

Persistent Memory Transactions

Virendra J. Marathe¹ Achin Mishra² Ameer Trivedi³ Yihe Huang⁴ Faisal Zaghoul⁵
Sanidhya Kashyap⁶ Margo Seltzer^{1,4} Tim Harris¹ Steve Byan¹ Bill Bridge⁷ Dave Dice¹

¹Oracle Labs ²Dell Inc. ³University of Massachusetts, Amherst ⁴Harvard University
⁵Yale University ⁶Georgia Institute of Technology ⁷Oracle Corporation

Abstract

This paper presents a comprehensive analysis of implementation choices for transaction runtime systems optimized for persistent memory. Our work focuses on performance implications of *persist barriers*, primitives required to persist writes on persistent memory. In the process we introduce a new taxonomy of *persistence domains* (portion of the memory hierarchy where the data is effectively persistent) that has a significant impact on persist barrier latencies. We present algorithms for *undo logging*, *redo logging*, and *copy-on-write* based transactions, as well as a memory allocator, all optimized to reduce persist barriers per transaction. Our microbenchmarking does a comprehensive sweep of read-write mix ratios in transactions, showing performance trends of the transaction runtimes under different assumptions about persist barrier latencies. No single runtime dominates the rest across the board. However, we pinpoint approximate read-write ratio ranges where specific runtimes outperform the rest. Our analysis highlights the significant influence of multiple factors on performance – transaction runtime specific bookkeeping overheads, persist barrier latencies, and cache locality. We find similar performance trade offs on three “real world” workloads – a key-value store of our own, SQLite, and memcached.

1 Introduction

Byte-addressable persistent memory technologies (e.g. *spin-transfer torque MRAM (STT-MRAM)* [14, 15], *memristors* [32]), that approach the performance of DRAM (100-1000x faster than state-of-the-art NAND flash) are coming, as evidenced by Intel and Micron Technologies’ recent announcement of their 3D XPoint persistent memory technology [1]. While the simple load/store based interface of these technologies is appealing, it introduces new challenges; a simple `store` does not immediately persist data, because processor state and various layers or the memory hierarchy (e.g., store buffers, caches) are expected to remain *nonpersistent* for the foreseeable future. Prior research [9, 19, 27] and processor vendors, such as Intel, have proposed new hardware instructions [16]

```
void *p; // pointer to persistent memory
...
// obj and clone(obj) are persistent
p = clone(obj);
```

Figure 1: Example illustrating complexities of programming with just the hardware instructions for persisting data. `p` is a pointer hosted in persistent memory. `clone` clones its argument object (`obj`). The programmer must persist the clone before `p`’s assignment, otherwise an untimely failure could result in a state where the clone is not persisted but `p`’s new value is persisted.

to flush or write cache lines back to lower layers in the memory hierarchy and new forms of *persist barrier* instructions that can be used to order persistence of stores. However, even with these new instructions, correctly writing programs to use them remains a daunting task. Figure 1 illustrates this challenge – the programmer must carefully reason about the order in which updates to various pieces of the application’s persistent data structures are persisted. Omission of even a single flush, write back, or persist barrier instruction can result in persistent data inconsistencies in the face of failures.

These programming challenges have resulted in investigations of transactional mechanisms to access and manipulate data on persistent memory [5, 6, 8, 10, 20, 28, 36]. However, we observe that research on transaction runtime implementations in this new context is in its early stages. In particular, the combination of a load/store interface to storage and high overheads of persist barriers is a significant change from the traditional transaction settings of both databases [12] and memory transactions [13]. As such, it has significant implications on performance trade offs between various implementation approaches. Most early work has largely ignored this consideration [5, 6, 8, 28], while some recent efforts are attempting to explicitly reduce the number of persist barriers in transactions [10, 20, 21, 36]. Our objective is similar, but we take a more holistic approach:

- We consider the implications of *persistence domains* [33] on persist barrier overheads (see § 2). Briefly, a persistence domain is the portion of the memory hierarchy that is considered to be “effectively persistent” – the underlying hardware/software

system ensures that data that reaches its persistence domain is written to the persistent media before the system is shut down, either planned or due to failures. We introduce a new taxonomy of persistence domain choices enabled by different hardware systems.

- We present a comprehensive investigation of implementation choices for persistent memory transactions with the focus on how transaction writes are implemented (see § 4). We present *undo logging*, *redo logging*, and *copy-on-write (COW)* based transaction runtimes. Our runtimes are carefully designed to reduce the number of persist barriers executed within a transaction. While we reduce the number of persist barriers needed for our undo logging implementation by almost 50% relative to prior approaches, committing transactions in both of our redo logging and COW implementations requires only four persist barriers. Our memory management algorithm (see § 5) is also designed to reduce persist barriers in a transaction (2 per undo logging transaction, and effectively zero for redo logging and COW transactions).
- Our microbenchmarking (see § 6) comprehensively sweeps through read-write mix ratios within transactions and shows how performance trends in the transaction runtimes change as the read-write mix within transactions changes for different persist barrier latencies. Our analysis reveals the significant influence of a combination of factors – transaction runtime specific bookkeeping overheads, persist barrier latencies, and cache locality – which determines the performance of a runtime. Because of the interplay of these factors, no single runtime’s performance dominates the rest in all settings. We find similar performance trade offs in three “real world” workloads: (i) a key-value store we developed from scratch, (ii) a port of SQLite, and (iii) a port of memcached.

2 Persistence Domain

While data hosted in persistent memory DIMMs is expected to survive power failures, the rest of the memory hierarchy (e.g. processor caches, memory controller buffers, etc.) is fundamentally not persistent. However, system solutions do exist that make various parts of the memory hierarchy “effectively persistent”. For instance, in battery backed systems [7, 18, 25, 26], where the processor state and caches can be flushed out to persistent memory DIMMs on power failure, the whole memory hierarchy effectively becomes persistent. Another example is the *asynchronous DRAM refresh (ADR)* feature provided by

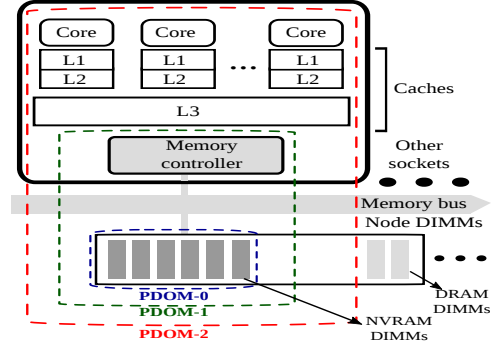


Figure 2: Persistence domains of a near future processor socket that hosts persistent memory DIMMs.

Operations	Persistence domains		
	PDOM-0	PDOM-1	PDOM-2
Writes	store clwb/clflush-opt	store clwb/clflush-opt	store
Ordering persists	sfence pcommit	sfence	nop

Table 1: Persistent memory writes and primitives needed on Intel’s upcoming processors [16] for different persistence domains.

modern processors, where the memory controller buffers are flushed out to memory DIMMs on power failure [17]. With the ADR feature, the memory controller buffers can be considered effectively persistent since the data is guaranteed, discounting ADR hardware failures, to persist. There may be other ways to slice the memory hierarchy in persistent and nonpersistent parts; however, we focus on 3 specific partitioning strategies that we believe will capture most future system configurations.

We define *persistence domain* [33] as the portion of memory hierarchy where data is effectively persistent. As shown in Figure 2, we classify persistence domains in three categories: (i) PDOM-0, which contains only the persistent memory DIMMs. (ii) PDOM-1, which included PDOM-0 and memory controller buffers. Modern processors with ADR capabilities and persistent memory DIMMs effectively support PDOM-1. (iii) PDOM-2, which includes the entire memory hierarchy as well as processor state, such as store buffers, containing persistent data. Battery backed systems support PDOM-2.

The persistence domain affects the instruction sequence needed to persist updates. Table 1 depicts the instructions needed to persist these updates on (near future) Intel processors with persistent memory [16]. There are two phases to the persistent update process: (i) The actual write (i.e., store) and (ii) the persist barrier. PDOM-0 and PDOM-1 require a flush instruction in addition to the store to move data into the persistence domain. Both the clwb and clflush-opt trigger asynchronous cache-line

sized writes to the memory controller; they differ in that `clflush-opt` invalidates the cache line while `clwb` does not. In principle, the flush instructions can be delayed, and almost certainly should be for multiple store instructions to the same cache line. In practice, as they are asynchronous, starting the writeback sooner speeds up the persist barriers in the second phase of this process. In PDOM-2, flush instructions are not needed, since store buffers and caches are part of the persistence domain.

In PDOM-0, the persist barrier needs to ensure that all flushes have completed (the first sfence), and then force any updates in the memory controller to be written to the DIMMs (`pcommit`). As the `pcommit` is asynchronous, persistence requires the second sfence to indicate when the `pcommit` has completed. In PDOM-1, the persist barrier need only ensure that prior flushes have completed, since the memory controller now resides inside the persistence domain. PDOM-2 requires no further action as data is persisted as soon as it has been stored. Intel has recently deprecated the `pcommit` instruction [17]. However, we include it in our discussion as a concrete example of a PDOM-0 persistence domain. Note that `clwb`, `clflush-opt`, and `pcommit` have store semantics in terms of memory ordering, and applications must take care to avoid problematic reordering of loads with these instructions, using sfence or other instructions with fence semantics.

3 Programming Model

Our programming model is based on the abstractions of persistent regions, persistent data types, and transactions, similarly to prior works [5, 6, 8, 36].

Persistent regions. We assume that a file system hosts all persistent data accessible to applications [29, 35, 37]. Applications access persistent data directly via a persistent-region abstraction, using the `mmap` interface [33]. Each region has its own heap and a user instantiated root pointer. Transactions do not span regions, but can access nonpersistent objects residing outside a region.

Our persistent heap provides `pm_alloc` and `pm_free` functions to allocate and deallocate memory from the persistent region. To avoid memory leaks due to program crashes and other failures, these functions can be used only within transactions.

Transactions and accessors. Various semantic models for persistent memory transactions have been explored in prior work: full ACID transactions in the spirit of transactional memory [8, 36], failure-atomic critical sections [6], and failure-atomic transactions [5, 10, 20, 28]. Our work focuses on programming models for skilled system software developers who are adept to manage concurrency control on their own and require just *failure atomicity* of updates to persistent memory. As a result, we focus

```

struct foo {
    int cnt;
};
// pm_foo, the persistent version of type foo
DEFINE_PM_TYPE(foo);
// x points to an instance of pm_foo
pm_foo *x;
// failure-atomic transaction for x->cnt++;
pm_txn_t txn;
do {
    TXN_BEGIN(txn);
    int counter; // temporary
    TXN_READ(txn, x, cnt, &counter);
    counter++;
    TXN_WRITE(txn, x, cnt, &counter);
    status = TXN_COMMIT(txn);
} while (status != TXN_COMMITTED);

```

Figure 3: Example of a simple transaction that increments a counter in a persistent object.

on failure-atomic transactions that guarantee atomicity of writes to persistent memory, i.e., “all-or-nothing” semantics across failure boundaries, relegating isolation to applications. Failure-atomic transactions allow us to ensure consistency of persistent data, while allowing applications to leverage semantic information in their synchronization approach. In the example from Figure 1, the assignment statement can be delineated by a transaction to guarantee failure atomicity. If isolation is needed, the application might enclose the transaction or its enclosing context in a critical section.

Similar to prior work [5, 6, 8, 20, 36], we advocate the use of language-level constructs for transactions in the form of lexically scoped *failure-atomic blocks*. However, for prototyping purposes, we have implemented a library-based interface to our failure-atomic durable transactions. Our library contains APIs for accessing complex data structures, and beginning and ending transactions. Figure 3 illustrates the use of the API. Using this API, we have implemented various forms of transactional reads and writes (e.g., for individual fields, pointers, arbitrary sized buffers, etc.) as well as common memory buffer operators, e.g., `memcpy`, `memcmp`, and `memset`.

In the example, we create a persistent version of a type using `DEFINE_PM_TYPE`. This creates an opaque persistent data type that encloses an argument type instance.

We support nested transactions with flat-nesting semantics. Successful completion of a transaction (`TXN_COMMIT` returns `TXN_COMMITTED` status) ensures that all updates made by the transaction to persistent memory have been persisted. If a failure occurs before the transaction commits, the transaction’s effects are rolled back during recovery. We consider only the fail-stop failure model in our work. A transaction can

voluntarily roll back via the `TXN_ABORT` call. The programmer will explicitly manage the control flow in all cases (i.e., there is no automatic restart of the transaction after abort).

`PM_UNWRAP` returns a reference to the type instance wrapped within a persistent type instance. Programmers can directly access the “unwrapped” object’s fields using loads/stores. This avoids unnecessary transactional instrumentation overheads (e.g., when initializing a newly allocated object). The `pm_foo` instance from Figure 3 is unwrapped as follows:

```
foo *direct_x = PM_UNWRAP(x);
// no transactional instrumentation
direct_x->cnt = 0;
// write back and persist the update
pm_async_wb_loc(&direct_x->cnt,
               sizeof(direct_x->cnt));
pm_persist_barrier();
```

Modifying the unwrapped object makes the programmer responsible for writing and persisting affected cache lines, using the primitives we provide.

Concurrency control. While our failure-atomic transactions do not ensure isolation, we recognize that occasionally these transactions may invoke critical sections that modify persistent objects. This can lead to scenarios where a transaction executes a critical section, modifying persistent objects, while the prior transaction that executed the same critical section, modifying the same persistent objects, has still not completed. The prior transaction’s uncommitted writes may end up racing with the new transaction’s reads/writes of these persistent objects. To aid isolation for such circumstances, we provide the facility of a *deferred lock release*, where the lock acquired for a critical section’s execution is released when the enclosing durable block completes its execution. Furthermore, we also provide a hook for *deferred operations*, where the operation is executed at the end of the enclosing transaction (e.g. entire critical sections). Any issues related to deadlocks while using deferred operations are to be programmatically addressed by the programmer.

4 Implementing Persistent Transactions

We have developed three different transaction runtime systems: (i) undo logging, (ii) redo logging, and (iii) copy-on-write (COW). All the runtimes store transaction metadata in a persistent data structure called the *transaction descriptor*, which is assigned to a thread as part of `TXN_BEGIN`. A descriptor is always in one of four states: IDLE, RUNNING, ABORTED, or COMMITTED. A descriptor that is not in use is in the IDLE state. `TXN_BEGIN` transitions the descriptor into the RUNNING state. A transaction commits by entering the COMMITTED state and aborts

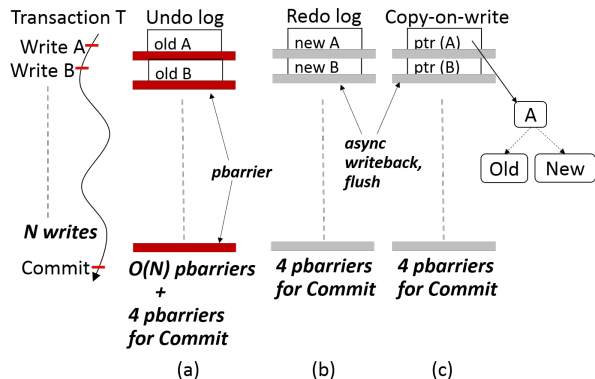


Figure 4: Transaction write implementations. Transaction T applies N distinct writes. `pbarrier` is a potentially high latency persist barrier, `async wb/flush` are asynchronous cache line write-back or flush instructions respectively.

by entering the ABORTED state. After the runtime cleans up a descriptor’s internal state and buffers, the descriptor returns to the IDLE state. During its execution, a transaction may read, write, allocate, and deallocate persistent objects using our API.

4.1 Undo Log based Transactions

Figure 4(a) shows transaction T ’s undo logging activities. The undo log is implemented as a simple chunked list. Transaction T writes w , using `TXN_WRITE` producing a log record containing the original value of w . This log record must be persisted before w is modified. A typical implementation of the undo log append requires two persist barriers – one to persist the new log record and one to change the log’s tail pointer. Although correct, this approach leads to $2N$ persist barriers for N appends, which results in high overheads.

Our implementation requires only a single persist barrier per record. Instead of relying on a correct tail pointer during recovery, we infer the tail of the log. We assign each transaction a monotonically increasing persistent 64-bit version number. Each undo record contains a transaction version number, a 64-bit checksum, and a 64-bit *prolog* sentinel value that appears at the beginning of the undo record. So, we write the prolog sentinel value, the transaction’s version number, the record contents, and then compute and write the checksum. Then we issue a single persist barrier. If a failure occurs before or during execution of the persist barrier, and only part of the undo record has been persisted, we will detect a checksum mismatch during recovery. We also maintain the log tail pointer, but update it *after* the persist barrier, so the tail update is guaranteed to persist on or before the next record write and persist. Recovery can rely on the following invariant: the tail pointer will be at most one record behind the actual end of log. So log recovery requires

only that we examine the record after the current end of log and determine if there is a valid log record present.

One of the most compelling benefits of undo logging is that transactional reads are implemented as uninstrumented loads [8, 5, 20, 28, 36].

An undo log transaction commits in four steps: (i) First it ensures that all transactional writes are persisted, which requires one persist barrier. (ii) Then it logically commits the transaction by appending the commit record to the transaction’s undo log. (It also switches the transaction’s state to COMMITTED, but that does not have to persist.) Steps (iii) and (iv) are largely related to transactional metadata cleanup, which requires persistence only if the transaction allocated or deallocated persistent memory (see § 5). (iii) Persist the allocation/deallocation calls’ effects and cleans up the transaction’s metadata. (iv) Mark the transaction IDLE; this state change needs to be persisted only if the transaction did allocations/deallocations.

4.2 Redo Log based Transactions

Figure 4(b) shows transaction T ’s redo logging activities. Like the undo log, the redo log is implemented as a simple chunked list. Transaction T writes w , using `TXN_WRITE` producing a log record containing the new value of w . The record need not persist at the time of the write; if a failure occurs, the entire redo log can be discarded. However, an implementation, like ours, may proactively schedule a low latency asynchronous writeback/flush of the record.

The challenge for redo logging schemes is handling read-after-write accesses. As the new value appears only in the log, a subsequent read must consult the log for the latest value. A naive implementation could walk down the whole log looking for the latest value. Furthermore, the redo log lookup could be done for *every* subsequent read done by the transaction. The resulting overhead can be significant. We apply two optimizations to overcome these overheads.

First, we add a 64-bit bitmap to each persistent object’s header to indicate current writers of that object. A writer first sets its corresponding bit during the write (the bit is cleared after the transaction completes). A read checks the object’s writers bitmap to determine if a redo log lookup is necessary, and does on if so. If a lookup is not necessary, the read becomes an uninstrumented load. Each transaction maps to a unique bit in the object’s writers bitmap. Up to 64 transactions can concurrently “own” a bit in the writers bitmap. This can be easily extended to larger bitmaps, but at present we force additional transactions to consult the redo log for all their reads.

Second, we avoid scanning the entire log by maintaining a per-transaction hash table indexed by persistent

object base addresses. The hash table record points to the latest redo log record for that object. Each such record also contains a pointer to the previous redo log record for that object, if one exists. We effectively superimpose a linked stack of records for each object within the redo log. This avoids unnecessary traversal of unrelated log records during a redo log lookup.

Committing a transaction requires persisting the redo log. After the persist completes, the transaction logically commits by updating its state to COMMITTED, and then persists the new state with a second persist barrier. After the logical commit, the runtime applies the redo log to each modified object and issues a third persist barrier. Finally, we mark the transaction IDLE, and persist it. In total, the redo logging implementation requires four persist barriers for commit, but none on abort.

4.3 Copy-on-Write based Transactions

Our copy-on-write (COW) implementation introduces an extra level of indirection between a persistent type instance (the wrapper) and the real data type instance (payload) it encloses. As shown in Figure 4(c), the persistent type contains pointers to *old* and *new* versions of the enclosed type’s instances. Before modifying an object, a transaction creates a new copy of the payload. We provide a special `TXN_OPEN` API that applications can use to obtain read-only or read-write access to a persistent object:

```
TXN_OPEN(txn, obj, mode, copy_ctor);
```

where *mode* is either read-only or read-write, and *copy_ctor* is the copy constructor. The copy constructor can be used to clone specialized objects (e.g. linked structures, self-relative pointers, etc.). A NULL copy constructor will default to using `memcpy`.

Each transaction descriptor maintains a *write set* containing the list of objects the transaction has written. Objects are added to the write set in `TXN_OPEN` invocations with read-write mode. Object wrappers also contain the writing transaction’s ID (assuming at most one writer per persistent object), which is used to direct transactional reads to the appropriate payload copy.

Payload copies, as well as writes to their wrappers, need not be persisted during the writer’s transaction. The transaction’s write set and the objects it writes to are persisted using a single persist barrier at the beginning of the commit operation. Then, the runtime updates the transaction’s state to COMMITTED and persists it.

The post-commit cleanup requires four steps: (i) make the modified (new) object payload the real (old) payload, (ii) reset *new* to NULL, (iii) discard (deallocate) the old payload, and (iv) clear the writer’s ID from the wrapper. This process is susceptible to memory leaks: a failure

between steps (i) and (iii) can result in the reference to the old payload being lost. We avoid this leak by adding another field in the wrapper, called `old_backup`, that is set to point to the old payload in `TXN_OPEN`. This update is persisted during the first persist barrier in the commit operation. `old_backup` is used to deallocate the old payload. Next, the transaction’s allocations/deallocations are all persisted. The third persist barrier is issued after all this cleanup. Then, the transaction updates its state to IDLE and persists it using a fourth persist barrier. This ensures that no further cleanup is needed. Finally, we clear the transaction’s ID from all the objects to which it wrote. If a transaction aborts, only the last two clean up related persist barriers are needed for correct rollback.

5 Persistent Memory Management

Memory management is a foundational tier in any software stack. We anticipate that applications using transactions to access persistent data will routinely allocate and deallocate persistent objects within these transactions. Most previous work on persistent memory management focuses either on wear-leveling [24] or techniques for correct allocation that tolerates failures [5, 8, 29], disregarding the overhead due to persist barriers. Volos et al. [36] present an algorithm that effectively eliminates persist barriers for memory allocation/deallocation calls within a transaction. But that works only in their redo logging transactions. Our algorithm is similar in nature, but works with all of our transaction runtimes.

We build our algorithm on previous approaches that separate the allocator’s metadata in persistent and non-persistent halves [29, 36]. Our allocator is modeled after the Hoard allocator [4], where the heap is divided in shared and thread-private superblocks. Each superblock, hosted in persistent memory, contains a persistent bitmap indicating allocation status of corresponding blocks, and non-persistent metadata (free and used lists) hosted in DRAM. A superblock is protected by a nonpersistent lock.

Each transaction maintains a persistent private *allocation log* that consists of all the allocation/deallocation requests made by the transaction. In a `pm_alloc` call, the nonpersistent metadata of the superblock is updated by the transaction and a corresponding record is appended to its allocation log. `pm_free` simply appends an entry to the allocation log.

The first persist barrier in a transaction’s commit operation persists the allocation log as well. Once the transaction persists its COMMITTED state, operations in the allocation log are reflected in the persistent metadata (bits are flipped using compare-and-swap instructions to avoid data races, and then the cache lines are written back or flushed). The post-commit cleanup phase’s first persist

barrier persists these flipped bits, and the last persist barrier marks the transaction as IDLE. Note that `pm_free` calls’ nonpersistent heap metadata (free and used lists of a superblock) is updated *after* the cleanup persists.

6 Empirical Evaluation

Our performance evaluation comprises two parts: (i) microbenchmarking, where we sweep through a comprehensive range of read/write mixes within transactions to identify performance patterns of the transaction runtimes over changing read/write proportions under varied assumptions about persist barrier latencies; and (ii) evaluation of three “real-world” applications – a new persistent key-value store we developed, a port of SQLite [31] that uses our transactions to persist the database, and a persistent version of memcached [22] – that confirms our findings reported in the micrbenchmarking part.

We conducted all our experiments on Intel’s Software Emulation Platform [29, 38]. This emulator hosts a dual socket 16-core processor, with 512GB of DRAM. 384GB of that DRAM is configured as “persistent memory” and 128GB acts as regular memory. Persistent memory is accessible to applications via `mmap`ing files hosted in the PMFS instance [29] installed in the emulator.

The emulator emulates the `clflush-opt` and `pcommit` instructions. The `clflush-opt` is implemented using the `clflush` instruction. Since `clflush-opt` evicts the target cache line, we expect it to lead to significant increase in cache miss rates, thereby degrading application performance. We therefore focus our evaluation largely on the `clwb` instruction, which does not evict the target cache line, and is likely to be the instruction of choice for applications on Intel platforms. We emulate `clwb` behavior with a `nop`; the actual latency of persisting the data is incurred by the subsequent persist barrier, which we emulate by using the emulator’s `pcommit` instruction. We note that Intel recently announced deprecation of the `pcommit` instruction [17] from future Intel processors; the persist barrier in that case is simply an `sfc` instruction (see Table 1). However, in the emulator, the `pcommit` instruction simply stalls the calling thread for a configurable amount of time [29, 38], which lets us experiment with different latencies of the real persist barrier – it will be difficult to know the real latency of a persist barrier until real hardware becomes available, so any evaluation needs to target a broad range of latencies, which we do in our experiments. Latency of loads from persistent memory is a configurable parameter as well. `store` latency in the emulator is the same as DRAM store latency.

We conducted experiments over a wide range of latency parameters for persist barriers (0 – 1000 nanoseconds) and loads (100 – 500 nanoseconds). We report results for

a load latency of 300 nanoseconds, and 3 different persist barrier latencies – 0, 100 and 500 nanoseconds, labeled as PDOM-2, PDOM-1, and PDOM-0 respectively to map them to the different persistence domains from our taxonomy. These latencies represent the overall performance trends we observed over the broader range of latencies.

6.1 Transaction Latency

Our latency microbenchmarking focuses on understanding performance of the transaction runtimes under different read/write loads. The synthetic Array microbenchmark developed by the SoftWrAP work [10] suffices this purpose. Array contains a 2-dimensional array of 64-bit integers hosted in persistent memory. The first dimension contains 10 million slots (each is an array of integers); the second dimension’s (slot’s) size is configurable, we vary it from 1 (8 bytes) to 64 (512 bytes) entries. Array continuously runs transactions, each of which randomly accesses a contiguous set of 20 slots. The slot sizes cover a broad range of access granularities. Each slot access can be a simple read of all the integers in a slot, or a read-write that increments all integers in the slot. We vary the number of slots accessed in read-only or read-write mode for different test runs. In addition, we implemented two versions of slot writes: (i) a “one shot” update version, called Array, where the transaction copies the slot integers in a private (nonpersistent) buffer, increments the integers in that buffer and then writes it back to the persistent slot; and (ii) a “read-after-write intensive” version, called Array-RAW, where each integer in the slot is individually incremented “in-place”. The second version helps us understand overheads related to read-after-write accesses in the redo logging transaction runtime. We report results as the mean of three 10-second test runs preceded by a 10-second warmup phase (we observed less than 5% deviation from the mean in all results). Array microbenchmark is single threaded and is sufficient for latency measurements.

Figure 5 shows latency results of our experiments for slot sizes of 4 and 64 (they capture the performance patterns of the configurations with other slot sizes we tested), and over the three different persist barrier latencies discussed above. Latency graphs for Array appear in Figure 5(a)–(f). The first takeaway of these is that COW performs worst across the board. In fact the margin grows from 2X to 50X compared to the best performing alternative when we move from 4-integer slots to 64-integer slots. The overheads of COW are largely related to worse cache locality (proportionally high cache miss rates) compared to the undo and redo logging alternatives – (i) the extra level of indirection, and (ii) the constant cloning of objects as they are updated. This gets worse with increasing

granularity of objects.

As expected, undo logging performance degrades with increasing persist barrier latencies. This impact is most noticeable for PDOM-0, where the persist barrier overheads are high – the barrier overhead in undo logging tends to increase the latency gap (up to 2X) with redo logging as the percentage of writes per transaction increases (Figure 5(a),(d)). The same behavior manifests in PDOM-1 configurations (Figure 5(b),(e)), albeit at a lower scale (up to 20% higher latency than redo logging), since the persist barrier latencies are lower.

For PDOM-2 however, the persist barrier is a nop. This shows in Figure 5(c), where the slot size is 4. Undo logging either performs as well, or better than redo logging. On further observation, we realized that redo logging actually performs increasingly worse than undo logging as the write percentage grows (up to 25% worse at 100% writes). Furthermore, at 0% writes, redo and undo logging are comparable in performance. This implies that the source of overheads is in the writes done by redo logging transactions. We determined that the operations related to maintaining the lookup structure for accelerating read-after-write lookups was the source of these overheads. Figure 5(f) presents a contrasting result, where, in spite of zero persist barrier latency, undo logging performs increasingly worse than redo logging as the write percentage per transaction grows. We determined that the overheads in undo logging were related to the checksum computations we needed to avoid an extra persist barrier per undo log append. Eliminating the checksum brought undo and redo logging performance at parity in these test runs. This indicates an interesting trade off in performance of undo logging implementations, where it is best to use checksums for transactions that do fine grain writes, whereas it may be best to use 2 persist barriers per undo log append for transactions that make coarse grained updates.

Notice that for read-only test runs (leftmost points in these graphs), redo logging is either at parity with undo logging, or slightly better (Figure 5 (a)). Our persistent object header checks to detect read-after-write scenarios appear to have no effect on redo logs performance. In the absence of real read-after-write scenarios, all the overhead in redo logging appears to be related to writes that do the bookkeeping needed for fast read-after-write lookups.

Figure 5(g)–(l) show performance of the transaction runtimes on Array-RAW as the percentage of read-after-write instances increases. While the performance of COW transactions remains more or less identical to their performance on Array, both redo and undo logging transactions perform relatively worse. This is directly attributable

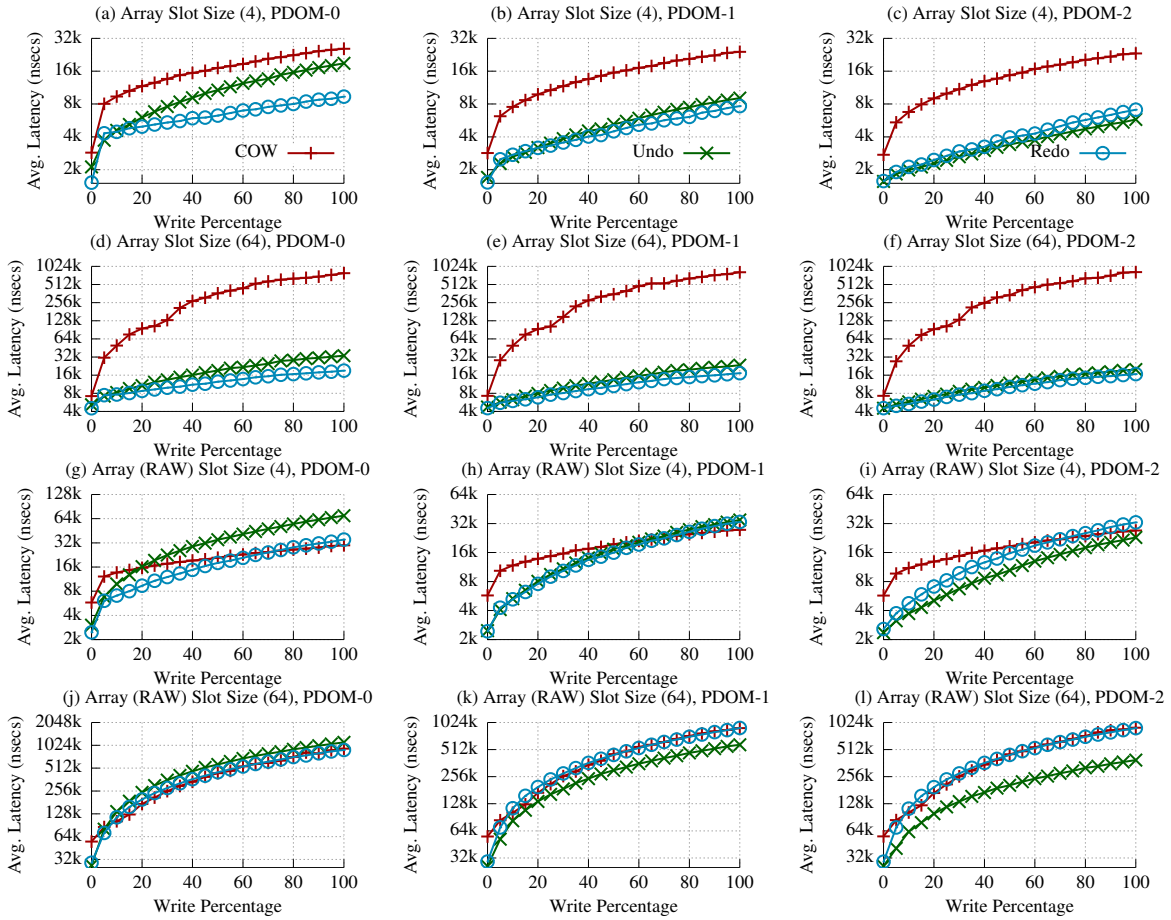


Figure 5: Average latency of transactions with increasing write percentage. Y-axes are in log scale. Each transaction accesses 20 slots. X-axes start with read-only access runs; we progressively increase read-write accesses as we move right.

to the proportional amount of churn happening on the redo/undo logs (one log record per integer increment). For PDOM-0 test runs, the high persist barrier latency combined with amplified number of persist barriers (4X or 64X) in undo logging, leads to worst performance (even in comparison with COW transactions) beyond a modest percentage of writes. Redo logging, on the other hand, incurs overheads related to read-after-write lookups (through the transaction’s lookup table and per-object update lists), which are more modest than the persist barrier overheads, but are nonetheless high enough to force redo logging perform as badly as COW transactions for a modest percentage of writes. At PDOM-1 persist barrier latencies however, these lookups turn out to be relatively more expensive, because of which undo logging performs comparably or better than redo logging. This difference increases furthermore for PDOM-2 where the persist barrier is a nop. These results affirm the overall intuition of read-after-write lookup overheads in redo logging in

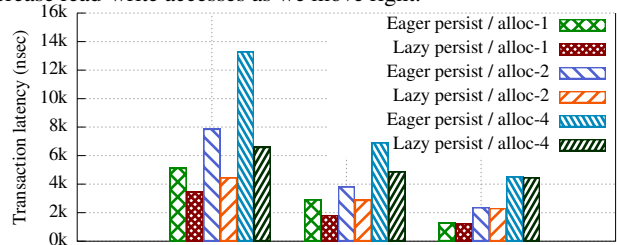


Figure 6: Latency per transaction that simply allocates the said number of blocks (Alloc-X, each of a random size from 1 to 512 bytes).

prior work [5, 6, 8, 28]. The critical question of how often do such instances arise in real world applications is something we address later in our evaluation.

6.2 Memory allocation performance.

Figure 6 shows memory allocation latency, comparing the Eager Persist approach that uses persist barriers per allocation/deallocation call, to our Lazy Persist approach that avoids persist barriers altogether during al-

location/deallocation calls. There is no performance difference in PDOM-2, because the persist barrier is a nop. However, for PDOM-1, the optimization produces a 20–30% latency improvement. In PDOM-0, the improvement grows to 30–100%, because the persist barrier latency is much higher.

6.3 Persistent Key-Value Store

We implemented a persistent key-value (K-V) store from scratch using our transactional interface. The implementation served as a vehicle to test the programmability limits of our transaction runtimes. Our K-V store’s central data structure is a concurrent, closed addressed, hash table that uses chaining to resolve hash collisions. The K-V store supports string-type keys and values, and provides a simple get/put interface. Clients connect to the K-V store via UNIX domain socket connections. The K-V store spawns a thread for each connected client (we plan to extend our implementation to let server threads handle multiple clients concurrently).

We started from an implementation that makes use of our transactional API to perform all persistent data accesses. We introduced persistent types for all the persistent data structures. This introduces “wrapper” objects for all the persistent objects hosted in our K-V store. The wrapper objects introduce a level of indirection and hence overhead. We also implemented a hand-optimized version of the K-V store that avoids use of persistent wrapper types altogether. Our optimized version also aggressively bypasses the transactional accessors wherever possible – e.g. for read accesses, we can bypass the `TXN_READ` accessors and fetch the data directly from the target address in cases where the transaction has not yet written to the address. These optimizations can be trivially supported with the undo and redo log runtimes. COW, however, appears to have a fundamental limitation – it relies on the persistent wrappers to perform the copy-on-write. As a result, it is not possible to build such a version of our K-V store using COW transactions. This is a significant limitation of the COW transaction interface.

Figure 7 shows performance of our various K-V store versions for client (hence worker) thread count ranging from 1 to 8. All client threads are bound to 1 processor socket of the emulator, while the worker threads are bound to the other socket. Due to space restrictions, we report just the PDOM-2 numbers. COW experiences cache locality overheads and performs worse than redo and undo logging, which perform comparably. We observed 12% and 14% consecutive drops in performance for both redo and undo logging when we ran the same experiments with PDOM-1 and PDOM-0 configurations. COW also experiences similar performance drops. Our

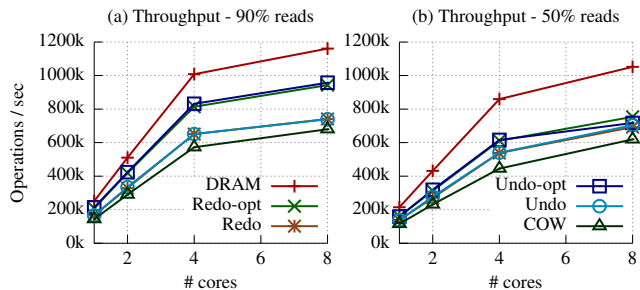


Figure 7: Persistent K-V store throughput. Suffixes “-opt” denote optimized versions. Performance is measured assuming PDOM-2. Each utilized core serves one connected client, who issues 1 million operations according to the specified read-write ratio. The keys are randomly drawn from a space of 238,328 keys, and the K-V store is pre-populated by 1 million random puts. Averages of 3 consecutive runs are reported. We omit error bars because of low variabilities in these benchmarks.

optimizations of bypassing accessors, persistent wrappers, and locality-friendly data placement deliver performance improvements of as much as 27% under read-dominated workloads. This difference comes down to 8% for redo logging and 2% for undo logging under the 50% reads workload.

Overall, we observe negligible difference between redo and undo logging versions (both base and optimized versions). The common case hash table accesses (gets and puts) are extremely short transactions, accessing a few cache lines, and updating even fewer (1–2) cache lines transactionally (e.g. linking or unlinking a node from a hash table bucket). Additionally, they do not contain any read-after-write accesses. Furthermore, the workers receive requests from clients over a TCP connection, which itself dominates the latency of client operations. This represents a possibly significant class of real world workloads, where differences in these lower level abstraction implementations may not matter to overall performance of the application.

6.4 SQLite

SQLite [31] is a popular light-weight relational database. It hosts the entire database in a single file, with another file for logging (rollback or write-ahead log) that is used to ensure ACID semantics for its transactions. SQLite can also be configured to an “in-memory” mode where the storage tier is completely removed. The database does not provide durability guarantees in that configuration. We have extended this configuration to use our transactional API for persistence.

SQLite stages all changes to the database at a page granularity in a *page cache*, and writes them out to the database file at the commit phase. For in-memory databases, the page cache is still populated, but it does not get persisted during the commit. We built a persistent

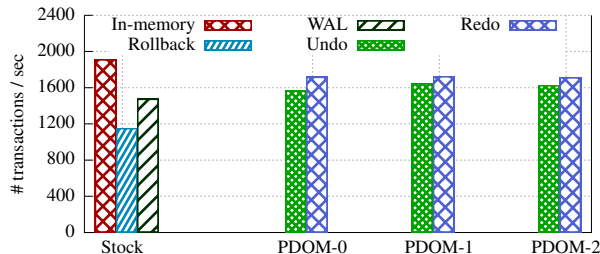


Figure 8: TPC-C, as implemented by PyTpcc [34], results on SQLite. For each persistence domain, we present both undo and redo logging performance. In-memory databases and stock databases – default unmodified SQLite with files hosted in PMFS – are presented for comparison. In our transactional version, all transactions are write-only (no reads). version of in-memory SQLite by creating a “region file” based on our persistent region abstraction that is written to, and persisted, during the commit phase using our transactional API. The dirty list is essentially applied in a single transaction, thus making the transaction’s effects durable in a failure atomic way. Our transaction effectively plays the role of rollback or write-ahead logs in the stock SQLite configuration.

Since the region file is a single large object, COW transactions would entail prohibitively high overheads, which is why we did not implement SQLite with COW transactions. We however tested our SQLite port with both undo and redo log transactions, and compared them with the in-memory configuration, and default SQLite databases whose files – both the database file, and journal file – were hosted in PMFS (we tested memory-mapped files, which is a feature supported by SQLite, but the results were identical to the default SQLite configurations). Our modifications and tests were carried out on SQLite 3.13.0.

On all mentioned configurations, we ran the TPC-C [34] benchmark, as implemented by PyTpcc [34]. Figure 8 shows the results. For each configuration, we took the average of 3 runs. As expected, the in-memory version has highest throughput. All our redo logging transactions have comparable throughput, just 10% under the in-memory version. Undo logging transactions in PDOM-1 and PDOM-2 configurations have 2-3% lower throughput than the redo logging ones; all this overhead is attributable to the checksum overhead incurred for the page-granular transaction writes. Furthermore, the PDOM-0 throughput of undo logging is 3% lower than the throughput of undo logging transactions in PDOM-1 and PDOM-2; this is where the 500 nanosecond persist barrier latency shows up. This performance nicely tracks the performance we observed in the Array microbenchmark with slot size of 64. We expect such performance pattern to emerge in applications that do lots of coarse grain writes. The unmodified SQLite with a rollback journal on PMFS are

about 36% to 50% slower than the in-memory database.

6.5 Memcached

We have used our transactional API to “persistify” memcached [22], a widely used, high performance, in-memory key-value store. The motivation for building a persistent version of memcached is to accelerate the shutdown-restart-warmup cycle, which can take several hours [11] in some instances because memcached is a nonpersistent key-value store, hosted in DRAM. A persistent memcached can significantly accelerate the warmup time. However, the cache’s state must be correctly persisted across failure events such as power failures, kernel crashes, etc. To that end, we have ported memcached to our transactional API.

We originally started the effort with the goal to simply port memcached’s central data structure, a concurrent, growable, closed addressed hash table. However, we quickly realized that we would need to modified other parts of memcached (LRU cache management, slab allocator, lazy memory reclamation, etc.) to persist memcached’s entire internal state for warm restart. As a result, this simple effort evolved into a major port of memcached to our transactional API. Transactions ended up encompassing fairly complex code paths that led to some interesting scenarios such as: tiny critical sections within transactions (supported using deferred operations and lock releases), semantically nonpersistent data located within persistent data objects (supported using persistent generation numbers), memory leaks caused by failures because memcached tracks nonpersistent references to key-value pairs, which are now persistent, using reference counts (addressed by persistently tracking key-value pairs, with non-zero reference counts, that are removed from the hash table), etc.

Additionally, the COW transaction runtime system relies on our optimized allocator, and does not work with memcached’s slab allocator. Furthermore, each key-value pair contains various fields protected by different locks so as to allow multiple threads to concurrently modify those fields. Our COW implementation allows only one writer at any given time. All these considerations lead us to a significantly different data layout, so we decided not to develop a COW based implementation of persistent memcached. We report results for our undo and redo logging implementations.

To evaluate memcached we ran the mutilate workload [3] fixing the number of client threads to 8. We varied the number of memcached worker threads from 1 to 8. Figure 9 shows our persistent memcached’s performance with 90/10% and 50/50% get/put ratios. First, note that, for all thread counts, at 10% puts, the best perform-

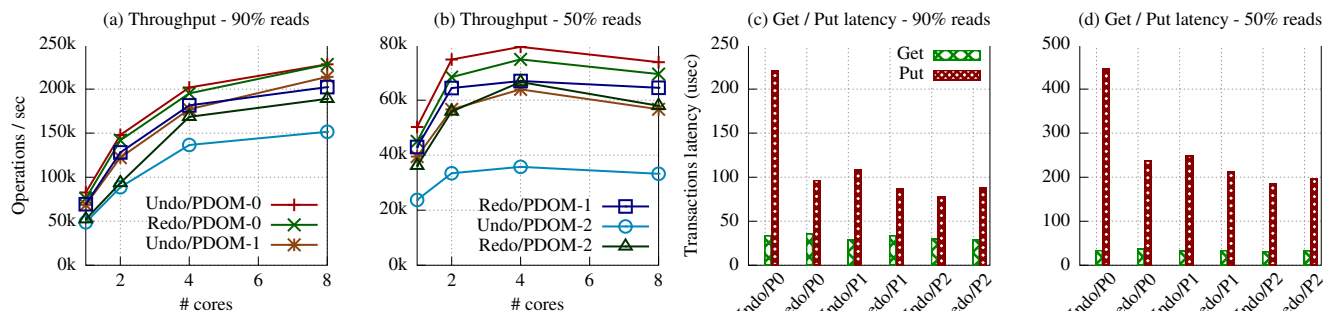


Figure 9: memcached scalability and latency results with 90/10% and 50/50% get/put ratios. Latency bar charts are for the 8 thread runs. (P0 = PDOM-0, P1 = PDOM-1, P2 = PDOM-2.)

ing runtime Undo/PDOM-0 has about 10–30% lower throughput than the original memcached, whereas the same runtime has about 45–60% lower throughput than the original memcached for 50% puts. This largely highlights the instrumentation and bookkeeping overheads of our transaction runtimes for transactional reads and writes.

Similar to our microbenchmark results, undo logging performance is better than redo logging for PDOM-2 since persist barriers are nops, and redo logging does appear to incur extra overheads for read related lookups in the redo log. However, redo logging catches up with undo logging on the read dominated, 10% put, test for PDOM-1, while slightly outperforming undo logging in the 50% put test. PDOM-0 results are more interesting. While they perform comparably at low thread counts for the read-dominated test, undo logging does not scale as well as redo logging. The explanation appears in the corresponding latency bar chart in Figure 9(c), where we clearly see the latency of put operations go up significantly for undo logging (over 200 microseconds). Our transactions end up inflating some of the critical sections of memcached. The higher latency of puts leads to greater lock hold intervals, which in turn hinders scalability. In particular, the lock that experiences contention is the slab allocator lock that protects the entire memcached heap. Every put allocates one key-value pair, and may deallocate an older key-value pair that is removed from the hash table (the key-value pair updates are done using the copy-on-write idiom). Once the slab allocator lock is acquired it needs to be released at the end of the enclosing transaction. Since puts are so slow with undo logging, the lock hold times are much greater, leading to higher lock contention and lower scalability. With 50% puts, for PDOM-0, the higher write latency (see Figure 9(d)) leads to significant slowdown in undo logging transactions at all thread counts.

7 Related Work

While most early work on persistent memory transaction runtimes ignores persist barrier overheads [5, 6, 8, 28], a growing number of efforts [10, 20, 21, 36] is addressing the problem in different ways.

Volos et al. [36] implement redo logging transactions in their Mnemosyne runtime. Giles et al. [10] go further with a “lazy” cleanup proposal that moves redo log application and cleanup to a background thread. They additionally add a DRAM-based aliasing mechanism to permit fast access to the committed data resident in redo logs. Our evaluation in ?? suggests that it may not translate to performance gains in practice due to the extra cache pressure. Lu et al. [21] optimize out Mnemosyne transactions’ last persist barrier, and propose a full processor cache flush as a technique to checkpoint committed transactions’ results. Kolli et al. [20] propose an undo logging based transaction runtime that introduces just four persist barriers per transaction, assuming that transactions know the data they need to modify in advance.

In contrast to these works that focus on one or two specific transaction runtime implementations, we perform a far more comprehensive analysis of the design space. Our analysis not only considers workload characteristics but also performance implications of persistence domains. In addition, we also consider and evaluate cache locality.

Moraru et al. [24] present a memory allocator optimized for wear leveling, which can plug easily into our allocation log technique. Volos et al. [36] presented an allocator that uses the transaction’s redo log to track persistent memory allocations. This is similar to our memory allocator. However we splice out the allocation records into a separate allocation log that lets us use it in undo and COW transaction runtimes.

Work similar to ours has emerged in the in-memory database setting [2], where the authors compare database transactions based on “in-place” updates (similar to our undo/redo logging runtimes), copy-on-write, and write-ahead logging [23] implementations. While their results

align with ours – in-place updates tend to dominate over the other two approaches – the two settings are significantly different. Their runtimes are designed to optimize database processing, with hand optimized implementations of core database data structures, whereas ours are much lower level runtimes developed to track individual loads and stores to arbitrary data structures hosted in persistent memory.

Different memory persistency models have been proposed in academia over the past few years [9, 19, 27, 30]. However, there seems to be a convergence emerging on the thread-local *epoch persistency* model [9, 27] with Intel’s recent deprecation of the `pcommit` instruction from future Intel processors. Our work applies to all these persistency models.

8 Conclusion

We presented a new taxonomy of persistence domains based on our observations of support for persistent memory in past and future systems. We also presented three novel transaction runtime systems based on undo logging, redo logging and copy-on-write implementations of transactional writes. Our runtimes are designed with the goal to reduce persist barriers needed in a transaction. Our allocator is also designed with the same goal and plugs into all our transaction runtimes. Our comprehensive evaluation of microbenchmarks and “real world” workloads suggests that while the nature of the workload does affect performance of a transaction runtime, the size of the persistence domain also has significant performance impact. At a high level, undo logging appears to perform best in all workload settings in PDOM-2 systems. It also performs best in read dominated workloads as long as there are no big write transactions. Large numbers of writes, which entail large number of persist barriers in undo log transactions, are costly for systems that have high persist barrier latencies (e.g. PDOM-0 systems in our taxonomy). Redo logging performance appears to dominate others on write-heavy workloads in PDOM-0 and PDOM-1 systems. Because of its poor cache locality, COW ends up not doing well in any of our workloads. COW runtimes also impose a significantly different (object oriented) programming model that may not align well with the application’s requirements.

References

- [1] 3D XPoint Technology Revolutionizes Storage Memory. <http://www.intel.com/content/www/us/en/architecture-and-technology/3d-xpoint-technology-animation.html>, 2015.
- [2] J. Arulraj, A. Pavlo, and S. Dullloor. Let’s talk about storage & recovery methods for non-volatile memory database systems. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, pages 707–722, 2015.
- [3] B. Atikoglu, Y. Xu, E. Frachtenberg, S. Jiang, and M. Paleczny. Workload analysis of a large-scale key-value store. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems*, pages 53–64, 2012.
- [4] E. D. Berger, K. S. McKinley, R. D. Blumofe, and P. R. Wilson. Hoard: A scalable memory allocator for multithreaded applications. In *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS IX*, pages 117–128, New York, NY, USA, 2000. ACM.
- [5] B. Bridge. Nvm-direct library. <https://github.com/oracle/nvm-direct>, 2015.
- [6] D. R. Chakrabarti, H. Boehm, and K. Bhandari. Atlas: leveraging locks for non-volatile memory consistency. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014, part of SPLASH 2014, Portland, OR, USA, October 20-24, 2014*, pages 433–452, 2014.
- [7] P. M. Chen, W. T. Ng, S. Chandra, C. Aycock, G. Rajamani, and D. Lowell. The rio file cache: Surviving operating system crashes. In *Proceedings of the Seventh International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 74–83, 1996.
- [8] J. Coburn, A. M. Caulfield, A. Akel, L. M. Grupp, R. K. Gupta, R. Jhala, and S. Swanson. NV-Heaps: Making Persistent Objects Fast and Safe with Next-generation, Non-volatile Memories. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 105–118, 2011.
- [9] J. Condit, E. B. Nightingale, C. Frost, E. Ipek, B. Lee, D. Burger, and D. Coetzee. Better i/o through byte-addressable, persistent memory. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*, pages 133–146, 2009.

- [10] E. Giles, K. Doshi, and P. J. Varman. Softwrap: A lightweight framework for transactional support of storage class memory. In *IEEE 31st Symposium on Mass Storage Systems and Technologies, MSST 2015, Santa Clara, CA, USA, May 30 - June 5, 2015*, pages 1–14, 2015.
- [11] A. Goel, B. Chopra, C. Gereia, D. Mátáni, J. Metzler, F. Ul Haq, and J. Wiener. Fast database restarts at facebook. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data, SIGMOD '14*, pages 541–549, 2014.
- [12] J. Gray and A. Reuter. *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 1992.
- [13] T. Harris, J. Larus, and R. Rajwar. *Transactional Memory, 2nd Edition*. Morgan and Claypool Publishers, 2nd edition, 2010.
- [14] M. Hosomi, H. Yamagishi, T. Yamamoto, K. Bessho, Y. Higo, K. Yamane, H. Yamada, M. Shoji, H. Hachino, C. Fukumoto, H. Nagao, and H. Kano. A novel nonvolatile memory with spin torque transfer magnetization switching: Spin-RAM. *International Electron Devices Meeting*, pages 459–462, 2005.
- [15] Y. Huai. Spin-Transfer Torque MRAM (STT-MRAM): Challenges and Prospects. *AAPPS Bulletin*, 18(6):33–40, 2008.
- [16] Intel Architecture Instruction Set Extensions Programming Reference. <https://software.intel.com/sites/default/files/managed/0d/53/319435-022.pdf>, 2015.
- [17] NVDIMM Block Window Driver Writer’s Guide. http://pmem.io/documents/NVDIMM_DriverWritersGuide-July-2016.pdf, 2016.
- [18] J. Izraelevitz, T. Kelly, and A. Kolli. Failure-atomic persistent memory updates via justdo logging. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 427–442, 2016.
- [19] A. Joshi, V. Nagarajan, M. Cintra, and S. Viglas. Efficient persist barriers for multicores. In *Proceedings of the 48th International Symposium on Microarchitecture, MICRO-48*, pages 660–671, 2015.
- [20] A. Kolli, S. Pelley, A. Saidi, P. M. Chen, and T. F. Wenisch. High-Performance Transactions for Persistent Memories. In *21th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2016.
- [21] Y. Lu, J. Shu, and L. Sun. Blurred Persistence: Efficient Transactions in Persistent Memory. *ACM Transactions on Storage*, 12(1):3:1–3:29, 2016.
- [22] Memcached – a distributed memory object caching system. <https://memcached.org/>.
- [23] C. Mohan, D. Haderle, B. Lindsay, H. Pirahesh, and P. Schwarz. Aries: A transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging. *ACM Transactions on Database Systems*, 17(1):94–162, 1992.
- [24] I. Moraru, D. G. Andersen, M. Kaminsky, N. Tolia, P. Ranganathan, and N. Binkert. Consistent, Durable, and Safe Memory Management for Byte-addressable Non Volatile Main Memory. In *Proceedings of the First ACM SIGOPS Conference on Timely Results in Operating Systems*, pages 1:1–1:17, 2013.
- [25] D. Narayanan and O. Hodson. Whole System Persistence. In *Proceedings of the 17th International Conference on Architectural Support for Programming Languages and Operating Systems*, 2012.
- [26] F. Nawab, D. R. Chakrabarti, T. Kelly, and C. B. M. III. Procrastination beats prevention: Timely sufficient persistence for efficient crash resilience. In *Proceedings of the 18th International Conference on Extending Database Technology, EDBT 2015*, pages 689–694, 2015.
- [27] S. Pelley, P. M. Chen, and T. F. Wenisch. Memory persistency. In *ACM/IEEE 41st International Symposium on Computer Architecture, ISCA 2014, Minneapolis, MN, USA, June 14-18, 2014*, pages 265–276, 2014.
- [28] pmem.io: Persistent Memory Programming. <http://pmem.io/>, 2015.
- [29] D. S. Rao, S. Kumar, A. Keshavamurthy, P. Lantz, D. Reddy, R. Sankaran, and J. Jackson. System software for persistent memory. In *Ninth Eurosys Conference 2014, EuroSys 2014, Amsterdam, The Netherlands, April 13-16, 2014*, page 15, 2014.
- [30] J. Ren, J. Zhao, S. Khan, J. Choi, Y. Wu, and O. Mutlu. Thynvm: Enabling software-transparent crash consistency in persistent memory systems. In *Proceedings of the 48th International Symposium on Microarchitecture*, pages 672–685, 2015.

- [31] SQLite. <https://www.sqlite.org/>.
- [32] D. B. Strukov, G. S. Snider, D. R. Stewart, and R. S. Williams. The missing Memristor found. *Nature*, 453:80–83, 2008.
- [33] The SNIA NVM Programming Technical Working Group. NVM Programming Model (Version 1.0.0 Revision 10), Working Draft. http://snia.org/sites/default/files/NVMProgrammingModel_v1r10DRAFT.pdf, 2013.
- [34] TPC-C, <https://www.tpc.org/tpcc>.
- [35] H. Volos, S. Nalli, S. Panneerselvam, V. Varadara-
jan, P. Saxena, and M. M. Swift. Aerie: flexible file-
system interfaces to storage-class memory. In *Ninth
Eurosys Conference 2014, EuroSys 2014, Amster-
dam, The Netherlands, April 13-16, 2014*, page 14,
2014.
- [36] H. Volos, A. J. Tack, and M. M. Swift. Mnemosyne:
lightweight persistent memory. In *Proceedings of
the 16th International Conference on Architectural
Support for Programming Languages and Operating
Systems*, pages 91–104, 2011.
- [37] M. Wilcox. DAX: Page cache bypass for filesystems
on memory storage. <http://lwn.net/Articles/618064/>,
2014.
- [38] Y. Zhang and S. Swanson. A study of application
performance with non-volatile main memory. In
*IEEE 31st Symposium on Mass Storage Systems and
Technologies, MSST 2015, Santa Clara, CA, USA,
May 30 - June 5, 2015*, pages 1–10, 2015.