Design Guidelines for Correct, Efficient, and Scalable Synchronization using One-Sided RDMA

TOBIAS ZIEGLER, Technische Universität Darmstadt, Germany
JACOB NELSON-SLIVON, Lehigh University, USA
VIKTOR LEIS, Technische Universität München, Germany
CARSTEN BINNIG, Technische Universität Darmstadt & DFKI, Germany

Remote data structures built with one-sided Remote Direct Memory Access (RDMA) are at the heart of many disaggregated database management systems today. Concurrent access to these data structures by thousands of remote workers necessitates a highly efficient synchronization scheme. Remarkably, our investigation reveals that existing synchronization schemes display substantial variations in performance and scalability. Even worse, some schemes do not correctly synchronize, resulting in rare and hard-to-detect data corruption. Motivated by these observations, we conduct the first comprehensive analysis of one-sided synchronization techniques and provide general principles for correct synchronization using one-sided RDMA. Our research demonstrates that adherence to these principles not only guarantees correctness but also results in substantial performance enhancements.

CCS Concepts: • **Information systems → Parallel and distributed DBMSs**; • **Networks → Network protocols**.

Additional Key Words and Phrases: Distributed Database Management Systems; RDMA; Synchronization;

ACM Reference Format:
Tobias Ziegler, Jacob Nelson-Slivon, Viktor Leis, and Carsten Binnig. 2023. Design Guidelines for Correct, Efficient, and Scalable Synchronization using One-Sided RDMA. *Proc. ACM Manag. Data* 1, 2, Article 131 (June 2023), 26 pages. https://doi.org/10.1145/3589276

1 INTRODUCTION

RDMA & disaggregated databases. Remote Direct Memory Access (RDMA) has quickly become one of the indispensable tools for building disaggregated database systems. Not only does RDMA provide single-digit microsecond network latencies, but it also provides efficient primitives for remote memory access. In particular, RDMA's *one-sided verbs* allow a compute server to read or write directly to a remote memory server while bypassing the remote CPU. Since memory servers frequently possess near-zero computational capacity [48] and most computational power is concentrated within the compute layer, one-sided RDMA proves to be well-suited for disaggregated DBMSs. Consequently, recent literature has explored how disaggregated DBMSs can leverage one-sided verbs [3, 8, 10, 22, 27, 44, 54–56, 58].

Synchronization of remote data structures. A key building block of these disaggregated DBMSs are remote data structures such as one-sided hash tables [30, 49, 61], B-Trees [47, 60],

Authors' addresses: Tobias Ziegler, tobias.ziegler@cs.tu-darmstadt.de, Technische Universität Darmstadt, Germany; Jacob Nelson-Slivon, jjn217@lehigh.edu, Lehigh University, USA; Viktor Leis, leis@in.tum.de, Technische Universität München, Germany; Carsten Binnig, carsten.binnig@cs.tu-darmstadt.de, Technische Universität Darmstadt & DFKI, Germany.

2836-6573/2023/6-ART131 $15.00
https://doi.org/10.1145/3589276

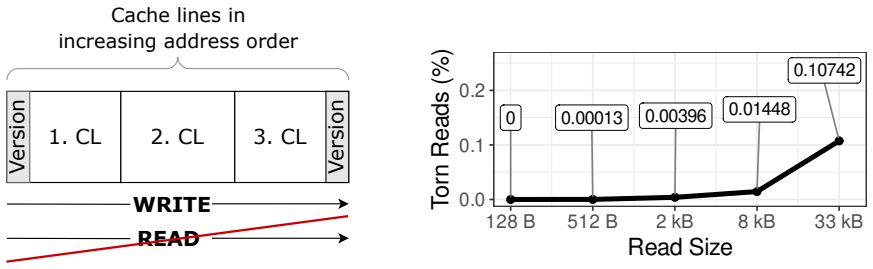

(a) Faulty Optimistic Scheme (b) Percentage of Torn Reads

Fig. 1. Incorrect optimistic synchronization

or SkipLists [29] which enable efficient access to remote data. But because one-sided operations bypass the remote CPU, traditional storage server-side synchronization techniques where the remote CPU is in charge[1] do not work. Instead, various one-sided synchronization techniques have been proposed [8, 32, 53, 60]. Those techniques can be categorized into *pessimistic* and *optimistic* schemes. While pessimistic schemes *prevent* concurrent modifications, optimistic schemes *detect* (and handle) concurrent modifications. These approaches fundamentally differ in their scalability and performance characteristics.

Performance is key. Remote data structures may need to serve thousands of clients connecting from several compute servers. With such a high degree of concurrency, the performance depends on how well the implemented one-sided synchronization scheme performs. While individual papers have proposed various one-sided synchronization schemes [8, 31, 60], it is surprising that there has not yet been a systematic study of these schemes under comparable workloads and conditions. This paper provides the first in-depth performance analysis. We show that small design choices when implementing a scheme can severely impact its performance and lead to performance bottlenecks. For example, contrary to expectations, data alignment hinders the scalability of pessimistic one-sided latches, even in uncontended workloads. In fact, if not carefully implemented, the performance for an uncontended workload can be as dismal as that for a highly contended one. To this end, this paper proposes design principles to mitigate those pitfalls and presents several optimizations that improve the performance of a well-known disaggregated RDMA-optimized DBMS [54] by 2×.

Correctness is hard. Achieving high performance in synchronization is unquestionably valuable, but ensuring correctness is mandatory. We have discovered that early techniques proposed in the literature fail to accurately synchronize concurrent operations, potentially resulting in hard-to-detect data inconsistencies. For example, consider an optimistic synchronization scheme as implemented in [32] (shown in Figure 1(a)), which presumes that RDMA operations occur in ascending address order — a common assumption in many papers. This scheme implements an update by writing the head version first, then the data modification, and eventually updating the tail version. Under the assumption of operations being performed in increasing address order, a concurrent reader can detect concurrent modifications by comparing the head and tail versions.

Incorrect assumptions. Unfortunately, in contrast to the general assumption in many papers [32, 47], RDMA reads are not guaranteed to be performed in increasing address order. In fact, the RDMA specification does not state the ordering within a single RDMA read. As a result, in the previously mentioned synchronization scheme, an RDMA read may first read both versions (i.e., the first and third cache line) and then retrieves the data from the second cache line. At the same time, a writer

[1]In this paper, *synchronization techniques* refer to the low-level synchronization mechanism, i.e., latching, not higher-level concurrency control.

concurrently modifies the data in address order. The concurrent data update may not be detected because the versions were read first before the concurrent writer started. However, the reader and the writer overlapped at the second cache line leading to inconsistent data.

We validate the existence of this behavior with a simple experiment consisting of one storage node and two compute nodes. A single-threaded remote writer on one compute node repeatedly fills a block of its local memory, e.g., 512 bytes, with the same 8-byte version number and then writes it to a remote buffer (50 MB) on the storage node with a single RDMA write. The version number is incremented on each iteration, and the new block is written to the next slot in the buffer. Concurrently a reader on the other compute node reads a block in the remote buffer with a single RDMA read and then checks whether the header and footer version numbers are identical. If the header and footer version numbers match, then the intermediate values are examined to determine if an inconsistency exists. "Torn" reads – having an identical header and footer version but inconsistent intermediate values – are *undetectable* by a validation scheme that checks the block's leading and trailing version numbers.

Figure 1(b) shows the percentage of how many such torn reads are undetectable due to changes in the read order. Note that no torn read appears with 128 bytes as only the header and footer cache lines are read, i.e., if they are inconsistent, this can be detected. However, inconsistencies happen for more than 128 bytes, and while not frequently, often enough to corrupt the data. Surprisingly, this problem is not widely known, and techniques that assume ordering are still very popular [47]. We believe the main reason for this assumption is that a single RDMA request requires many protocols — not only RDMA but also PCIe, and cache coherence — to work in concert. Thus, it is challenging to understand which guarantees are provided by the respective specification.

Contribution. The primary goal of this paper is to distill general principles for correct one-sided synchronization techniques. To our knowledge, this paper is the first principled analysis comparing the performance, scalability, and correctness of one-sided synchronization techniques. Our work demonstrates that understanding the specification and low-level hardware details is crucial for correct and efficient synchronization. Our underlying goal is to provide researchers and developers with guidelines on how and when to use the different synchronization techniques. Finally, we open-sourced benchmarks[2] that help to transfer our findings to different hardware setups and future developments.

2 BACKGROUND AND METHODOLOGY

While many systems [8, 31, 32, 60] implement various synchronization schemes, they do not isolate the impact of their synchronization on the overall performance. However, as we will show, synchronization techniques significantly affect scalability and system performance. This section describes the necessary background on RDMA and the typical RDMA hardware stack, existing synchronization techniques, and our experiment methodology.

2.1 Remote Direct Memory Access

Remote Direct Memory Access (RDMA) is a networking protocol that provides high bandwidth and low latency access to a remote node's main memory [36], using zero-copy transfers from the application space. Several RDMA implementations are available – most notably InfiniBand [15], RDMA over Converged Ethernet [3] (RoCE) [16], and iWARP [36]. RDMA offers four transport types: reliable or unreliable, which can be connected or unconnected. This paper focuses on the *reliable connected* transport as it is the only configuration that fully supports one-sided primitives.

[2]https://github.com/DataManagementLab/RDMA_synchronization
[3]RoCE is an attempt to combine RDMA with Ethernet. Refer to [12] for the shortcomings of RoCE in modern datacenters.

RDMA one-sided verbs. RDMA implementations provide two communication paradigms (called verbs) (1) *two-sided* and (2) *one-sided* verbs. Two-sided verbs are similar to traditional socket-based programming in that both sides (sender and receiver) are involved. In contrast, one-sided verbs (read, write, and atomics) provide remote memory access semantics, in which a process specifies the memory address of the remote node that should be accessed. The CPU of the remote node is not actively involved in the data transfer, i.e., only one side is involved. In this paper, we solely focus on one-sided primitives. RDMA read and write enable applications to read and write remote memory directly without remote CPU involvement. To support highly concurrent applications, RDMA specifications provide atomic operations [15, 39]. One-sided compare-and-swap (CAS) and fetch-and-add (FAA) operations atomically read, modify, and write memory at a remote destination. Those operations work similarly to local CPU CAS and FAA instructions. CAS atomically swaps the current value with a new one if it equals the expected value. FAA increments the current value with some user-defined value and then returns the original value to the caller. RDMA atomics are limited to 64-bit, 8-byte aligned values.

RDMA two-sided verbs. Two-sided verbs are widely used in high-performance RDMA systems [18, 20, 52] to send messages between distributed processes. The sender issues an RDMA send request, which consumes a waiting RDMA receive request at the destination. The receiving process issues the receive request to dictate the destination address for the sent payload. After the payload is written, the receiving process is notified of its arrival. Since the receiving process is explicitly involved in the communication, synchronization of remote processes is managed using request handlers and traditional multiprocessing approaches. In contrast, disaggregated memory systems leveraging one-sided verbs must synchronize processes directly through remote memory access, which requires careful consideration due to the behaviors we highlight in this paper. In addition, two-sided RDMA requires explicit processing on the storage side, which is typically not ideal with limited computational resources on the storage servers. Therefore, two-sided RDMA is not in the scope of this paper; however, several works do explore the relative merits of one-sided vs. two-sided RDMA [18–20, 52, 60]

Interface. Modern network interfaces typically provide asynchronous networking. This means that a network operation is dispatched to the NIC, which notifies the application once the operation is completed. RDMA's interface uses so-called send/receive queues to post operations: (1) While a *send queue* is used by the requester to issue operations such as read, write, send as well as atomics (2) the *receive queue* is used by the responder to issue receive requests. With RDMA, a connection between a requester and a responder bundles these two queues and is therefore called queue pair (QP). More precisely, the application must create queue pairs on both ends and connect them to initiate a connection between a requester and a responder. To issue RDMA operations, a host creates a *work queue element* (WQE). The WQE specifies parameters such as the verb and other metadata (e.g., the remote target address). The requester then adds the WQE to its send queue and informs the local RDMA NIC (RNIC) via Programmed IO (PIO) to process the WQE. For a signaled WQE, the local RNIC pushes a completion event into the *completion queue* (CQ) via a DMA WRITE once the remote side has processed the WQE. To enforce synchronous communication, the application can also block until the network card generates the completion event. Typically, RDMA is used asynchronously, meaning multiple WQE are simultaneously registered with the RNIC, and later the application checks for completions in the CQ. This technique is often referred to as *doorbell batching*.

2.2 RDMA Hardware Stack

As mentioned in Section 1, synchronization techniques rely on certain hardware guarantees; thus, we briefly present the required hardware in Figure 2. An RNIC connects via the PCIe bus to the

memory controller hub (MCH) of a multi-core CPU [35]. The MCH is responsible for handling memory access by both the CPU and peripheral devices. Among others, the MCH contains the memory controller, the coherency engine, and acts as the root complex for the PCIe bus. Because modern server architectures implement cache coherent I/O, e.g., Intel DDIO [17] and ARM CCI [1], knowing that an RDMA access to system memory will snoop (lookup) CPU cache for conflicting addresses is essential. If a conflicting address is found, it can be served from the CPU cache; otherwise, it is fetched from main memory. Cache coherent I/O is a core enabler for many of the techniques described in Section 4.

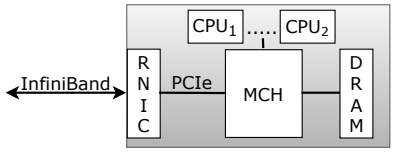

Fig. 2. Hardware components involved in RDMA

	Variant	Ops.	Systems	Correct
Pess.	Reader/Writer	2	[3], [53], [54]	Yes
	Exclusive	2	[46], [5],[33]	Yes
Opt.	CRC	2	[41],[61],[31],[45]	Prob.
	Versioning	2	[60], [32], [47]	No†
	Cache line ver.	2	[8], [51]	Yes

† can be fixed with an additional RDMA read (see Section 4.3)

Table 1. Pessimistic and optimistic techniques for one-sided RDMA synchronization.

2.3 Existing Synchronization Techniques

The one-sided synchronization techniques of modern distributed systems can be categorized into pessimistic and optimistic approaches, as shown in Table 1. Pessimistic approaches mirror traditional latching with the distinction of having only a single latch mode or two for reader-writer support; meanwhile, optimistic approaches can be further subdivided. Each optimistic technique offers different characteristics, like memory and computation overhead, but all embed metadata in the object, which is updated by writers and validated by readers. We now briefly introduce existing optimistic techniques from Table 1: *CRC*, *versioning*, and *cache line versioning* before we discuss them in more detail in Section 4.

Checksums. A common detection technique used to identify potential inconsistencies [31, 41, 61]. After updating the object, writers write back a new checksum (typically a 64-bit CRC). Readers then recompute a checksum locally and compare it to the observed checksum. While effective in practice, there is still a non-zero probability of a collision causing validation to succeed on corrupted data.

Versioning. Versioning is a strategy that augments the data with a single version. During execution, writers increment the version before and after updating the data. Readers read the version before and after their operation to validate whether an update occurred concurrently. Optimistic lock coupling is a special case of versioning in which sequence locks are used for both write-write synchronization and validating reads (e.g., [23, 60]). As we will explain in Section 4.3, there are pitfalls with this approach that impact its correctness.

Cache line versioning. Finally, cache line versioning maintains the object's version in each constituent cache line [8]. Writers increment all versions before updating the data, then increment them upon completion. Unlike versioning, readers detect inconsistencies with a single remote read by validating that all versions match.

To understand the differences between the sundry approaches, we perform a principled analysis and evaluation of one-sided synchronization techniques. To the best of our knowledge, this paper is the first to do so.

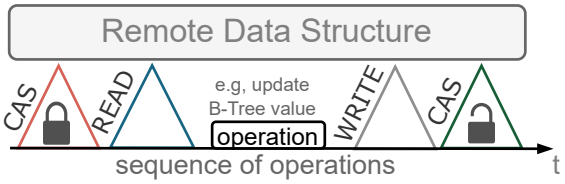

Fig. 3. Example of an exclusive latch acquisition.

2.4 Evaluation Methodology

Framework. We implemented all techniques in the same code base to isolate the fundamental properties of the synchronization schemes from incidental differences. We highlight each synchronization scheme's overhead using a perfectly-sized remote hash table to avoid hash collisions. In addition to the hash table, we use a remote B-Tree to show how the synchronization schemes behave when contention is inherent to the data structure due to its hierarchical nature (all workers start their traversal at the root node). We use multi-threading and execute one operation on a single thread (worker) until completion, i.e., no batching or asynchronous execution. This allows a fair comparison of the approaches.

Setup. We conducted all reported experiments on an 8-node cluster running Ubuntu 18.04.1 LTS, with a Linux 4.15.0 kernel. All nodes are connected to a SB7890 InfiniBand switch using one Mellanox ConnectX-5 MT27800 RNIC (InfiniBand EDR 4x, 100 Gbps) per node. Each node has two Intel Xeon Gold 5120 CPUs (14 cores) and 512 GB main-memory split between sockets.

Since the ConnectX-5 is connected to one NUMA socket, we inevitably have NUMA effects when using more than 14 cores (i.e., 28 threads). We allocate the memory on the socket of the NIC, and to alleviate the NUMA effects, we assign threads round-robin (interleaved) to both sockets. More details of NUMA effects on RDMA can be found in [34, 37]. Besides ConnectX-5, we also validate our results for different generations of RNIC, namely, ConnectX-3 and ConnectX-6. Mellanox is widely used; from 30 papers, we recently analyzed 28 used Mellanox cards. Besides on-premises, only Microsoft offers RDMA of the top three cloud providers and they use Mellanox cards in their instances. [21]. That being said, we believe that most findings are independent of the network card. The correctness discussion mainly depends on the protocol specifications underpinning the RDMA communication. In particular, our testbed leverages the widely-used InfiniBand specification [15] which shares commonalities with alternative RDMA protocols and make our results applicable to a broader range of deployments. The performance considerations shed light on possible pitfalls worth investigating when building an RDMA application. We open-source our benchmarks to help developers uncover performance and correctness bottlenecks in other RDMA system configurations.

3 PESSIMISTIC SYNCHRONIZATION

In order to prevent concurrent modifications of remote data structures such as a B-tree or a hash table, one-sided pessimistic synchronization techniques implement latches using one-sided RDMA operations. In this section, we present the basic implementation of such a latch and use it as a running example to discuss possible optimizations. Since RDMA atomics are the fundamental building block of these one-sided latches, we will then drill down into the performance and scalability characteristics of these RDMA primitives. Afterward, we outline and evaluate possible optimizations for one-sided latches.

3.1 Running Example of a Pessimistic Latch

Table 1 shows that existing pessimistic schemes are subdivided into two types of latches: While some latches such as RCC-NOWAIT [46] support only one latch mode (latched and unlatched),

Fig. 4. Scalability of contended and uncontended atomic RDMA operations when increasing the number of workers (4 compute nodes and 1 storage node)

others distinguish between shared and exclusive modes, i.e., reader/writer latches. Both latch types can be implemented with atomic RDMA operations. We will use a reader/writer latch for the running example, but all optimizations generalize to both latch types.

We implement a typical reader/writer latch using an 8-byte value [28, 53]. A worker uses an RDMA compare-and-swap (CAS) operation to set the latch bit (usually the trailing bit) for exclusive access. Readers increment the reader count, encoded in the remaining bits, with a fetch-and-add (FAA). In the basic variant, all operations are executed synchronously, i.e., the worker blocks after every operation until its result is returned.

Figure 3 gives an intuition on how this latch is used to access a remote data structure exlusively. First, the remote data structure is latched with an RDMA CAS operation on the 8-byte latch by setting the lock bit. Afterward, the desired data is read from the data structure, modified locally, and written back with an RDMA write operation. Finally, the remote data is unlatched with another RDMA CAS operation on the 8-byte latch.

To access the remote data structure in shared mode, the clients use RDMA FAA to increment the reader count of the latch speculatively. The return value (i.e., the full 8-byte) of the operation allows the worker to check if the latch is in the exclusive mode, in which case the worker decrements the reader count and retries. Otherwise, the worker reads the data using an RDMA read and then decrements the counter to unlatch.

3.2 Performance of RDMA Atomics

Because every pessimistic approach relies on RDMA atomics (CAS and FAA), it is important to understand their isolated performance before discussing how our basic latch can be optimized.

Uncontended vs. contended RDMA atomics. In the first experiment, we examine the scalability behavior of contended and uncontended RDMA atomics. Both scenarios are equally vital, and while heavy contention is typically rare, it is unavoidable in some workloads, e.g., having hot tuples. To show the effect of contention, we perform an experiment in which all workers issue an RDMA CAS operation of the same 8-byte atomic counter. To understand how uncontended atomics scale, we assign a private 8-byte remote atomic counter to each worker. For reference, we compare uncontended atomics to the performance of an RDMA read.

Figure 4 shows the scalability behavior of both uncontended and contended RDMA atomic operations when increasing the number of workers. As we can observe, uncontended atomics scale like one-sided RDMA read operations, peaking at around 51.2 million operations per second with 128 workers. This suggests that the parallel uncontended atomic requests do not interfere, but we will demonstrate later that this does not hold for all scenarios. Note that the performance drop of uncontended atomics at 512 workers has nothing to do with the atomic operation but can be attributed to QP-thrashing [8] on the client machines. QP-thrashing means that the QP state

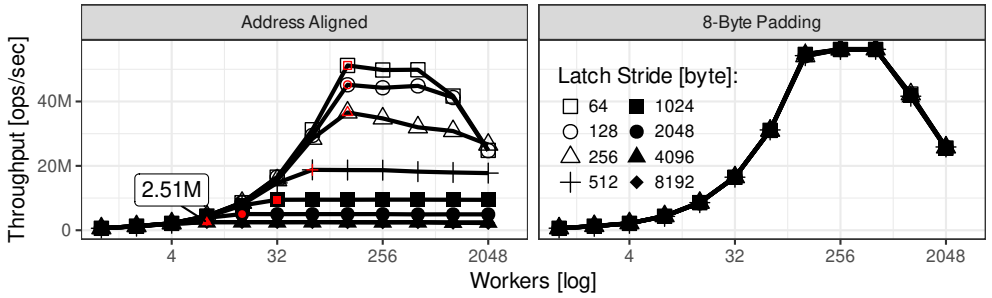

Fig. 5. Scalability of uncontended atomics with varying strides between the latches (4 compute nodes and 1 storage node).

cannot be cached on the RNIC and is swapped in and out to host memory. This occurs at around 128 utilized parallel QPs per client NIC in our hardware, i.e., 512 workers.

As expected, when running the contended atomic workload, the peak is significantly lower at 2.32 million operations per second and atomic operations only scale until 8 workers.

In the literature, RDMA atomics seem to have a bad reputation for being fundamentally unscalable [19] – even for uncontended workloads. However, the above experiments demonstrate that uncontended atomics scale well w.r.t. the number of workers. In fact, they show comparable scalability to RDMA read operations. While this experiment offers valuable insights into the scalability of atomics, it is not the complete picture, as we will demonstrate.

Atomic stride size and alignment. Obviously, the scalability of RDMA atomics depends on the degree of parallelism. However, subtle details such as the data alignment can also interfere with the scalability. So far, in our experiments, values are placed in a 64-byte stride, i.e., a cache line. We only used the first 8 bytes for the atomic counter and ignored the remaining 56 bytes. However, in practice, RDMA atomics are placed at larger strides as they protect data of various sizes, e.g., a 4 KB B-Tree node.

In the following experiment, we measure the effect of larger stride sizes by varying the distance between the atomic counters. As in the previous uncontended experiment, each worker has a private latch to avoid contention. Consequently, the expected outcome should be similar to the uncontended result in Figure 4. Surprisingly, Figure 5 shows that the stride size impairs the scalability tremendously. That is, with larger stride sizes, the inflection points w.r.t. throughput (highlighted in red) are reached earlier. With a 64-byte stride, the peak performance is 50M operations with 128 workers, i.e., the same upper-bound as in Figure 4. With a 256-byte stride, the peak performance is at 40M operations with 128 workers. With a 512-byte stride, the peak performance is halved and reached with fewer workers (20M operations with 64 workers). Remember that there is no true latch contention, and we only vary the distance between the atomic counters and nothing else. The observed behavior must be based on a physical contention point in the RNIC.

NIC internals - physical contention. Through reverse engineering, we believe that the RNIC uses an internal locking table to serialize atomic operations. Since the locking table works similarly to a hash table, collisions can happen. Unrelated atomic operations can be assigned to the same slot, severely limiting the throughput of concurrent uncontended atomic operations. The lock slots are determined based on the last 12 bits of the atomic operation's target address. For instance, if we use a 4 KB-aligned address as illustrated in Figure 6(b) i), the last 12 bits of the address are zeros, and they are assigned to the same lock slot. Even though the atomics do not contend on the same latch, the way the RNIC handles atomics results in physical contention inside the locking table, as exemplified by the arrows in Figure 6(a). The two CAS operations target different latches

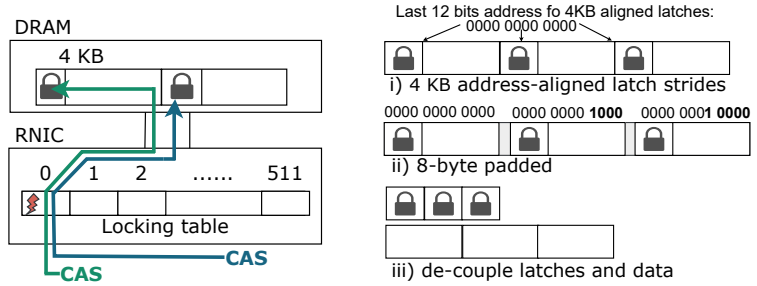

(a) RNIC Internal Locking Table (b) Stride Size and Alignment

Fig. 6. Pessimistic synchronization performance can be impacted by RNIC architecture and by host memory layout.

and are still serialized in the same lock slot, negatively impacting performance. Consequently, logically uncontended atomic operations do not scale very well when the data alignment is not carefully chosen, as shown in Figure 5. When we compare the results of our initial contended scalability experiment from Figure 4 with the performance of 4 KB-aligned stride sizes, we can see that the performance and scalability characteristics are very similar. Given the underlying hardware mechanism, this performance is now explainable: the operations are serialized in the same lock slot, whether through a collision of the address or the operations targeting the same latch. We validated the existence of a performance signature matching our hypothesized 12-bit lock table in three RNIC generations: ConnectX-3, ConnectX-5, and ConnectX-6 RNICs. Also, Kalia et al. [19] observed similar findings for an older network card (Connect-IB). However, it is hard to generalize our findings to all NICs since the implementation details are not made publicly available by the RNIC vendors. Therefore, like other papers, we can only infer implementation details [19, 47]. Instead, we emphasize that NIC hardware details are essential and demonstrate potential bottlenecks that should be carefully evaluated when building a high-performance system.

Improving scalability of uncontented atomics. To improve the scalability of uncontended operations, we must avoid collisions in the locking table. The only way one can control this is through the data layout, i.e., the addresses of our latches. The goal is to place the latches so that the last 12 bits (used for the lock-slot calculation) are different. Consider the example in Figure 6(b) ii); instead of allocating the 4 KB blocks back-to-back, a padding of 8-byte is placed before the latches. Now, the last 12 bits of the latch addresses are not all zeros and, thus, will be assigned to different lock slots. The effect of this mitigation technique is illustrated in Figure 5. We can observe that all stride-sizes scale equally well (all measurements happen to be on the same line). Another possibility to better utilize the lock-slots is to decouple latches from the data as depicted in Figure 6(b) iii). In this technique, the latches are placed back-to-back. Since the latches are only 8 bytes, the last 12 bits of the latch address will vary and achieve the same scalability behavior as for the 8-byte padding. Note, in this experiment, we test the performance of the atomic operations, i.e., we do not read the data. This allows us to isolate the atomic contention effect, but in practice separating the latch from the data may have adverse effects due to locality. In particular, the translation from physical-to-virtual memory could suffer as every data access invokes two translations, i.e., one for the latch and one for the data. However, this depends on other parameters, such as the working set, tuple size, and NIC-cache size [19], and thus requires a holistic evaluation.

To conclude, despite contrasting beliefs, RDMA atomics can scale well w.r.t. the number of workers. However, the scalability depends on the number of concurrent accesses (contention) and the data alignment. While the first is workload-dependent, the latter can be carefully tuned to best

utilize the internal RNIC hardware. While we demonstrate that padding is beneficial for RDMA atomics, it can also have consequences elsewhere, such as making page table lookup less efficient. Therefore, we argue that it is crucial to understand the internals of the RNIC and holistically optimize the DBMS system.

3.3 Optimized Pessimistic Latches

Equipped with our findings on how to utilize atomics optimally, we can now focus on how to design optimized pessimistic latches. Recall that the basic latch variant executes all operations synchronously, as illustrated in Figure 3. After every operation, the completion is awaited by polling the completion queue (cf. Section 2). While this is certainly correct, it is inefficient and increases the per-operation latency.

Latch optimizations. To reduce the operation latency, many systems use optimizations. Unfortunately, those optimizations are often only briefly discussed by the authors. We compiled a list of existing optimizations following numerous small hints in related work and a careful study of published source code. This is the first paper that describes these optimizations in detail and discusses why those optimizations guarantee correctness. In the following, we present those optimizations and highlight possible pitfalls.

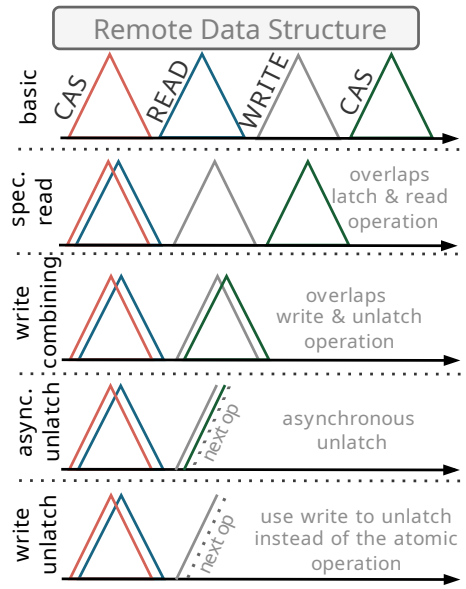

Fig. 7. Evolution of exclusive latch optimizations.

The *speculative read* optimization overlaps the latch with the read operation (cf. Figure 7) to hide the latency of the read. If the latch request is successful, we can proceed, and if it is unsuccessful, the latch request is restarted. However, the correctness of this optimization relies on the order of the operations. That is, it must be guaranteed that the latch operation takes place before the read happens. Fortunately, RDMA atomics are executed prior to subsequent RDMA operations on the same QP as per the InfiniBand specification [15].

Similarly, *write combining* overlaps the write and the unlatch operation (CAS) as illustrated in Figure 7. The intuition is that the unlatch operation can mask the write latency. This optimization is only applicable for exclusive access since read-only operations do not update the data and thus do not involve an RDMA write.

The *asynchronous unlatch* optimization is an optimization that goes even further and does not wait for the last completion event synchronously and thus immediately processes the next operation. In contrast to the synchronous variants, however, the memory buffer containing the write's payload cannot be reused immediately for the next operation. The issue is that when immediately re-using the buffer, the data from the previous operation could be overwritten as the previous RDMA write may still be outstanding; as such, we need to ensure that the operations are finished before re-using the buffers. The typical solution to avoid overwriting in-flight buffers is to use multiple buffers per QP. For instance, if the worker wants to modify two tuples, then the first tuple is written to the first (local) buffer. Subsequently, the remote content of the first tuple is updated and asynchronously

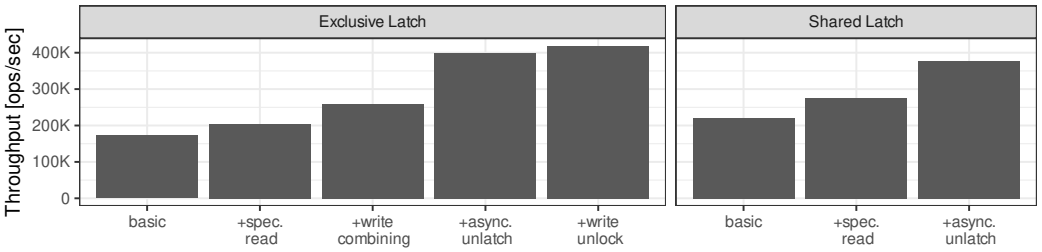

Fig. 8. Single-threaded ablation study of latch optimizations (1 compute node and 1 storage node).

unlatched. Since the remote operations are executed asynchronously, the RDMA operations (CAS and write) may still be outstanding, and the first buffer should not be re-used for the second tuple. Therefore, the second tuple is is written to a second (local) buffer. Using separate buffers gives the outstanding operation from the first buffer time to complete without any risks of overwriting the content. The first buffer can be safely re-used when the second operation on the same QP generates a completion event, i.e., after the payload is read.

Write unlatch is an optimization that relies on the fact that RDMA writes are performed in increasing address order. This optimization often works in practice but is similar to RDMA reads, not specified by the RDMA standard. However, because many essential applications such as MPI [26] and many other systems rely on last-bit polling [8, 11, 29, 59], RNIC vendors typically implement RDMA writes in increasing address order [8] for compatibility reasons. Note that the latch must be located at the highest address (typically as a footer) to protect the data until the full write has been completed. For example, the last 8 bytes of the value of an item residing in an RDMA-enabled key-value store could encode the lock protecting it. The payload of an RDMA write to update the value would be suffixed by an unlocked value to also release the lock. Since this optimization uses the write to unlatch the data, it saves a complete atomic operation. Unfortunately, there is a catch; the optimization only works in *certain cases*. The InfiniBand specification makes no guarantees that non-atomic and atomic RDMA operations are ordered and visible to each other when issued from different QPs [15]. That means that the RDMA write that unlatches the data item may not be visible to subsequent atomic operations from other workers. This can lead to behavior that might seem surprising because sequential consistency is not guaranteed.

We believe this is because the RNIC can buffer atomic operations for fast completion in a special on-chip buffer. When an atomic operation arrives at the RNIC, the cache line in which the value is stored, is read into this buffer from the memory. But because there is no guarantee w.r.t. interference of non-atomic operations, it may happen that once the atomic operation reads the current value in the buffer, an RDMA write changes the value on the cache line, i.e., unlatch. To make this more concrete, assume the data item is exclusively latched. Now, an atomic FAA operation wants to increment the reader count. The value is read to the RNIC buffer, but in the meantime, the latch-holder unlatches the data item with an RDMA write. Unfortunately, there are now two incoherent states: (1) in the RNIC buffer, the latch is still latched (2) in the cache line, the latch is unlatched. The RNIC increments the old (latched) value and overwrites the unlatch state on the cache line. Therefore, the original unlatch operation is lost, and the latch will remain latched, leading to a deadlock. Hence, the write unlatch optimization cannot be used with FAA operations and only works in combination with CAS operations, as used in RCC-NOWAIT [46]. The reason is that CAS operations conditionally overwrite the state. The operation detects that the old state is still latched and does not modify the latch. Thus, a concurrent worker may detect the unlatch operation delayed for the incoherent states, but eventually, the cache coherence protocol ensures that the RNIC sees the newest version. Consequently, creating a reader/writer latch in combination with the

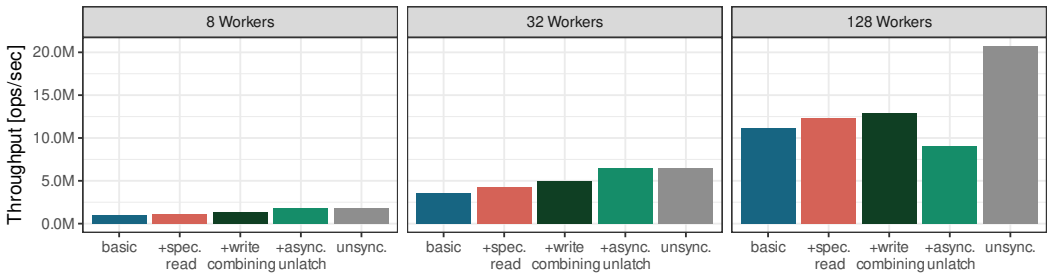

Fig. 9. Multi-threaded performance of optimized latches (4 compute nodes and 1 storage node).

write unlatch optimization is impratical and only one latch mode can be supported (no distinction between reader and writer) (cf. Table 1 RCC-NOWAIT).

3.4 Ablation Study of Latch Optimizations

To better understand how the latch optimizations perform, we first enable them step-by-step in a single-threaded ablation study. We show the throughput for exclusive and shared latch acquisitions in Figure 8. The numbers include all necessary operations: *CAS, read, write, and CAS* for a modification, and *FAA, read, FAA* for a read operation. As mentioned, some optimizations are only applicable for the exclusive latch acquisition. For the analysis, we repeatedly latch and perform the operations on a 256 byte-sized tuple placed on a storage node.

We can see that all exclusive latch optimizations in Figure 8 (left) increase the single-threaded performance. One of the most effective optimizations is asynchronous unlatch since we can immediately start the next operation. In contrast, write unlatch does not significantly increase the performance further. Combined with the fact that it is impractical to implement a reader/writer latch, we will not further consider this optimization in this section. However, we will leverage this optimization again in Section 4.

When focusing on the shared latch performance in Figure 8 (right), we can observe that the basic performance is higher since the read consists of only three sequential operations. The advantage of fewer operations vanishes with the higher optimization levels as they aggressively overlap messages (cf. Figure 7). Eventually, the performance of exclusive and shared latch acquisitions converges with higher optimizations and becomes latency bound.

Scalability of latches. So far, we have focused on the single-threaded performance for pessimistic latch acquisitions – however, most disaggregated database systems run with multiple workers dispersed across several compute nodes. Thus in the following, we discuss how the optimizations perform with an increasing number of workers equally distributed across 4 compute nodes. We use a data set with 20 thousand 256 byte-sized tuples stored on one storage node. We also measured with 20 million tuples, but the results behaved similarly. The accesses are randomly distributed across the data set.

We only show the performance of the write-only workload, i.e., the exclusive latch optimizations, since the read-only workload behaves very similarly in this experiment. As the upper bound, we include an unsynchronized version (an RDMA read and RDMA write). Figure 9 shows the results for 8, 32, and 128 workers. The first observation is that the performance is on par with the unsynchronized version when all optimizations are enabled, i.e., *asynchronous unlatch*, for 8 and 32 workers. The optimizations effectively hide the latch operations with the data movement operations. Furthermore, the performance scales near-linear when increasing the worker count from 8 to 32. However, with 128 workers, the highest optimization level seems to be counterproductive. We attribute this behavior to the fact that the next operation may start sooner. Acquiring the next

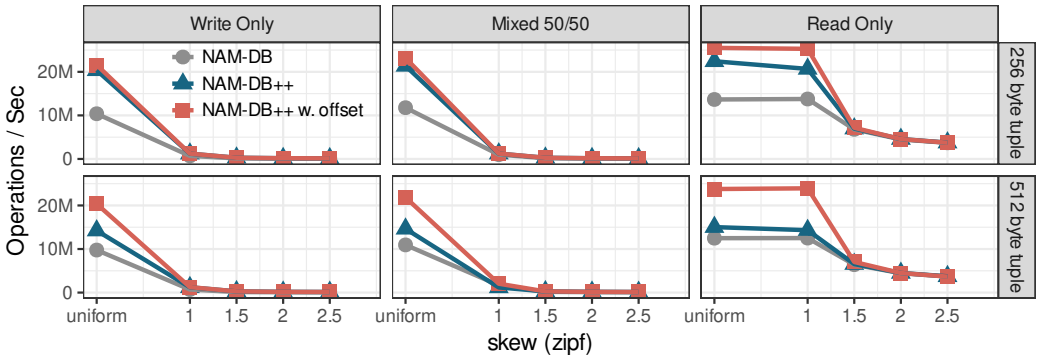

Fig. 10. Effect of optimizations in NAM-DB (4 compute nodes and 4 storage nodes).

lock earlier introduces additional lock contention in the presence of more workers. For instance, one worker already acquires the next lock even though the old lock is still latched due to the asynchronous nature, i.e., a worker can hold two locks for a limited period. It can also lead to increased RNIC contention on the storage node, with more in-flight operations. In practice, the *asynchronous unlatch* optimization is still worthwhile since there are typically more storage nodes, as we will see in the following experiment.

3.5 Effect on a Disaggregated DBMS

As stated earlier, latch optimizations improve performance tremendously. However, the previous results were only obtained in micro-benchmarks. Let us substantiate that claim by integrating them into an existing disaggregated DBMS. We use NAM-DB [54], which is a disaggregated RDMA-optimized DBMS. In our setup, we use 4 compute nodes with 28 workers each (112 total workers) and 4 storage nodes. Among those storage nodes, we distribute 20 million tuples. We measure the throughput in operations per second for representative tuple sizes of 256 and 512 byte. We implemented the highest optimization level., i.e., *asynchronous unlatch*, and call this version NAM-DB++. In addition, we show the effect of manipulating the latches' data layout, as we discussed in Section 3.2. We compare the original NAM-DB with NAM-DB++ in a write-only, mixed, and read-only workload. To show the effect of contention, we vary the access skew.

Read-only performance.. Let us first focus on the 256 byte-sized tuples and the read-only workload (top right in Figure 10). The optimizations double the performance with uniform and slightly skewed (Zipf 1) access patterns. When the contention increases, the hardware limits the performance, and both systems converge. The dramatic performance degradation is nevertheless surprising in the read-only scenario. In theory, the read-only performance should not dramatically collapse since multiple readers can acquire a latch simultaneously. Unfortunately, the RNIC cannot handle the concurrent atomic RDMA operations necessary to acquire the reader latch in the first place. As mentioned in the previous experiment, the RDMA atomics are serialized in the RNIC, which does not scale well when the lock slot is contended. Despite the hardware limitations, the read-only workload still performs much better than the write-only workload under high contention. For instance, with a Zipf factor of 2 the write-only performance is $170K$, whereas the read-only performance is $4.5M$.

Write performance. The contention exacerbates the performance issue in the write-only and mixed workload. The locks now logically contend in addition to the physical contention on the RNIC. In other words, the performance of the atomic operations decreases since those operations often target the same latch and are serialized in the same RNIC lock slot. Once the RNIC processes

the atomic operation, the latch may have been already latched exclusively, which requires a restart and aggravates the problem.

Larger tuples. When looking at the 512 bytes-sized tuples in the read-only workload, we can see that the effect of applying the padding is more pronounced. As mentioned earlier, the larger latch strides lead to more contention inside the RNIC's lock-table, thus limiting the performance. We use 4 storage nodes, as opposed to the experiment in Section 3.2, and still, the collisions inside the RNICs become the primary bottleneck. Therefore, with the padding optimization, the performance significantly improves by removing the bottleneck of physical contention, allowing the latch optimizations to achieve their potential.

4 OPTIMISTIC SYNCHRONIZATION

In the previous section, we have seen that pessimistic synchronization scales when the latches are uncontended. However, in some data structures, latch contention is unavoidable and, in fact, inherent to the data structure design. For instance, B-Tree operations have to traverse the B-Tree root node. Although the root node is mainly latched in shared mode, this creates (physical) contention on the RNIC when using pessimistic synchronization, negatively impacting the performance. Figure 11 shows this effect for a B-Tree with 4 KB nodes stored on a single storage machine and accessed read-only from 4 client machines. As we can observe, the unsynchronized B-Tree can saturate the bandwidth with only 16 workers, whereas the pessimistic synchronized version stagnates much earlier and never achieves the full bandwidth. Note that the pessimistically synchronized version includes all the optimizations presented above, and even then, the performance is disappointing in this experiment.

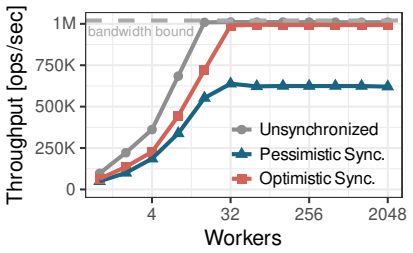

Fig. 11. Lookups in one-sided B-Tree (4 compute nodes and 1 storage node).

This is why many papers eschew RDMA atomics and propose an optimistic synchronization scheme in which reads do not physically latch the data but detect concurrent modifications. In contrast to pessimistic synchronization, Figure 11 demonstrates that optimistic synchronization can achieve the full bandwidth. We start this section by providing an intuition on how optimistic synchronization works. After that, we discuss PCIe's ordering guarantees and its devastating effects on some optimistic schemes leading to an incorrect synchronization. We then present correct optimistic synchronization schemes andgt evaluate their performance.

4.1 Intuition for Optimistic Synchronization

Optimistic reads. The intuition for optimistic synchronization is that readers proceed optimistically and then validate, while writers physically acquire an underlying pessimistic lock to avoid write-write conflicts. To achieve the same guarantees as shared pessimistic latches[4] readers must check that the data item did not change during their operation. This is typically realized by augmenting the data item with a version counter and incrementing it upon each modification, which allows readers to detect concurrent writes. The layout of such an augmented optimistic lock is shown in Figure 12(a). It consists of a pessimistic latch used for exclusive access and the version counter. Using this version counter, readers validate that the version did not change during their operation. If the validation fails, the operation is restarted.

[4]We will discuss relaxed guarantees in Section 5

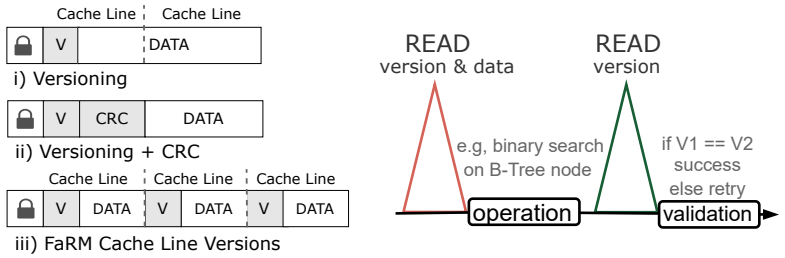

(a) Optimistic Layouts (b) Naïve Optimistic Technique

Fig. 12. Implementation of optimistic synchronization

Naïve implementation. One way of implementing a one-sided optimistic lock is depicted in Figure 12(b), which we call the *naïve implementation*. This approach uses a single RDMA read to copy the latch, version, and data to the worker. The worker can then check if the item is exclusively latched and possibly restart. Otherwise, if the check on the latch succeeds, the operation, e.g., a binary search in a B-Tree node, is performed optimistically. Once the operation is finished, the version is read once more (via RDMA) and compared to the initial value.

This after-the-fact validation is crucial to detect concurrent modifications and to get the same guarantees as a pessimistic latch. Thus, the validation is effectively equivalent to an unlatch operation. These semantics are critical for higher-level synchronization techniques such as optimistic lock-coupling [24, 60] on a B-Tree since we need to ensure that the B-Tree node did not split and the child node is still valid.

4.2 PCIe's Ordering Guarantees

Intermediate protocols. Unfortunately, the naïve implementation in Figure 12(b) is not correct. As pointed out by Taranov et al. [43] there are three factors that can affect data deliver order of transmitted messages: (1) Message ordering, (2) packet ordering within a single message, and (3) DMA ordering. The first two factors are generally ensured by InfiniBand and RoCE[5]. But even though message ordering is guaranteed, it is not guaranteed that DMA operations are performed in-order. The InfiniBand RDMA specification [15, 16] does not itself provide any ordering guarantees concerning the order of bytes read by an RDMA read. In other words, RDMA operations are not designed with concurrency in mind (except for RDMA atomics). However, in practice, many academic and industry-grade systems use RDMA concurrently on the same memory regions. Since RDMA does not specify the behavior of those concurrent accesses, the intermediate protocols are important. For example, the lack of order for reads permits intermediate protocols involved in the operation, e.g., PCIe, to retrieve host memory as they see fit. Therefore, to fully understand why an RDMA read may not execute in increasing address order, it is critical to investigate the impact of PCIe [35] and cache coherence protocol. Understanding the underlying protocols subsequently allows us to extract correct synchronization techniques.

As shown in Section 2, an RDMA read request is sent over the network to the remote node. The RNIC then dispatches the request to the PCIe controller, which fetches the requested data region from the host memory. The data is transferred via PCIe to the RNIC and then sent back to the requester in one or more RDMA packets. An important aspect is that PCIe requests serviced by the host are cache coherent on modern server architectures. Modern architectures provide *direct cache access* [14, 25, 42] to allow high-performance I/O such as RDMA to access processor caches

[5]There are RNICs that deviate from this and offer out-of-order delivery [6]

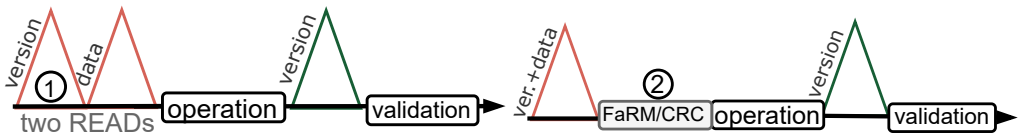

(a) Two RDMA reads (b) FaRM and CRC technique

Fig. 13. Correct optimistic techniques

directly. For example, the x86 machines in our testbed are equipped with Intel's DDIO [17], and a similar mechanism is available for ARM [2].

PCIe reordering. Since the actual data transfer from the remote memory to the remote RNIC is initiated and carried out via PCIe, we must look to the PCIe specification. RDMA requests are translated to PCIe transactions consisting of reads and writes that are processed by the so-called PCIe root complex. Hence, the guarantees provided at this layer of the hardware stack play a pivotal role. The root complex services PCIe read requests, which are coherent at a cache line granularity. Once a request is initiated, the PCIe protocol transfers that data to the endpoint via so-called *completions*.

Multiple completions are used for reads larger than a given size, e.g., 64 bytes. Herein lies the problem. The PCIe specification states, "Memory Read Request serviced with multiple completions, the completions will be returned in address order." [35]. This is hard to interpret, but it only determines the order of the completions and not in which order the data is retrieved from memory. In fact, an implementation note permits that a "single Memory Read that requests multiple cachelines from host memory is permitted to fetch multiple cachelines concurrently" [35].

Implications on correctness. Due to the concurrent cache line retrieval, we cannot reliably detect concurrent modifications with our naïve implementation from Figure 12(b). For example, assume a concurrent writer and a reader access a data item as depicted in Figure 12(a) i). With the lack of order, the reader could first retrieve the second cache line in which the writer currently modifies data. In the meantime, the writer increments the version and unlatches the data item again. Only then does the reader retrieve the first cache line containing the latch and the version counter. Consequently, despite the validation step at the end in which it will retrieve the version again, the reader will falsely assume that the version did not change. In our introduction, we have already shown that this is a real risk for modern RNICs (cf. Figure 1(b).) This may be surprising for two reasons: (1) The PCIe standard guarantees that RDMA writes are performed in increasing address order while reads can be read out of order, as discussed before. (2) many papers rely on this unspecified behavior. In fact, an earlier paper of ours [60] suffered from this wrong ordering assumption. More details on the memory semantics of RDMA can be found in [7].

Referring back to Figure 1(b), we can see that data is rarely corrupted, but it happens, and if data corruption occurs, it is almost undetectable. We tested this behavior on three generations of modern RNICs, including the RNICs available in the cloud, and observed that no RNIC provides additional ordering guarantees beyond the PCIe specification. Fortunately, there exist techniques that prevent this issue, but, as we will see in the following section, they are not for free and come with different performance characteristics and trade-offs.

4.3 Correct Optimistic Synchronization

Optimistic synchronization relies on the initial RDMA read to observe a consistent view of the version and data. Due to the lack of ordering at the PCIe bus, we must rely on an additional

mechanism to provide this capability. Only then the after-the-fact validation will correctly detect concurrent modifications. We will discuss existing mechanisms[6] in the following.

Versioning (using two RDMA reads). This technique is not very different from the naïve version; however, it requires two dedicated RDMA reads in the beginning, as shown in Figure 13 a). The first RDMA read targets only the latch and the version to ensure correct serialization, and only the second reads the data. Because the version is always read before the data, every concurrent update will be detected. When looking at Figure 13 one may wonder if the two reads can be overlapped similarly to the operations in Figure 7. Unfortunately, this is not possible with two RDMA reads since they could be re-ordered at the PCIe level or even in the network. This is due to another unexpected pitfall: Although RDMA operations are ordered in a connected queue pair, the ordering only holds for the RDMA completion events. There are some exceptions in which the order is ensured, e.g., for atomic operations. On the other hand, the two reads must be issued sequentially.

Besides the correct order of the two reads, the latch and version must be stored in one cache line to exploit the cache coherence protocol and read both consistently. Only this enables a reader to detect concurrent modifications reliably. Unfortunately, this technique needs two RDMA reads, which may be expensive. The following two approaches only require a single RDMA read by embedding additional metadata in the object as shown in Figure 12(a)

Checksum by CRC. This scheme detects inconsistencies by using a checksum in the data, e.g., CRC64, that allows the worker to verify the data with high probability. The CRC will not match the corrupt data if there is a concurrent writer. Therefore, in the best case, only one RDMA read is required (cf. Figure 13 b)). The downside, however, is that (1) the CRC generation is computationally expensive and (2) it is only probabilistic. Admittedly, the probability of a collision is low for CRC64. However, if a collision occurs, the data corruption is hard-to-detect.

FaRM cache line versions. As with CRC, FaRM [8, 9] proposes a technique that requires one RDMA read in the best case. However, instead of computing a checksum, FaRM relies on coherent cache DMA (specified by x86) to detect if a single RDMA is consistent. FaRM stores a version at the beginning of *every* cache line as illustrated in Figure 12(a). Therefore, we can detect if a concurrent modification happens by comparing all cache line versions. To enable this, writers first latch the data item, read it, modify the data, increment the version locally, and then write the data back via RDMA in address order. Although not as computationally expensive as CRC, every cache line must be accessed to validate the versions introducing additional cache misses, i.e., stalled CPU cycles. In addition, this technique imposes additional storage overhead to accommodate a version in every cache line.

4.4 Single-threaded Performance

We initially focus on single-threaded performance to understand the different trade-offs of the correct schemes. Figure 14(a) compares the single-threaded read-only performance with varying tuple sizes. We also include the "broken" optimistic scheme and the optimized (and correct) pessimistic scheme from Section 3.

Optimistic vs. pessimistic. Maybe unexpectedly, the pessimistic synchronization scheme performs better than all optimistic schemes, including the incorrect one. Although the broken optimistic scheme only requires 2 messages as opposed to the pessimistic scheme that requires 3 messages, the pessimistic scheme can exploit the asynchronous unlatch optimization. This optimization unlatches asynchronously, and the subsequent operation begins immediately (cf. Figure 7). Analogous behavior

[6]Note of caution: Some optimistic techniques rely on hardware details that may not hold. E.g., when data is not aligned or larger than the MTU.

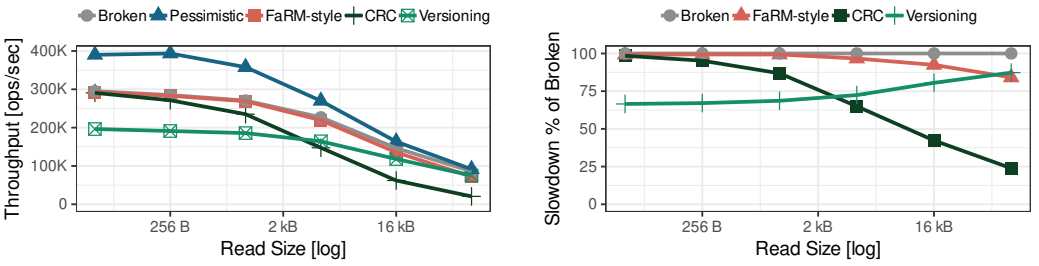

(a) Throughput comparison

(b) Relative slowdown of correct schemes

Fig. 14. Single-threaded performance of optimistic reads (1 compute node and 1 storage node).

is not possible for optimistic schemes as the validation, i.e., the unlatch operation, determines if the operation failed or succeeded. Consequently, the validation must be synchronous (see Figure 13).

Correct schemes vs. broken. Let us now quantify the induced overhead of the correct schemes compared to the broken scheme, which only consists of a single RDMA request and no consistency checks to acquire the data and version information (i.e., the first round trip in Figure 12(b)). From Figure 14(a), we can observe that the broken scheme performs better than the correct schemes. The closest competitor is the FaRM scheme, which performs nearly as well as the broken scheme and only drops with larger message sizes. CRC can only compete with small tuple sizes, contrary to the versioning scheme, which catches up with large tuple sizes.

To highlight those effects, Figure 14(b) shows the slowdown in percentage compared to the broken scheme. We can see that both CRC and FaRM suffer from larger tuple sizes. The computational overhead for both schemes is $O(n)$ since CRC calculates the checksum for every byte, and FaRM checks the versions in every cache line. Consequently, both schemes get linearly more expensive with increasing tuple sizes, more precisely, the consistency check (phase ②) as shown in Figure 13. However, because the CRC calculation is considerably more expensive than checking every cache line sequentially, the performance degrades more severely. For FaRM the sequential HW prefetcher and the out-of-order execution hide the cache misses. In contrast, the versioning approach incurs a substantial *constant* overhead by requiring an additional read for the version. Therefore, the versioning approach amortizes the initial cost with larger tuples sizes and performs better than FaRM.

To summarize, looking at the results of the single-threaded experiments, we found (1) the pessimistic schemes out-performs the optimistic schemes in the single-threaded setup, (2) FaRM performs better with smaller tuple sizes ($\leq$ 64 KB), (3) while versioning is beneficial with larger tuples.

4.5 Scalability of Optimistic Techniques

While the optimistically latched single-threaded performance was worse than the pessimistic scheme, it is not indicative of its scalability behavior. After all, the main benefit of optimistic schemes is that they avoid RDMA atomic operations during a read and, thus, avoid physical contention on the RNIC generated by multiple workers.

Read-only scalability. Figure 15 shows the scalability w.r.t. the number of workers for small, i.e., 256-byte sized, tuples. Again with fewer workers (8), the pessimistic approach performs better than the optimistic schemes. In line with the previous results, CRC and FaRM are close to the broken scheme, whereas the versioning scheme is far behind. However, under the highest contention when using 128 workers, all the optimistic schemes perform better than the pessimistic scheme.

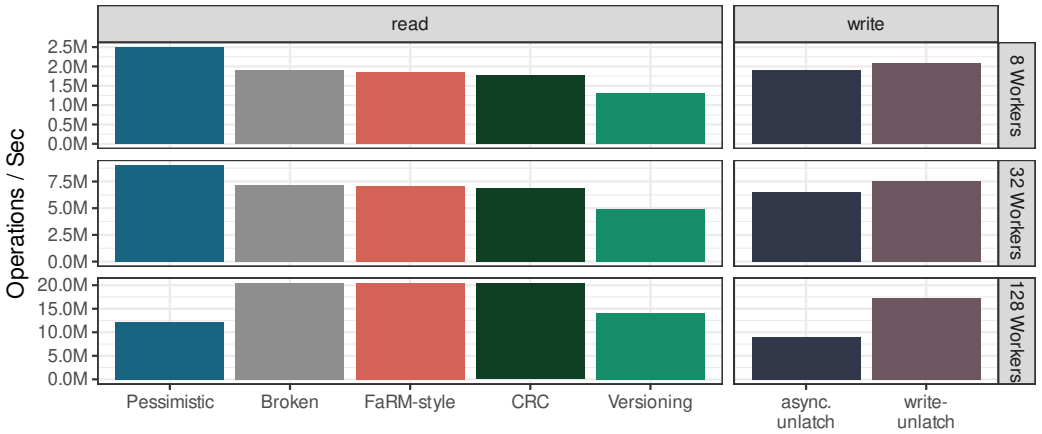

Fig. 15. Read-only and write-only optimistic scalability (4 compute nodes and 1 storage node)

In particular, FaRM and CRC perform almost twice as well as pessimistic synchronization. This confirms our intuition that by avoiding RDMA atomics, the physical contention effect does not limit the performance in optimistic schemes.

We also compared the scalability of larger, 16 KB sized tuples and found that the pessimistic scheme performs better in this scenario. The reason is that with larger tuples, the workload becomes network bound before the RNIC contention limits the performance. In other words, only 8 workers are required to reach the full bandwidth (12 GB/s) with the pessimistic synchronization, and the RNIC can easily sustain RDMA atomic operations at such a rate.

Write-only scalability. Remember that only reads are performed optimistically, while modifications are still latched pessimistically. However, since the optimistic readers do not use RDMA atomics, we can now leverage the write-unlatch optimization, which uses a write instead of an atomic operation to unlatch (see Figure 7). To remind ourselves, we have not considered them before because the write-unlatch optimization does not work with concurrent FAA operations required to implement reader-writer latches. But because optimistic schemes only require one latch mode, i.e., exclusively latched or unlatched, write-unlatch can be used to further reduce the number of RDMA atomic operations.

Figure 15 (right side) shows the effect of this optimization when increasing the number of workers up to 128. We can observe that the write-unlatch optimization enables better scalability by avoiding one RDMA atomic operation. When comparing the performance of the write-unlatch optimization with 128 worker and the results from Figure 9 we can see that it approaches the performance of the unsynchronized variant. Important to note is that the layout as shown in Figure 12(a) must be reversed so that the version and latch are placed at the end of the data to exploit address-ordered RDMA writes. That is, the last byte written must unlatch the record. The reversed layout does not affect the performance of the optimistic approaches since it merely changes the location of the version.

Qualitative discussion. So far, we have focused on performance numbers. However, the correct techniques come with trade-offs, such as the number of required messages. In the best case, the versioning approach requires 3 messages and can be a sub-optimal strategy when the workload is message-bound. In contrast, CRC requires only 2 messages but is computationally more expensive, which may be detrimental if the threads can do other valuable work. Moreover, CRC is probabilistic, and although the collision probability is low, the question remains: why risk data corruption when reliable schemes exist? Lastly, although FaRM is certainly efficient, it adds $2 - 8$ bytes storage

overhead [8] per cache line, which means it is not only $O(n)$ in terms of computing but also storage overhead. In addition, the higher-level logic must handle the interleaved cache line versions. For instance, any string operation (comparison, regexp etc.), e.g., on a value in a B-Tree, must explicitly deal with strings chunked across cache lines. Because only a few existing libraries support chunking, one would have to copy larger strings into a contiguous memory before being able to use them and potentially offset the gains by the efficient scheme. Consequently, choosing the correct scheme has trade-offs that must be carefully weighed against each other and co-designed with the system.

Summary. To conclude, we have found that the optimistic techniques scale and perform better than pessimistic techniques for workloads in which the RNIC's capability of handling RDMA atomic is the limiting factor. For instance, this happens in a read-only workload with small tuples and 128 or more workers. When using larger tuples, the workload tends to become network bound, reducing the number of possible parallel operations and shifting the bottleneck for pessimistic schemes from the RNIC contention to the network. Thus, a pessimistic scheme often performs better in such scenarios. Another interesting finding is that in combination with optimistic reads, we can improve the scalability of writes by using the write-unlatch optimization. In contrast to pessimistic schemes, optimistic schemes come with different trade-offs regarding computational and storage requirements.

4.6 Effect on a Disaggregated DBMS

Finally, we compare the optimistic schemes by implementing them in NAM-DB and call this system OPT-DB. Similar to the experiment in Section 3.5, we vary the contention and show write-only, mixed, and read-only workloads. The configuration is unchanged with 4 compute and storage nodes storing $20M$ tuples. Unlike the experiment in Section 3.5, we stick to 512 byte but vary the number of workers from 112 to 224.

Contended shared latches. In the read-only workload in Figure 16, we can observe that all optimistic techniques are pretty robust against contention, contrary to the pessimistic synchronization in NAM-DB++, which already degrades with light skew (1.5). Furthermore, we can back up the previous findings that the pessimistic scheme performs better with fewer workers than the optimistic schemes. The optimistic approaches perform better when increasing the number of workers to 224. Overall, the OPT-DB implemented with FaRM performs the best, followed by CRC and Versioning (for 512 byte tuples).

Reader/writer contention and aborts. There is a well-known trade-off for optimistic approaches in reader-writer contention: optimistic approaches typically lead to more restarts, resulting in wasted work. Nevertheless, the mixed workload shows that the optimistic approaches out-perform NAM-DB++ despite the restarts.

The reason lies again in the fewer RDMA atomic operations. Especially under contention (skew), the atomics fall in the same lock slot, which creates physical contention in the RNIC. Moreover, in this experiment, the logical contention exacerbates this physical effect. Once the RDMA atomic operation is executed, it does not necessarily mean that the lock operation was successful, e.g., it may be already latched exclusively by another worker. Thus, for a pessimistic scheme, the operations may restart, which incurs additional atomic operations. To this end, for the mixed workload, the optimistic approaches, which get away with a single atomic operation for exclusive latches and none for optimistic (read) acquisitions, perform much better.

The same holds for the write-only benchmark; since the optimistic approaches harmonize with the write-unlatch optimization, they perform on par or sometimes even better (e.g., with low contention and 224 workers as shown in Figure 16). Moreover, in the write-only workload with 112 workers, the versioning approach is slightly faster since the writers do not need to perform

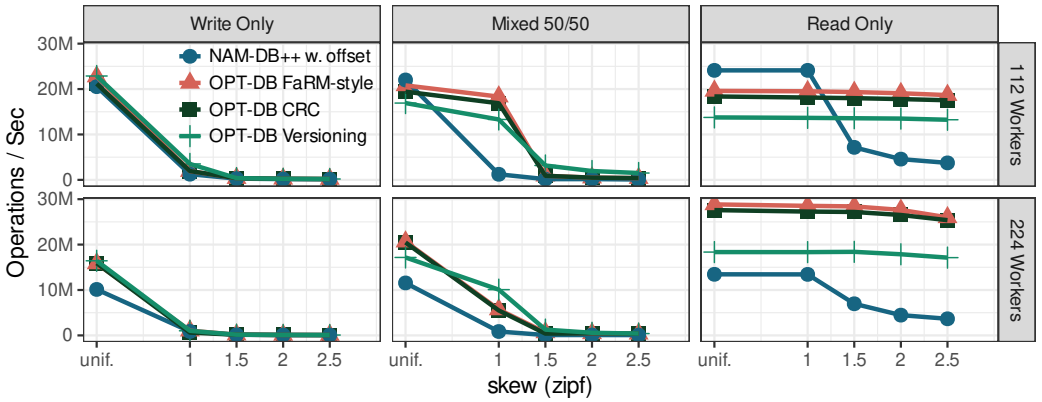

Fig. 16. Optimistic techniques in a disaggregated DBMS (4 compute nodes and 4 storage nodes).

any additional computation, such as CRC checks or incrementing the cache line versions. When contention increases, all the approaches converge at merely $100K$ operations per second.

5 DISCUSSION OF OTHER APPROACHES

This section discusses other synchronization techniques with relaxed guarantees.

Consistent optimistic read. The synchronization schemes that we discussed before captured latch semantics. That is, they include an additional validation step after retrieving a data value and operating on it to make sure that the data did not change in the meantime. However, these strong semantics are not required for all use cases [31, 32, 61]. A classic example is a key-value store that is only concerned about returning consistent data to clients and does not consider the clients' operations on the data. Our presented optimistic schemes can be easily adapted to support such relaxed semantics by removing the additional validation step at the end.

Other incorrect schemes. Another common technique is called bookend versioning which we briefly showed in Figure 1(a). While incorrect, the intuition of this approach is similar to versioning, where the second version (embedded in the data) validates that the data was not concurrently updated during the RDMA read. This approach suffers from the ordering problem, and its use in recent literature (e.g., [47]) is a testament to the subtleties involved in one-sided synchronization.

Out-of-place updates. An alternative design to the data consistency techniques described in Section 4.3 is to leverage out-of-place updates [4, 38, 50], eliminating concurrent data modifications using a copy-on-write strategy. This approach has the same communication overhead as the (broken) single RDMA read versioning technique but comes at the obvious cost of additional storage overhead. Out-of-place updates also enable the use of a technique we call *marking*. Marking is a simple detection technique that can alert readers by relying on a logical flag to indicate that a concurrent write is taking place. This approach manifests as logical insertion or deletion and a busy or pending update flag. However, when combined with in-place updates, marking suffers from the same re-ordering problem as versioning and can fail to detect an inconsistency because of the PCIe bus. Unfortunately, this incorrect use of marking has found its way into existing literature [13, 57].

6 LESSONS LEARNED AND CONCLUSION

As the first analytical study of RDMA synchronization primitives, we aim to leave the reader with a distilled perspective on the lessons learned from our work. These lessons are cast as *anti-patterns*, which reflect common mistakes that lead to incorrect designs, and important design *considerations*, which do not impact correctness but can be detrimental to performance and system complexity.

Anti-pattern #1.
Reliance on cache line order within a single RDMA read.

The first anti-pattern stems directly from the lack of ordering in the PCIe bus, which was first introduced in Section 1 and discussed in more detail in Section 4.2. Designs fall prey to this anti-pattern and thus incorrectly assume that an RDMA read observes the same order in which writes to different cache lines are made. Examples that relied on this anti-pattern include bookend versioning [32, 47] and marking with in-place updates [13, 57]. For instance, in [57], a writer first marks data as invisible to readers, then updates the value. As we have already established (cf. Section 4.2), a reader may observe these two steps out-of-order and, therefore, could accidentally assume a corrupted read is consistent. Worse, the version retrieved during this initial read could later be (incorrectly) validated. This behavior is identical to the broken versioning scheme discussed in Section 4.2.

Anti-pattern #2.
Reliance on the ordering of overlapping reads.

One solution to address the first anti-pattern is to manually enforce an order by issuing multiple reads. However, there are also pitfalls to this approach because distinct RDMA read operations also have subtle ordering guarantees. The second anti-pattern captures the lack of ordering between overlapping reads from the same connection. While RDMA operations can be fenced to enforce ordering in specific cases, this does not pertain to RDMA reads. In other words, there is currently no mechanism to enforce the order in which a remote machine processes RDMA reads without waiting for completion. Optimistic techniques like versioning hence require two sequential reads to ensure that the version is read before the data, as explained in Section 4.3. One-shot optimistic approaches, such as FaRM, do not suffer from this because they detect inconsistency at the granularity of cache lines.

Anti-pattern #3.
Reliance on the atomicity of RDMA writes and RDMA atomics.

The third anti-pattern results from the lack of atomicity guarantees by RDMA write and RDMA atomic operations in the RDMA specification. As Section 3.3 describes, various optimizations to improve pessimistic synchronization techniques exist. Recall that the write-unlatch optimization is only permissible when the locking primitive exclusively utilizes CAS operations. Again, this is because there is a store buffer in the RNIC. CAS operations are a special case because the write-back to memory is conditional. While it is possible to correctly utilize the write-unlatch optimization [46], as shown in Section 4.5, it requires a careful understanding of the underlying hardware. Therefore, we generally caution against designs that combine RDMA writes with RDMA atomic operations but acknowledge that in some instances where the lack of atomicity does not break semantics, it can yield better performance.

Consideration #1.
Data alignment impacts the RDMA atomic scalability.

The first consideration is described in detail in Section 3.2, but we review it here. Because of the lock table present in current RNICs, physical contention between independent latches can arise. We demonstrate that this is not only easily demonstrated by microbenchmarks that highlight the behavior but also have an important role in the scalability of a real system (i.e., NAM-DB). Hence, we advocate that system designers pay close attention to their data alignment when leveraging RDMA atomic operations to avoid this physical contention.

Consideration #2.
Optimistic synchronization performs best under high contention.

Conventional wisdom suggests that pessimistic synchronization pays off in high-contention write-heavy workloads because it eliminates excessive retries. However, our analysis demonstrates that this perspective does not hold for RDMA-based synchronization. Notably, in Section 4.5, we show that in write-heavy workloads, the optimistic approaches can match and even surpass pessimistic ones. While readers often retry in optimistic schemes, pessimistic approaches are subject to RNIC contention. Our results suggest that pessimistic approaches are beneficial when there is less skew, and the number of concurrent workers is small. Interestingly, this also applies in the read-only case since the sequential overhead is low compared to the optimistic strategies, which require validation. Therefore, optimistic synchronization should be preferred for highly skewed workloads to avoid RDMA atomic bottlenecks. As we established with the B-tree performance in Figure 11, this is particularly beneficial when there is a single point of contention, e.g., the root node of a B-tree.

Consideration #3.
Optimistic synchronization has non-negligible overheads.

Although they outperform pessimistic synchronization in many scenarios, there are computational and storage trade-offs among the various optimistic approaches, which we discussed in Section 4.5. Versioning requires an additional RDMA read compared to the other techniques, which increases operation latency and is most evident for small data sizes. On the other hand, CRC is computationally expensive, becoming untenable at large data sizes ($\geq$ 4 KB). Cache line versioning generally performs well but has an increased storage cost, not to mention that software must either handle the embedded versions or copy the data out. While optimistic synchronization is beneficial in many scenarios, it is not a silver bullet as they often have more complex designs compared to pessimistic schemes.

Consideration #4.
Pessimistic synchronization is more "future proof".

As demonstrated, subtle hardware characteristics can have a profound impact on correctness. We highlight this using a widespread deployment but point out that disaggregated memory technologies are in active development. For example, advancements like CXL [40] – a protocol built on PCI-e to support memory coherence across devices – are poised to disrupt the status quo. Changes to intermediate components of RDMA communication may alter the behavior of concurrent RDMA reads and writes, and therefore optimistic synchronization schemes, in unpredictable ways. In contrast, RDMA Atomic operations implement a well-established higher-level abstraction independent of hardware implementation. Hence, system designers should lean toward simple pessimistic synchronization when prioritizing production stability until the community has successfully converged on well-defined memory semantics for RDMA reads and writes.

Conclusion. This paper is the first paper that holistically (1) highlights the subtleties of RDMA-based synchronization regarding the correctness, (2) provides a robust performance analysis of existing synchronization techniques and optimizations, (3) demonstrates pitfalls of existing designs, and (4) offers a set of anti-patterns and design considerations for enabling developers to design correct and high-performance RDMA-enabled system.

ACKNOWLEDGMENTS

This work was partially funded by the German Research Foundation priority program 2037 (DFG) under the grants BI2011/1 & BI2011/2, the DFG Collaborative Research Center 1053 (MAKI), and the state of Hesse as part of the NHR Program. We also thank hessian.AI, DFKI, Mellanox, and 3AI for their support. We also like to thank Torsten Hoefler for an insightful discussion about the memory model of RDMA and RDMA's ordering guarantees.

REFERENCES

[1] ARM. 2018. Arm CoreLink CCI-550 Cache Coherent Interconnect Technical Reference Manual. https://developer.arm.com/documentation/100282/0100/?lang=en. https://developer.arm.com/documentation/100282/0100/?lang=en

[2] ARM. 2021. Introducing the AMBA Coherent Hub Interface. https://developer.arm.com/documentation/102407/0100

[3] Wei Cao, Yingqiang Zhang, Xinjun Yang, Feifei Li, Sheng Wang, Qingda Hu, Xuntao Cheng, Zongzhi Chen, Zhenjun Liu, Jing Fang, Bo Wang, Yuhui Wang, Haiqing Sun, Ze Yang, Zhushi Cheng, Sen Chen, Jian Wu, Wei Hu, Jianwei Zhao, Yusong Gao, Songlu Cai, Yunyang Zhang, and Jiawang Tong. 2021. PolarDB Serverless. In *Proceedings of the 2021 International Conference on Management of Data*. ACM. https://doi.org/10.1145/3448016.3457560

[4] Benjamin Cassell, Tyler Szepesi, Bernard Wong, Tim Brecht, Jonathan Ma, and Xiaoyi Liu. 2017. Nessie: A Decoupled, Client-Driven Key-Value Store Using RDMA. *IEEE Trans. Parallel Distributed Syst.* 28, 12 (2017), 3537–3552. https://doi.org/10.1109/TPDS.2017.2729545

[5] Yeounoh Chung and Erfan Zamanian. 2015. Using RDMA for Lock Management. *CoRR* abs/1507.03274 (2015). arXiv:1507.03274 http://arxiv.org/abs/1507.03274

[6] NVIDIA Coporation. 2021. NVIDIA InfiniBand Adaptive Routing Technology. Whitepaper WP-10326-001_v01.

[7] Andrei Marian Dan, Patrick Lam, Torsten Hoefler, and Martin T. Vechev. 2016. Modeling and analysis of remote memory access programming. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016, part of SPLASH 2016, Amsterdam, The Netherlands, October 30 - November 4, 2016*, Eelco Visser and Yannis Smaragdakis (Eds.). ACM, 129–144. https://doi.org/10.1145/2983990.2984033

[8] Aleksandar Dragojevic, Dushyanth Narayanan, Miguel Castro, and Orion Hodson. 2014. FaRM: Fast Remote Memory. In *NSDI*.

[9] Aleksandar Dragojevic, Dushyanth Narayanan, Edmund B. Nightingale, Matthew Renzelmann, Alex Shamis, Anirudh Badam, and Miguel Castro. 2015. No compromises: distributed transactions with consistency, availability, and performance. In *SOSP*.

[10] Jingwen Du, Fang Wang, Dan Feng, Changchen Gan, Yuchao Cao, Xiaomin Zou, and Fan Li. 2023. Fast One-Sided RDMA-Based State Machine Replication for Disaggregated Memory. *ACM Trans. Archit. Code Optim.* (mar 2023). https://doi.org/10.1145/3587096 Just Accepted.

[11] Philipp Fent, Alexander van Renen, Andreas Kipf, Viktor Leis, Thomas Neumann, and Alfons Kemper. 2020. Low-Latency Communication for Fast DBMS Using RDMA and Shared Memory. In *ICDE*.

[12] Torsten Hoefler, Duncan Roweth, Keith Underwood, Bob Alverson, Mark Griswold, Vahid Tabatabaee, Mohan Kalkunte, Surendra Anubolu, Siyan Shen, Abdul Kabbani, Moray McLaren, and Steve Scott. 2023. Datacenter Ethernet and RDMA: Issues at Hyperscale. arXiv:2302.03337 [cs.NI]

[13] Chenchen Huang, Huiqi Hu, Xuecheng Qi, Xuan Zhou, and Aoying Zhou. 2021. RS-store: RDMA-enabled skiplist-based key-value store for efficient range query. *Frontiers of Computer Science* 15, 6 (sep 2021). https://doi.org/10.1007/s11704-020-0126-6

[14] Ram Huggahalli, Ravi R. Iyer, and Scott Tetrick. 2005. Direct Cache Access for High Bandwidth Network I/O. In *32st International Symposium on Computer Architecture (ISCA 2005), 4-8 June 2005, Madison, Wisconsin, USA*. IEEE Computer Society, 50–59. https://doi.org/10.1109/ISCA.2005.23

[15] InfiniBand Trade Association 2007. *InfiniBand Architecture Specification Volume 1*. InfiniBand Trade Association. Release 1.2.1.

[16] InfiniBand Trade Association. 2010. RDMA Over Converged Ethernet (RoCE). https://cw.infinibandta.org/document/dl/7148.

[17] Intel. 2012. Intel Data Direct I/O Technology (Intel DDIO): A P rimer. https://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/data-direct-i-o-technology-brief.pdf

[18] Anuj Kalia, Michael Kaminsky, and David G. Andersen. 2014. Using RDMA efficiently for key-value services. In *SIGCOMM*.

[19] Anuj Kalia, Michael Kaminsky, and David G. Andersen. 2016. Design Guidelines for High Performance RDMA Systems. *login Usenix Mag.* 41, 3 (2016).

[20] Anuj Kalia, Michael Kaminsky, and David G. Andersen. 2016. FaSST: Fast, Scalable and Simple Distributed Transactions with Two-Sided (RDMA) Datagram RPCs. In *OSDI*.

[21] Tejas Karmarkar. 2015. Availability of Linux RDMA on Microsoft Azure. Online. https://azure.microsoft.com/en-us/blog/azure-linux-rdma-hpc-available/

[22] Dario Korolija, Dimitrios Koutsoukos, Kimberly Keeton, Konstantin Taranov, Dejan S. Milojicic, and Gustavo Alonso. 2022. Farview: Disaggregated Memory with Operator Off-loading for Database Engines. In *12th Conference on Innovative Data Systems Research, CIDR 2022 , Chaminade, CA, USA, January 9-12, 2022*. www.cidrdb.org. https://www.cidrdb.org/cidr2022/papers/p11-korolija.pdf

[23] Viktor Leis, Michael Haubenschild, and Thomas Neumann. 2019. Optimistic Lock Coupling: A Scalable and Efficient General-Purpose Synchronization Method. *IEEE Data Eng. Bull.* 42 (2019), 73–84.

[24] Viktor Leis, Florian Scheibner, Alfons Kemper, and Thomas Neumann. 2016. The ART of practical synchronization. In *Proceedings of the 12th International Workshop on Data Management on New Hardware, DaMoN 2016, San Francisco, CA, USA, June 27, 2016*. ACM, 3:1–3:8. https://doi.org/10.1145/2933349.2933352

[25] Edgar A. León, Kurt B. Ferreira, and Arthur B. Maccabe. 2007. Reducing the Impact of the MemoryWall for I/O Using Cache Injection. In *15th Annual IEEE Symposium on High-Performance Interconnects, HOTI 2007, Stanford, CA, USA, August 22-24, 2007*, John W. Lockwood, Fabrizio Petrini, Ron Brightwell, and Dhabaleswar K. Panda (Eds.). IEEE Computer Society, 143–150. https://doi.org/10.1109/HOTI.2007.8

[26] Jiuxing Liu, Jiesheng Wu, Sushmitha P. Kini, Pete Wyckoff, and Dhabaleswar K. Panda. 2003. High performance RDMA-based MPI implementation over InfiniBand. In *ICS*.

[27] Simon Loesing, Markus Pilman, Thomas Etter, and Donald Kossmann. 2015. On the Design and Scalability of Distributed Shared-Data Databases. In *SIGMOD*.

[28] Teng Ma, Kang Chen, Shaonan Ma, Zhuo Song, and Yongwei Wu. 2021. Thinking More about RDMA Memory Semantics. In *IEEE International Conference on Cluster Computing, CLUSTER 2021, Portland, OR, USA, September 7-10, 2021*. IEEE, 456–467. https://doi.org/10.1109/Cluster48925.2021.00033

[29] Teng Ma, Dongbiao He, and Gordon Ning Liu. 2021. HybridSkipList: A Case Study of Designing Distributed Data Structure with Hybrid RDMA. In *IEEE 45th Annual Computers, Software, and Applications Conference, COMPSAC 2021, Madrid, Spain, July 12-16, 2021*. IEEE, 68–73. https://doi.org/10.1109/COMPSAC51774.2021.00021

[30] Christopher Mitchell, Yifeng Geng, and Jinyang Li. 2013. Using One-Sided RDMA Reads to Build a Fast, CPU-Efficient Key-Value Store. In *2013 USENIX Annual Technical Conference, San Jose, CA, USA, June 26-28, 2013*, Andrew Birrell and Emin Gün Sirer (Eds.). USENIX Association, 103–114. https://www.usenix.org/conference/atc13/technical-sessions/presentation/mitchell

[31] Christopher Mitchell, Yifeng Geng, and Jinyang Li. 2013. Using One-Sided RDMA Reads to Build a Fast, CPU-Efficient Key-Value Store. In *USENIX ATC*.

[32] Christopher Mitchell, Kate Montgomery, Lamont Nelson, Siddhartha Sen, and Jinyang Li. 2016. Balancing CPU and Network in the Cell Distributed B-Tree Store. In *USENIX ATC*.

[33] Sundeep Narravula, A. Marnidala, Abhinav Vishnu, Karthikeyan Vaidyanathan, and Dhabaleswar K. Panda. 2007. High Performance Distributed Lock Management Services using Network-based Remote Atomic Operations. In *Seventh IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2007), 14-17 May 2007, Rio de Janeiro, Brazil*. IEEE Computer Society, 583–590. https://doi.org/10.1109/CCGRID.2007.58

[34] Jacob Nelson and Roberto Palmieri. 2020. Performance Evaluation of the Impact of NUMA on One-sided RDMA Interactions. In *SRDS*.

[35] PCI-SIG. 2014. PCI Express Base Specification Revision 4.0. (2014).

[36] R. Recio, B. Metzler, P. Culley, J. Hilland, and D. Garcia. 2007. *A Remote Direct Memory Access Protocol Specification*. Technical Report. https://doi.org/10.17487/rfc5040

[37] Yufei Ren, Tan Li, Dantong Yu, Shudong Jin, and Thomas G. Robertazzi. 2013. Design and performance evaluation of NUMA-aware RDMA-based end-to-end data transfer systems. In *HiPC*.

[38] Zhenyuan Ruan, Malte Schwarzkopf, Marcos K. Aguilera, and Adam Belay. 2020. AIFM: High-Performance, Application-Integrated Far Memory. In *14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2020, Virtual Event, November 4-6, 2020*. USENIX Association, 315–332. https://www.usenix.org/conference/osdi20/presentation/ruan

[39] H. Shah, F. Marti, W. Noureddine, A. Eiriksson, and R. Sharp. 2014. *Remote Direct Memory Access (RDMA) Protocol Extensions*. Technical Report. https://doi.org/10.17487/rfc7306

[40] Debendra Das Sharma. 2019. *Compute Express Link*. Technical Report. Compute Express Link.

[41] Arjun Singhvi, Aditya Akella, Maggie Anderson, Rob Cauble, Harshad Deshmukh, Dan Gibson, Milo M. K. Martin, Amanda Strominger, Thomas F. Wenisch, and Amin Vahdat. 2021. CliqueMap: Productionizing an RMA-Based Distributed Caching System. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. ACM. https://doi.org/10.1145/3452296.3472934

[42] Dan Tang, Yungang Bao, Weiwu Hu, and Mingyu Chen. 2010. DMA cache: Using on-chip storage to architecturally separate I/O data from CPU data for improving I/O performance. In *16th International Conference on High-Performance Computer Architecture (HPCA-16 2010), 9-14 January 2010, Bangalore, India*, Matthew T. Jacob, Chita R. Das, and Pradip Bose (Eds.). IEEE Computer Society, 1–12. https://doi.org/10.1109/HPCA.2010.5416638

[43] Konstantin Taranov, Fabian Fischer, and Torsten Hoefler. 2022. Efficient RDMA Communication Protocols. arXiv:2212.09134 [cs.NI]

[44] Konstantin Taranov, Salvatore Di Girolamo, and Torsten Hoefler. 2021. CoRM: Compactable Remote Memory over RDMA. In *SIGMOD '21: International Conference on Management of Data, Virtual Event, China, June 20-25, 2021*. 1811–1824. https://doi.org/10.1145/3448016.3452817

[45] Shin-Yeh Tsai, Yizhou Shan, and Yiying Zhang. 2020. Disaggregating Persistent Memory and Controlling Them Remotely: An Exploration of Passive Disaggregated Key-Value Stores. In *2020 USENIX Annual Technical Conference,*

USENIX ATC 2020, July 15-17, 2020, Ada Gavrilovska and Erez Zadok (Eds.). USENIX Association, 33–48. https://www.usenix.org/conference/atc20/presentation/tsai

[46] Chao Wang and Xuehai Qian. 2021. RDMA-enabled Concurrency Control Protocols for Transactions in the Cloud Era. *IEEE Transactions on Cloud Computing* (2021), 1–1. https://doi.org/10.1109/tcc.2021.3116516

[47] Qing Wang, Youyou Lu, and Jiwu Shu. 2022. Sherman: A Write-Optimized Distributed B+Tree Index on Disaggregated Memory. In *Proceedings of the 2022 International Conference on Management of Data*. ACM. https://doi.org/10.1145/3514221.3517824

[48] Ruihong Wang, Jianguo Wang, Stratos Idreos, M. Tamer Özsu, and Walid G. Aref. 2022. The Case for Distributed Shared-Memory Databases with RDMA-Enabled Memory Disaggregation. *CoRR* abs/2207.03027 (2022). https://doi.org/10.48550/arXiv.2207.03027 arXiv:2207.03027

[49] Tinggang Wang, Shuo Yang, Hideaki Kimura, Garret Swart, and Spyros Blanas. 2020. Efficient Usage of One-Sided RDMA for Linear Probing. In *International Workshop on Accelerating Analytics and Data Management Systems Using Modern Processor and Storage Architectures, ADMS@VLDB 2020, Tokyo, Japan, August 31, 2020*, Rajesh Bordawekar and Tirthankar Lahiri (Eds.). 1–13. http://www.adms-conf.org/2020-camera-ready/ADMS20_06.pdf

[50] Yandong Wang, Li Zhang, Jian Tan, Min Li, Yuqing Gao, Xavier Guerin, Xiaoqiao Meng, and Shicong Meng. 2015. HydraDB: a resilient RDMA-driven key-value middleware for in-memory cluster computing. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2015, Austin, TX, USA, November 15-20, 2015*, Jackie Kern and Jeffrey S. Vetter (Eds.). ACM, 22:1–22:11. https://doi.org/10.1145/2807591.2807614

[51] Xingda Wei, Rong Chen, Haibo Chen, and Binyu Zang. 2021. XStore: Fast RDMA-Based Ordered Key-Value Store Using Remote Learned Cache. *ACM Trans. Storage* 17, 3 (2021), 18:1–18:32. https://doi.org/10.1145/3468520

[52] Xingda Wei, Zhiyuan Dong, Rong Chen, and Haibo Chen. 2018. Deconstructing RDMA-enabled Distributed Transactions: Hybrid is Better!. In *OSDI*.

[53] Xingda Wei, Jiaxin Shi, Yanzhe Chen, Rong Chen, and Haibo Chen. 2015. Fast in-memory transaction processing using RDMA and HTM. In *SOSP*.

[54] Erfan Zamanian, Carsten Binnig, Tim Kraska, and Tim Harris. 2016. The End of a Myth: Distributed Transactions Can Scale. *CoRR* abs/1607.00655 (2016).

[55] Erfan Zamanian, Julian Shun, Carsten Binnig, and Tim Kraska. 2021. Chiller: Contention-centric Transaction Execution and Data Partitioning for Modern Networks. *SIGMOD Rec.* 50, 1 (2021).

[56] Ming Zhang, Yu Hua, Pengfei Zuo, and Lurong Liu. 2022. FORD: Fast One-sided RDMA-based Distributed Transactions for Disaggregated Persistent Memory. In *20th USENIX Conference on File and Storage Technologies, FAST 2022, Santa Clara, CA, USA, February 22-24, 2022*, Dean Hildebrand and Donald E. Porter (Eds.). USENIX Association, 51–68. https://www.usenix.org/conference/fast22/presentation/zhang-ming

[57] Ming Zhang, Yu Hua, Pengfei Zuo, and Lurong Liu. 2022. FORD: Fast One-sided RDMA-based Distributed Transactions for Disaggregated Persistent Memory. In *20th USENIX Conference on File and Storage Technologies (FAST 22)*. USENIX Association, Santa Clara, CA, 51–68. https://www.usenix.org/conference/fast22/presentation/zhang-ming

[58] Yingqiang Zhang, Chaoyi Ruan, Cheng Li, Jimmy Yang, Wei Cao, Feifei Li, Bo Wang, Jing Fang, Yuhui Wang, Jingze Huo, and Chao Bi. 2021. Towards Cost-Effective and Elastic Cloud Database Deployment via Memory Disaggregation. *Proc. VLDB Endow.* 14, 10 (2021), 1900–1912. https://doi.org/10.14778/3467861.3467877

[59] Tobias Ziegler, Carsten Binnig, and Viktor Leis. 2022. ScaleStore: A Fast and Cost-Efficient Storage Engine using DRAM, NVMe, and RDMA. In *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*. ACM, 685–699. https://doi.org/10.1145/3514221.3526187

[60] Tobias Ziegler, Sumukha Tumkur Vani, Carsten Binnig, Rodrigo Fonseca, and Tim Kraska. 2019. Designing Distributed Tree-based Index Structures for Fast RDMA-capable Networks. In *SIGMOD*.

[61] Pengfei Zuo, Qihui Zhou, Jiazhao Sun, Liu Yang, Shuangwu Zhang, Yu Hua, James Cheng, Rongfeng He, and Huabing Yan. 2022. RACE: One-sided RDMA-conscious Extendible Hashing. *ACM Transactions on Storage* 18, 2 (may 2022), 1–29. https://doi.org/10.1145/3511895

Received April 2022; revised July 2022; accepted August 2022