recycling kicks in and works well across varying percentages of long-running transactions. Neither long-running transactions nor capacity and threshold settings affect QPS in a noticeable way. Figure 8 shows the throughput under different read/write ratios. With more writes, throughput drops as handling writes is more complex. The impact on ERMIA is very small in Figure 8(a). When the read ratio drops from 80% to 60%, under 100% InnoDB, performance can drop by up to ~30%. Cross-engine transactions in Figure 8(b) with 50% InnoDB accesses still have the advantage over 100% InnoDB.

We use short transactions to explore the overhead of Skeena and transaction management. Each transaction issues two queries (one per engine for cross-engine cases). In Figure 10, when the transaction only accesses ERMIA, the performance remains similar across all workloads, as all data is in memory and the relative speed difference for read/write in ERMIA is small. For 100% InnoDB, with writes, the performance drops by up to ~25%. The cross-engine 50% InnoDB exhibits the lowest performance due to extra time spent on CSR and the commit protocol. However, it is only slightly slower than 100% InnoDB as it is more heavyweight to process writes in InnoDB than accessing CSR which is purely in-memory.

6.6 Skewed Accesses

We test the storage-resident workload with different ERMIA/InnoDB accesses mixes with 80% of read and 20% of write per transaction. As Figure 11 shows, throughput remains similar across




Figure 10: Throughput under memory-resident benchmarks with short transactions (two queries) under 80 connections.




Figure 11: Throughput of cross-engine read-write transactions under storage-resident microbenchmarks and varying skewness.

Table 4: Throughput (TPS) of a storage-resident cross-engine workload with 50% InnoDB accesses under varying buffer pool hit ratios.

Number of Connections	100%	99%	90%	70%
1	1,973	1,972	1,939	1,866
80	36,749	36,186	35,414	28,369

different skewness for ERMIA-only cases. The result seems counter-intuitive: a skewed workload leads to a smaller footprint which should change performance. The smaller footprint may lead to higher contention on locks/latches (lower performance). It could also give better buffer pool and CPU cache locality (higher performance). Our profiling results show that the actual CPU time spent on ERMIA is less than 5%. The remaining > 95% of CPU time is taken by MySQL’s SQL and networking layers, overshadowing the effect brought by the smaller footprint. As we add more accesses to InnoDB, 50% and 100% InnoDB show in general higher performance but not by a lot. The main reason is the storage stack’s overhead starts to dominate once we access InnoDB tables. Skewness (smaller footprint) therefore does not show an obvious impact.

6.7 Impact of Slower Storage

Now we run experiments in a more realistic environment that uses an actual SSD to store table and log data. We use the storage-resident setting and vary the buffer pool size such that the hit ratio is between 70% and 100%. In this experiment, to stress Skeena’s CSR structure, each 10-record transaction accesses 5/5 records in InnoDB/ERMIA with 80% reads and 20% writes. As shown in Table 4, throughput under a single connection remains stable across different hit ratios because the working set is relatively small which gives better locality. With more (80) connections, throughput drops as the hit ratio drops because more I/Os on the SSD are needed. Our profiling results (details not shown here for brevity) indicate

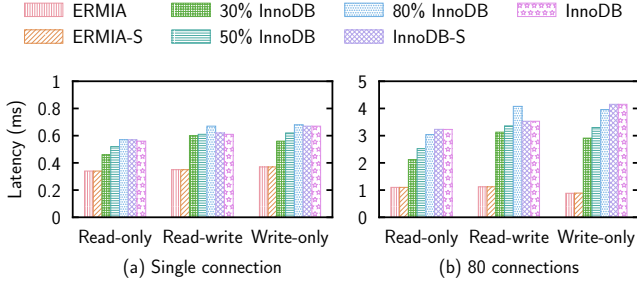


Figure 12: 95th percentile latency at a single connection and 80 connections under various storage-resident workloads.

that across all cases Skeena occupies less than 5% of total CPU time, most of which was on accessing the engine and the SQL layer.

6.8 Transaction Latency

Our final microbenchmark explores how Skeena impacts transaction latency using the storage-resident microbenchmarks. We run the same microbenchmarks done in Section 6.4 and present results obtained at one and 80 connections where the system is almost idle and fully saturated, respectively. In Figure 12, at 80 connections the absolute latency unavoidably increases. But in all cases Skeena does not increase latency noticeably for single-engine transactions: ERMIA and ERMIA-S do not use CSR, while the overhead for InnoDB-S is a constant. As expected, latency increases proportionally with more accesses in InnoDB.

6.9 Effectiveness of Cross-Engine Transactions

Realistic workloads may use cross-engine transactions to improve performance (using a main-memory engine) and/or reduce storage cost with the storage-centric engine. This is done by placing different tables in different engines. We use TPC-C to explore this aspect; we observed similar trends for memory- and storage-resident setups, so we only show the results from the storage-resident setup. As shown in Figure 13, we start with all tables in InnoDB (bottom) and gradually move them to ERMIA (up). Overall, throughput improves with more tables in ERMIA, but performance does not change much until the New-Orders table is moved to ERMIA ($\sim 10\times$ faster than 100% InnoDB).

Placing New-Orders in ERMIA should be the key to improving overall throughput, but it is unclear how each transaction benefits from this. We further run individual TPC-C transactions (instead of the full mix) under a variant that only places New-Orders in ERMIA (leaving the rest in InnoDB). Figure 14 compares its throughput (denoted as +New-Orders) to other variants, among which ++Orders and ++New-Orders refer to the corresponding rows in Figure 13, respectively. In Figures 14(a)–(b) and 14(d)–(e), placing New-Orders in ERMIA does not affect the New-Order, Payment, Stock-Level and Order-Status transactions. This is because Stock-Level, Order-Status and Payment do not access New-Orders, and the New-Order transaction only inserts one row into New-Orders. However, the Delivery transaction intensively accesses New-Orders with range scan and aggregation operations. In Figure 14(c), placing New-Orders in ERMIA accelerates Delivery by $\sim 30\times$ as InnoDB has to hold locks

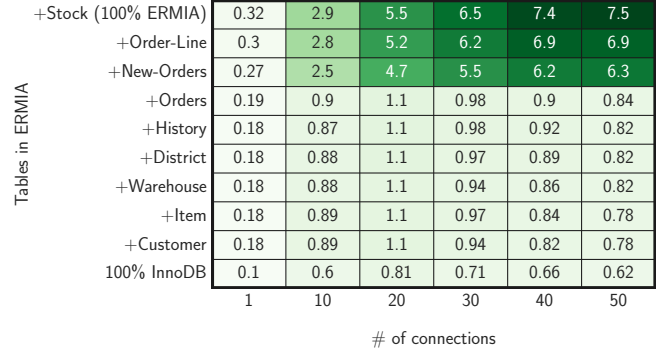


Figure 13: TPC-C throughput (kTPS) with tables gradually placed in ERMIA (cumulatively from bottom up). Placing New-Orders in ERMIA is the key to improve overall mix performance.

for records to be deleted. Therefore, improvement on Delivery is the main reason for the improved overall performance.

These results show that table placement affects different transactions in different scales. We analyze this effect in Figure 15 using individual transactions at 50 connections. Tables are gradually moved from InnoDB to ERMIA. As Figure 15 shows, placing Customer in ERMIA alone allows the Payment and Order-Status transactions to perform 6–7 \times better, whereas the Stock-Level transaction benefits the most when the Stock table is placed in ERMIA.

Recommended End-to-End Cross-Engine TPC-C. Based on the previous results, we recommend three schemes:

- New-Order-Opt: The Customer and Item tables are placed in ERMIA to optimize the New-Order transaction.
- Payment-Opt: Only Customer is placed in ERMIA to optimize the Payment transaction, which intensively accesses Customer.
- Archive: All the tables except History are placed in InnoDB, leveraging its cheaper storage cost compared to ERMIA.

The first two schemes aim to optimize database accesses with select use of main-memory tables, while Archive attempts to reduce storage cost of in-memory databases using a traditional engine. Figure 16 shows how they compare to baselines that place all tables in ERMIA and InnoDB. New-Order-Opt and Payment-Opt improve the performance of the affected transactions compared to InnoDB. Since Archive executes almost fully in ERMIA, its performance overlaps with ERMIA because History is never queried and only occupies less than 600MB of space. In reality, such workloads can run for much longer and accumulate much more data; placing it in InnoDB can drastically reduce storage cost as main memory is much more expensive than SSDs and disks.

Impact on Abort Rate. As we discussed in Section 4, Skeena can cause transactions to abort if they cannot find appropriate snapshots or fail commit check. We compare the abort rate using the preferred end-to-end TPC-C schemes. The workload runs under the memory-resident setup which exhibits low contention (each transaction works on a fixed home warehouse) so that snapshot selection and commit check in Skeena are the main source of aborts. As baselines, single-engine InnoDB and ERMIA exhibit an abort rate of 0.43% and 0.47%, respectively. With Skeena, New-Order-Opt, Payment-Opt and Archive respectively exhibit an abort rate of

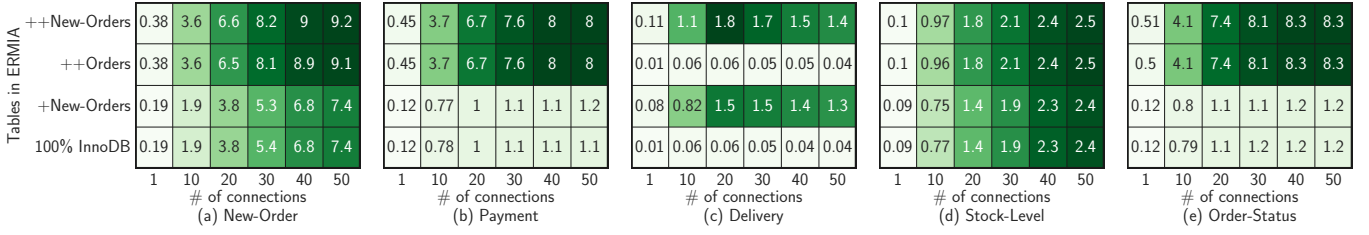


Figure 14: Throughput (kTPS) of individual TPC-C transactions when only New-Orders is placed in ERMIA (+New-Orders) compared to 100% InnoDB and two other variants that place cumulatively up to Orders and New-Orders in ERMIA.

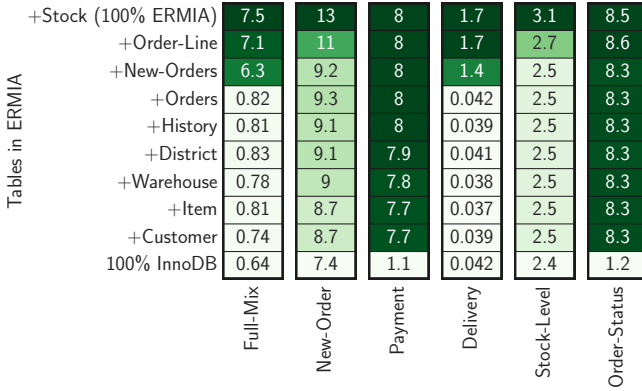


Figure 15: Throughput (kTPS) of TPC-C mix and individual transactions by varying table placement at 50 connections.

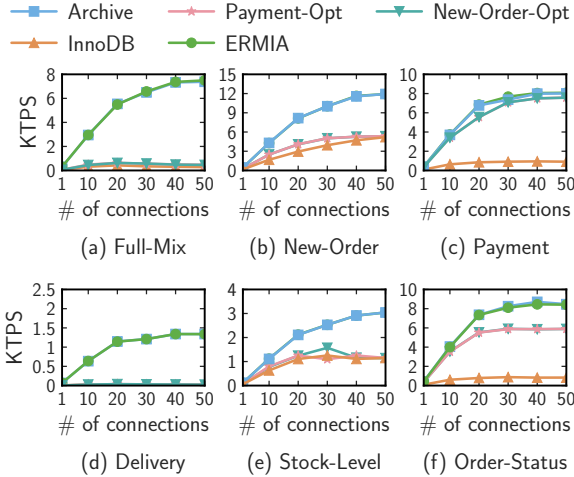


Figure 16: TPC-C throughput under select table placement schemes that optimize for different application scenarios.

0.61%, 0.54% and 0.45%. Across individual runs, we observed up to 0.3% higher abort rate than single-engine cases. The impact on abort rate is more pronounced for the read-write microbenchmarks, where up to ~5% additional cross-engine transactions are aborted due to Skeena. These results corroborate with prior work [8] that the impact on abort rate is very small for realistic workloads.

7 RELATED WORK

Our work builds upon rich literature on federated systems, distributed and replicated SI systems, and modern database engines.

Federated and Polystore Systems. Guaranteeing serializability is challenging due to the full autonomy and heterogeneity of member systems. Georgakopoulos et al. [29] proposed a ticket method and Superdatabase [56] exports transaction ordering for consistent updates. Breitbart et al. [12] proposed a protocol that detects cycles for consistency and resolving deadlocks. Myriad [31, 42] uses 2PL and 2PC for serializability and atomicity. Recent polystores [3, 4, 23] mainly focus on analytics, while our focus is OLTP.

DSI and Replication. Prior work identified the anomalies that would lead to inconsistent snapshots [58]. Binnig et al. [8] further identified more anomalies and proposed correctness conditions. Skeena is based on these results. Generalized SI [24] allows transactions to use older snapshots in replicated databases. Skeena shares the similar property by selecting snapshots using CSR.

Modern Database Engines. Multicore CPUs and large DRAM have led to numerous memory-optimized engines [21, 28, 35–37, 41, 43, 61, 63]. They feature new designs on indexing, concurrency control and logging protocols that drastically improve performance. Our work explores one of the possible ways to adopt them in practice. Some recent systems [40, 50] efficiently leverage modern fast SSDs to approach in-memory performance while keeping storage cost lower than main-memory systems. It is interesting future work to compare them with cross-engine systems.

8 CONCLUSION

Cross-engine transactions can be very useful in modern fast-slow multi-engine systems, but are poorly supported with various limitations. This paper proposes Skeena, a holistic approach to efficient and consistent cross-engine transactions. Skeena consists of a cross-engine snapshot registry (CSR) that tracks snapshots and a commit protocol for multi-engine systems. Skeena can be easily adopted by real systems, as shown by our experience with MySQL. Evaluation on a 40-core server shows that Skeena incurs negligible overhead and maintains the benefits of memory-optimized engines.

ACKNOWLEDGMENTS

We thank Chong Chen, Per-Åke Larson, Qiang Liu, Chengwei Zhang, Zongquan Zhang, Yanhui Zhong and Qingqing Zhou for their valuable discussions and comments on this project. We also thank the reviewers for their constructive feedback. This work received support from Huawei Cloud Database Innovation Lab.

REFERENCES

- [1] Atul Adya. 1999. Weak Consistency: A Generalized Theory and Optimistic Implementations for Distributed Transactions. PhD Thesis and Technical Report MIT/LCS/TR-786.
- [2] A. Adya, B. Liskov, and P. O’Neil. 2000. Generalized isolation level definitions. In *Proceedings of 16th International Conference on Data Engineering (Cat. No. 00CB37073)*. 67–78.
- [3] Divy Agrawal, Sanjay Chawla, Bertty Contreras-Rojas, Ahmed Elmagarmid, Yasser Idris, Zoi Kaoudi, Sebastian Kruse, Ji Lucas, Essam Mansour, Mourad Ouzzani, et al. 2018. RHEEM: enabling cross-platform data processing: may the big data be with you! *PVLDB* 11, 11 (2018), 1414–1427.
- [4] Rana Alotaibi, Damian Burszty, Alin Deutsch, Ioana Manolescu, and Stamatis Zampetakis. 2019. Towards scalable hybrid stores: constraint-based rewriting to the rescue. In *Proceedings of the 2019 International Conference on Management of Data*. 1660–1677.
- [5] Hillel Avni, Alisher Aliev, Oren Amor, Aharon Avitzur, Ilan Bronshtein, Eli Ginot, Shay Goikman, Eliezer Levy, Idan Levy, Fuyang Lu, Liran Mishali, Yeqin Mo, Nir Pachter, Dima Sivov, Vinoth Veeraraghavan, Vladi Vexler, Lei Wang, and Peng Wang. 2020. Industrial-Strength OLTP Using Main Memory and Many Cores. *PVLDB* 13, 12 (Aug. 2020), 3099–3111.
- [6] Hal Berenson, Phil Bernstein, Jim Gray, Jim Melton, Elizabeth O’Neil, and Patrick O’Neil. 1995. A Critique of ANSI SQL Isolation Levels. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data (SIGMOD ’95)*. 1–10.
- [7] Carsten Binnig, Andrew Crotty, Alex Galakatos, Tim Kraska, and Erfan Zamanian. 2016. The End of Slow Networks: It’s Time for a Redesign. *PVLDB* 9, 7 (March 2016), 528–539.
- [8] Carsten Binnig, Stefan Hildenbrand, Franz Färber, Donald Kossmann, Juchang Lee, and Norman May. 2014. Distributed Snapshot Isolation: Global Transactions Pay Globally, Local Transactions Pay Locally. *The VLDB Journal* 23, 6 (Dec. 2014), 987–1011.
- [9] Mihaela A. Bornea, Orion Hodson, Sameh Elnikety, and Alan Fekete. 2011. One-Copy Serializability with Snapshot Isolation under the Hood. In *Proceedings of the 2011 IEEE 27th International Conference on Data Engineering (ICDE ’11)*. 625–636.
- [10] Jan Böttcher, Viktor Leis, Thomas Neumann, and Alfons Kemper. 2019. Scalable Garbage Collection for In-Memory MVCC Systems. *PVLDB* 13, 2 (Oct. 2019), 128–141.
- [11] Yuri Breitbart, Hector Garcia-Molina, and Avi Silberschatz. 1992. Overview of Multidatabase Transaction Management. *The VLDB Journal* 1, 2 (Oct. 1992), 181–240.
- [12] Yuri Breitbart and Avi Silberschatz. 1988. Multidatabase Update Issues. In *Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data (SIGMOD ’88)*. 135–142.
- [13] Michael J. Cahill, Uwe Röhm, and Alan D. Fekete. 2009. Serializable Isolation for Snapshot Databases. *ACM Trans. Database Syst.* 34, 4, Article 20 (Dec. 2009), 42 pages.
- [14] M. A. Casanova and P. A. Bernstein. 1981. General Purpose Schedulers for Database Systems. *Acta Inf.* 15, 4 (Aug. 1981), 471.
- [15] Arvola Chan, Stephen Fox, Wen-Te K. Lin, Anil Nori, and Daniel R. Ries. 1982. The Implementation of an Integrated Concurrency Control and Recovery Scheme. In *Proceedings of the 1982 ACM SIGMOD International Conference on Management of Data (SIGMOD ’82)*. 184–191.
- [16] A. Chan and R. Gray. 1985. Implementing Distributed Read-Only Transactions. *IEEE Transactions on Software Engineering* SE-11, 2 (1985), 205–212.
- [17] Brian F Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking cloud serving systems with YCSB. In *Proceedings of the 1st ACM symposium on Cloud computing*. 143–154.
- [18] Khuzaima Daudjee and Kenneth Salem. 2006. Lazy Database Replication with Snapshot Isolation. In *Proceedings of the 32nd International Conference on Very Large Data Bases (VLDB ’06)*. 715–726.
- [19] Pinal Dave. 2019. SQL Server – Memory Optimized Tables, Transactions, Isolation Level and Error. <https://blog.sqlauthority.com/2019/06/11/sql-server-memory-optimized-tables-transactions-isolation-level-and-error/>
- [20] Kalen Delaney. 2016. SQL Server In-Memory OLTP Internals for SQL Server 2016. *Microsoft SQL Server Docs* (2016). https://download.microsoft.com/download/8/3/6/8360731A-A27C-4684-BC88-FC7B5849A133/SQL_Server_2016_In_Memory_OLTP_White_Paper.pdf
- [21] Cristian Diaconu, Craig Freedman, Erik Ismert, Per-Ake Larson, Pravin Mittal, Ryan Stonecipher, Nitin Verma, and Mike Zwilling. 2013. Hekaton: SQL Server’s Memory-Optimized OLTP Engine. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data (SIGMOD ’13)*. 1243–1254.
- [22] Deborah J. DuBourdieu. 1982. Implementation of Distributed Transactions. In *Proceedings of the Sixth Berkeley Workshop on Distributed Data Management and Computer Networks*. 81–94.
- [23] Aaron J Elmore, Jennie Duggan, Michael Stonebraker, Magdalena Balazinska, Ugur Cetintemel, Vijay Gadepally, Jeffrey Heer, Bill Howe, Jeremy Kepner, Tim Kraska, et al. 2015. A demonstration of the BigDAWG polystore system. *PVLDB* 8, 12 (2015), 1908.
- [24] Sameh Elnikety, Willy Zwaenepoel, and Fernando Pedone. 2005. Database Replication Using Generalized Snapshot Isolation. In *Proceedings of the 24th IEEE Symposium on Reliable Distributed Systems (SRDS ’05)*. 73–84.
- [25] Jason Evans. 2006. A Scalable Concurrent malloc (3) Implementation for FreeBSD. In *Proceedings of the BSDCan Conference*.
- [26] Jose M. Faleiro and Daniel J. Abadi. 2015. Rethinking Serializable Multiversion Concurrency Control. *PVLDB* 8, 11 (July 2015), 1190–1201.
- [27] Alan Fekete, Dimitrios Liarokapis, Elizabeth O’Neil, Patrick O’Neil, and Dennis Shasha. 2005. Making Snapshot Isolation Serializable. *ACM Trans. Database Syst.* 30, 2 (June 2005), 492–528.
- [28] H. Garcia-Molina and K. Salem. 1992. Main memory database systems: an overview. *IEEE Transactions on Knowledge and Data Engineering* 4, 6 (1992), 509–516.
- [29] D. Georgakopoulos, M. Rusinkiewicz, and A. Sheth. 1991. On serializability of multidatabase transactions through forced local conflicts. In *Proceedings of the Seventh International Conference on Data Engineering*. 314–323.
- [30] Goetz Graefe, Mark Lillibridge, Harumi Kuno, Joseph Tucek, and Alistair Veitch. 2013. Controlled Lock Violation. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data (SIGMOD ’13)*. 85–96.
- [31] S.-Y. Hwang, E.-P. Lim, H.-R. Yang, S. Musukula, K. Mediratta, M. Ganesh, D. Clements, J. Stenoien, and J. Srivastava. 1994. The MYRIAD Federated Database Prototype. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data (SIGMOD ’94)*. 518.
- [32] Intel Corporation. 2016. Intel 64 and IA-32 Architectures Software Developer Manuals. (Oct. 2016).
- [33] Ryan Johnson, Ippokratis Pandis, Nikos Hardavellas, Anastasia Ailamaki, and Babak Falsafi. 2009. Shore-MT: A Scalable Storage Manager for the Multicore Era. In *Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology (EDBT ’09)*. 24–35.
- [34] Ryan Johnson, Ippokratis Pandis, Radu Stoica, Manos Athanassoulis, and Anastasia Ailamaki. 2010. Aether: A Scalable Approach to Logging. *PVLDB* 3, 1 (Sept. 2010), 681–692.
- [35] Alfons Kemper and Thomas Neumann. 2011. HyPer: A Hybrid OLTP&OLAP Main Memory Database System Based on Virtual Memory Snapshots. In *Proceedings of the 2011 IEEE 27th International Conference on Data Engineering (ICDE ’11)*. 195–206.
- [36] Kangnyeon Kim, Tianzheng Wang, Ryan Johnson, and Ippokratis Pandis. 2016. ERMA: Fast memory-optimized database system for heterogeneous workloads. In *Proceedings of the 2016 International Conference on Management of Data*. 1675–1687.
- [37] Hideaki Kimura. 2015. FOEDUS: OLTP Engine for a Thousand Cores and NVRAM. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD ’15)*. 691–706.
- [38] H. T. Kung and John T. Robinson. 1981. On Optimistic Methods for Concurrency Control. *ACM Trans. Database Syst.* 6, 2 (June 1981), 213–226.
- [39] Per-Ake Larson, Spyros Blanas, Cristian Diaconu, Craig Freedman, Jignesh M. Patel, and Mike Zwilling. 2011. High-Performance Concurrency Control Mechanisms for Main-Memory Databases. *PVLDB* 5, 4 (Dec. 2011), 298–309.
- [40] 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 ICDE)*. 185–196. <https://doi.org/10.1109/ICDE.2018.00026>
- [41] Justin Levandoski, David Lomet, Sudipta Sengupta, Ryan Stutsman, and Rui Wang. 2015. High Performance Transactions in Deuteronomy. In *Conference on Innovative Data Systems Research (CIDR 2015)*.
- [42] Ee-Peng Lim, Sah-Yih Hwang, Jaideep Srivastava, Dave Clements, and M. Ganesh. 1995. Myriad: Design and Implementation of a Federated Database Prototype. *Softw. Pract. Exper.* 25, 5 (May 1995), 533–562.
- [43] Hyeontaek Lim, Michael Kaminsky, and David G. Andersen. 2017. Cicada: Dependably Fast Multi-Core In-Memory Transactions. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD ’17)*. 21–35.
- [44] Yandong Mao, Eddie Kohler, and Robert Tappan Morris. 2012. Cache craftiness for fast multicore key-value storage. In *Proceedings of the 7th ACM european conference on Computer Systems*. 183–196.
- [45] Prashanth Menon, Amadou Ngom, Lin Ma, Todd C. Mowry, and Andrew Pavlo. 2020. Permutable Compiled Queries: Dynamically Adapting Compiled Queries without Recompile. *Proc. VLDB Endow.* 14, 2 (Oct. 2020), 101–113.
- [46] Microsoft. 2016. *Microsoft SQL Documentation* (2016). <https://docs.microsoft.com/en-us/sql/relational-databases/in-memory-oltp/introduction-to-memory-optimized-tables?view=sql-server-ver15>
- [47] MySQL 8.0 Reference Manual. 2021. Alternative Storage Engines. <https://dev.mysql.com/doc/refman/8.0/en/storage-engines.html>
- [48] MySQL 8.0 Reference Manual. 2021. SET TRANSACTION Statement. (2021). <https://dev.mysql.com/doc/refman/8.0/en/set-transaction.html>
- [49] Thomas Neumann. 2011. Efficiently Compiling Efficient Query Plans for Modern Hardware. *PVLDB* 4, 9 (June 2011), 539–550.
- [50] Thomas Neumann and Michael J Freitag. 2020. Umbra: A Disk-Based System with In-Memory Performance. In *10th Conference on Innovative Data Systems*

- Research, CIDR 2020, Amsterdam, The Netherlands, January 12-15, 2020, Online Proceedings. <http://cidrdb.org/cidr2020/papers/p29-neumann-cidr20.pdf>
- [51] Thomas Neumann, Tobias Mühlbauer, and Alfons Kemper. 2015. Fast Serializable Multi-Version Concurrency Control for Main-Memory Database Systems. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD '15)*. 677–689.
 - [52] Oracle. 2021. SysBench Benchmark Tool. <https://dev.mysql.com/downloads/benchmarks.html>
 - [53] Ippokratis Pandis, Pinar Tözün, Ryan Johnson, and Anastasia Ailamaki. 2011. PLP: Page Latch-Free Shared-Everything OLTP. *PVLDB* 4, 10 (July 2011), 610–621.
 - [54] Percona. 2018. sysbench-tpcc. <https://github.com/Percona-Lab/sysbench-tpcc>
 - [55] PostgreSQL Wiki. 2021. Foreign Data Wrappers. https://wiki.postgresql.org/wiki/Foreign_data_wrappers
 - [56] Calton Pu. 1988. Superdatabases for Composition of Heterogeneous Databases. In *Proceedings of the Fourth International Conference on Data Engineering*. 548–555.
 - [57] Yoav Raz. 1992. The Principle of Commitment Ordering, or Guaranteeing Serializability in a Heterogeneous Environment of Multiple Autonomous Resource Mangers Using Atomic Commitment. In *18th International Conference on Very Large Data Bases Proceedings*. 292–312.
 - [58] Ralf Schenkel and Gerhard Weikum. 2000. Integrating Snapshot Isolation into Transactional Federations. In *Cooperative Information Systems*. 90–101.
 - [59] Ralf Schenkel, Gerhard Weikum, Norbert Weissenberg, and Xuequn Wu. 2000. Federated Transaction Management with Snapshot Isolation. In *Transactions and Database Dynamics*. 1–25.
 - [60] Amit P. Sheth and James A. Larson. 1990. Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases. *ACM Comput. Surv.* 22, 3 (Sept. 1990), 183–236.
 - [61] Michael Stonebraker, Samuel Madden, Daniel J. Abadi, Stavros Harizopoulos, Nabil Hachem, and Pat Helland. 2007. The End of an Architectural Era: (It's Time for a Complete Rewrite). (2007), 1150–1160.
 - [62] TPC. 2010. TPC Benchmark C (OLTP) Standard Specification, revision 5.11. <http://www.tpc.org/tpcc>
 - [63] Stephen Tu, Wenting Zheng, Eddie Kohler, Barbara Liskov, and Samuel Madden. 2013. Speedy Transactions in Multicore In-Memory Databases. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles (SOSP '13)*. 18–32.
 - [64] Tianzheng Wang and Ryan Johnson. 2014. Scalable Logging through Emerging Non-Volatile Memory. *PVLDB* 7, 10 (June 2014), 865–876.
 - [65] Tianzheng Wang, Ryan Johnson, Alan Fekete, and Ippokratis Pandis. 2017. Efficiently Making (Almost) Any Concurrency Control Mechanism Serializable. *The VLDB Journal* 26, 4 (Aug. 2017), 537–562.
 - [66] Tianzheng Wang, Ryan Johnson, and Ippokratis Pandis. 2017. Query Fresh: Log Shipping on Steroids. *PVLDB* 11, 4 (Dec. 2017), 406–419.
 - [67] Yingjun Wu, Joy Arulraj, Jiexi Lin, Ran Xian, and Andrew Pavlo. 2017. An Empirical Evaluation of In-Memory Multi-Version Concurrency Control. *PVLDB* 10, 7 (March 2017), 781–792.
 - [68] Yu Xia, Xiangyao Yu, Andrew Pavlo, and Srinivas Devadas. 2020. Taurus: Lightweight Parallel Logging for in-Memory Database Management Systems. *PVLDB* 14, 2 (Oct. 2020), 189–201.
 - [69] Xiangyao Yu, George Bezerra, Andrew Pavlo, Srinivas Devadas, and Michael Stonebraker. 2014. Staring into the Abyss: An Evaluation of Concurrency Control with One Thousand Cores. *PVLDB* 8, 3 (nov 2014), 209–220.
 - [70] Peter Zaitsev. 2020. Need to Connect to a Local MySQL Server? Use Unix Domain Socket! <https://www.percona.com/blog/2020/04/13/need-to-connect-to-a-local-mysql-server-use-unix-domain-socket>
 - [71] Erfan Zamanian, Carsten Binnig, Tim Harris, and Tim Kraska. 2017. The End of a Myth: Distributed Transactions Can Scale. *PVLDB* 10, 6 (Feb. 2017), 685–696.