THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

DHTM: Durable Hardware Transactional Memory

Citation for published version:

Joshi, A, Nagarajan, V, Cintra, M & Viglas, S 2018, DHTM: Durable Hardware Transactional Memory. in *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*. Institute of Electrical and Electronics Engineers (IEEE), Los Angeles, California, USA, pp. 452-465, The 45th International Symposium on Computer Architecture, Los Angeles, California, United States, 2/06/18. <https://doi.org/10.1109/ISCA.2018.00045>

Digital Object Identifier (DOI):

[10.1109/ISCA.2018.00045](https://doi.org/10.1109/ISCA.2018.00045)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

DHTM: Durable Hardware Transactional Memory

Arpit Joshi
University of Edinburgh
arpit.joshi@ed.ac.uk

Vijay Nagarajan
University of Edinburgh
vijay.nagarajan@ed.ac.uk

Marcelo Cintra
Intel, Germany
marcelo.cintra@intel.com

Stratis Viglas
Google
sviglas@google.com

Abstract—The emergence of byte-addressable persistent (non-volatile) memory provides a low latency and high bandwidth path to durability. However, programmers need guarantees on what will remain in persistent memory in the event of a system crash. A widely accepted model for crash consistent programming is ACID transactions, in which updates within a transaction are made visible as well as durable in an atomic manner. However, existing software based proposals suffer from significant performance overheads.

In this paper, we support both atomic visibility and durability in hardware. We propose DHTM (durable hardware transactional memory) that leverages a commercial HTM to provide atomic visibility and extends it with hardware support for redo logging to provide atomic durability. Furthermore, we leverage the same logging infrastructure to extend the supported transaction size (from being L1-limited to LLC-limited) with only minor changes to the coherence protocol. Our evaluation shows that DHTM outperforms the state-of-the-art by an average of 21% to 25% on TATP, TPC-C and a set of microbenchmarks. We believe DHTM is the first complete and practical hardware based solution for ACID transactions that has the potential to significantly ease the burden of crash consistent programming.

I. INTRODUCTION

The emergence of byte-addressable non-volatile memory technologies [1], [2], [3], [4], also known as *persistent memory*, is fast blurring the divide between memory and storage. Being directly attached to the memory bus, persistent memory provides a high-bandwidth and low-latency alternative for durability. However, merely providing a fast non-volatile medium will not suffice. Programmers need guarantees about what will remain in persistent memory upon a crash or a failure.

In a recent study, Marathe et al. [5] highlight the numerous challenges in designing crash consistent programs and advocate for systematic programming models such as transactions that provide ACID guarantees (Atomicity, Consistency, Isolation and Durability). ACID essentially implies that updates within a transaction are made visible (to other transactions) as well as durable (to a non-volatile medium), in an atomic manner. While the database community has developed a plethora of techniques to guarantee ACID efficiently, these techniques have predominantly been developed with slow block-based media in mind. When applied to in-memory settings, such techniques tend to spend a significant amount of time on concurrency control [6], [7], [8] and logging [6],

[9], [10]. This leads us to ask the question: How fast can we enforce ACID in the presence of fast persistent memory?

Related Work. Recently, there have been multiple proposals for providing ACID updates to persistent memory. These proposals are classified in Table I based on how they enforce atomic visibility and atomic durability. The first class of designs [11], [12], [13], [14] support atomic durability via software logging by employing flushing and ordering instructions. Ensuring atomic durability in software, however, comes at a significant performance cost [15], [16], [17], [18], [19] which motivated the development of the second class of designs that either employ hardware support for atomic durability [15], [17], [20], [21], [22], [23] or leverage hardware support for ordering to guarantee atomic durability [16], [18], [24], [25]. However, both of these classes enforce atomic visibility in software using software transactional memory (STM) or locks.

Another approach to ACID is to leverage commercially available Hardware Transactional Memory (HTM) to support atomic visibility, which is the focus of the remaining classes of designs. However, current commercially available HTM systems have two limitations. First, they efficiently support only small transactions [26], [27], [28], [29], [30]; if a cache line written within a transaction is evicted from the L1 cache, the transaction must abort. The severity of the problem has been highlighted by a recent study which finds that transactions whose write-set size is larger than 128 cache lines (quarter of the L1 size) are highly likely to abort [31]. This L1 limitation can significantly limit usability and efficiency for ACID transactions, which tend to have relatively large write working-set sizes (Section V). Second, HTM systems only provide ACI guarantees, i.e., atomic visibility but not atomic durability. To guarantee ACID, the third class of designs [13], [32], [33] leverages the HTM for atomic visibility and integrates it with software support for atomic durability. The latter requires the writing of a log entry for every modified object within the transaction, thereby increasing the transaction’s write set (and the abort rate). The fourth class supports ACID by integrating HTM with hardware support for durability. However, PTM [34] (the only proposal in this class) not only introduces significant changes to the cache hierarchy, but also continues to suffer from the L1 limitation.

Our Approach. Our primary goal is to design an HTM that

Designs	Atomic Visibility	Atomic Durability	Trans. Size	LLC Extensions
Atlas [11], REWIND [12], DudeTM [13], Mnemosyne [14]	Locks or STM	Software	Not limited	None
WrAP [15], DPO [16]*, LOC [17], HOPS [18]*, ATOM [20], [21], [22], Kiln [23], NVHeaps [24]*, DCT [25]*, [35]	Locks or STM	Hardware	Not limited	[15], [17], [21], [23], [24], [35]
DudeTM [13], PHyTM [32], cc-HTM [33], [36]	HTM	Software	L1 limited	None
PTM [34]	HTM	Hardware	L1 limited	Yes
DHTM	HTM	Hardware	LLC Limited	None

Table I: Classification of techniques supporting ACID updates on persistent memory. (* Leverage hardware support for ordering to provide atomic durability.)

can support ACID transactions efficiently. A secondary goal is to extend the supported transaction size by supporting overflows from the L1 cache to the last level cache (LLC) without adding significant complexity to the coherence protocol or the LLC.

One way of achieving these goals is to leverage existing unbounded HTM designs [37], [38], [39] that rely on logging to support overflows and make those logs durable [20]. However, such an approach, where durability is treated as a secondary consideration, will have poor performance as persisting the log and/or the data will be in the critical path.

We advocate an alternative approach in which durability is a first class design constraint. We propose Durable Hardware Transactional Memory (DHTM) in which we integrate a commercial HTM like RTM [28] with hardware support for redo logging. DHTM achieves atomic visibility by leveraging RTM. Whereas for achieving durability, DHTM provides architectural support for transparently and efficiently writing redo log entries to a durable transaction log maintained in persistent memory; the key efficiency enabler here is our novel mechanism for collating and flushing log entries without consuming excessive memory bandwidth. The redo log based design allows us to commit a transaction as soon as all the log entries have been written to persistent memory, without waiting for data to be made durable. DHTM then extends the supported transaction size, by leveraging the same logging infrastructure for also supporting L1 overflows. When the write set of a transaction overflows from the L1 cache, DHTM logs the address of the overflowed cache line and leverages the log to commit (or abort) the transaction. DHTM supports this with minor changes to the coherence protocol and without adding any additional transaction tracking hardware to the LLC. In summary, our key contributions are:

- We propose DHTM, the first complete hardware solution for an ACID compliant transactional memory system which is not bound by the size of the L1 cache.
- We enforce ACID efficiently by leveraging RTM [28] for atomic visibility and by providing atomic durability via hardware support for redo logging. We also propose a mechanism for coalescing log entries to reduce the required memory bandwidth.

- We extend the supported transaction size by allowing for the transaction’s write set to overflow from the L1 to the LLC by leveraging the same logging infrastructure for handling these overflows. We accomplish this with only minor changes to coherence protocol.
- Our evaluation shows that DHTM outperforms the state-of-the-art [20] by 21% to 25% on average across TATP, TPC-C and a set of micro-benchmarks.

II. BACKGROUND

A. Hardware Transactional Memory

From idea inception [40] to mainstream commercial adoption, HTMs have come a long way. Although prior work has explored unbounded transactions, current commercial HTMs predominantly provide only a best effort service, with transaction sizes being limited by the size and associativity of the L1 cache. Below, we briefly describe an HTM system which is similar to state-of-the-art commercial HTM designs [26], [27], [29] and is specifically modelled on Intel’s *Restricted Transactional Memory* (RTM) [28] design. Later we briefly discuss designs that support overflow from private caches. For a broader perspective, the reader is referred to Harris et al.’s book [41].

Commercial HTMs. HTMs primarily provide support for three functionalities: buffering the speculative state, tracking read and write sets and detecting conflicts. Commercial HTMs typically buffer speculative state in private caches (typically L1). Each L1 cache line is associated with a *write bit* to keep track of the write set of a transaction. If a cache line belonging to the write set of a transaction is evicted from the L1, the transaction is aborted. Thus, the supported write-set size is limited by the size and the associativity of the L1 cache. Commercial HTMs avoid supporting overflows from the private L1 caches to reduce the design complexity, and in particular that of the LLC.

Similar to the write bit, a *read bit* is also associated with each cache line in the L1 cache. This bit is set when the corresponding cache line is read within a transaction. When such a cache line is evicted, the transaction is typically not aborted, but the address of the cache line is added to a *read-set overflow signature* (also maintained in the L1 cache).

Thus, the read set of a transaction is tracked using both the read bits in the L1 cache and the read-set overflow signature.

Conflict detection happens at the L1 cache, with help from the cache coherence substrate. Specifically, when the L1 receives an invalidate request for a cache line in the read set, or an invalidate/data forwarding request for a cache line in the write set, a conflict is detected, triggering an abort of one of the transactions. What transaction must abort is determined by the conflict resolution policy. Two of the commonly used policies are the requester wins policy [28] and the (first) writer wins policy [29].

Overflow Support. Multiple techniques [37], [38], [39] have been proposed to support write set overflows from private caches. Techniques with lazy version management [38], [39] allow the write set to overflow into a redo log. On a commit, these values need to be copied in-place. Consequently, these techniques stall any transaction that conflicts with a committed transaction that is still copying its updates in-place. Techniques with eager version management, such as LogTM [37], allow the write set to overflow in-place in memory but maintain an undo log that is applied in case of an abort. Therefore, they have to stall transactions that conflict with an aborting transaction that is applying its undo log. Stalling adds significant design complexity as it requires support for retrying requests using a NACK based coherence protocol. Our goal with DHTM is to support overflows from the L1 cache to the LLC while maintaining the simplicity of an RTM like protocol (§III-C).

DHTM performs data updates in the cache and eager conflict detection in the same way as RTM. However, it additionally maintains a redo log in memory for atomic durability and also supports write set overflows to the LLC with minor modifications to the coherence protocol.

B. Crash Consistency

Storage systems [42] and more recently systems with persistent memory [12], [14], [24] have employed write-ahead logging to provide crash consistency. Write-ahead logging operates on the principle that the log entries be made persistent before data values can persist. Software implementations rely on instructions such as *non-temporal store*, *cache-line write-back*, *sfence* and *pcommit* to ensure the required ordering, but suffer from a significant performance overhead in the process [15], [16], [17], [18], [19]. To mitigate this overhead, prior work has proposed hardware support for accelerating ordering [15], [16], [18], [25], [35] and techniques for transparently performing logging in hardware [17], [20], [21]. With the former, programmers still need to insert appropriate barriers between log writes and data writes. In contrast, with hardware logging, the programmer is relieved from the burden of writing log entries. Instead, an interface is provided to demarcate the region of code that needs to execute in a crash consistent manner. The

Figure 1: Working set sizes for transactions (a) without including durability log and (b) with durability log.

hardware ensures that log entries are transparently written to persistent memory in the correct order.

III. DHTM DESIGN

In this section we present the design of our *durable HTM* (DHTM), that adds support for durability on top of a commercial RTM-like HTM design.

System Model. For the following discussion, we assume a multicore processor with a two level cache hierarchy consisting of private L1 caches and a shared last level cache (LLC). The private L1s are kept coherent using a MESI directory based coherence protocol with forwarding (similar to the one in section 8.2 in [43]). We assume the directory is held in the LLC with each cache line maintaining the coherence state and sharing vector. We assume a baseline HTM similar to Intel’s RTM. We assume that the HTM supports strong isolation. Finally, we assume that memory is non-volatile and byte addressable. It is worth noting that the above model is mainly to help anchor our description and as such, none of these choices are fundamental to DHTM.

Overview. At a conceptual level, adding durability to an HTM requires some form of logging. Figure 1a shows a volatile transaction at the top and the corresponding read and write sets at the bottom. The transaction reads *X* (read set) and writes to *X* and *Y* (write set). One way for this transaction to be made durable is by executing the code sequence shown on the top in Figure 1b, which additionally writes log entries for the data being modified. The resultant read set of the persistent transaction remains the same, but the write set consists of *Log X*, *X*, *Log Y* and *Y*. Thus, adding support for durability essentially doubles the write-set size of transactions. This is a challenge on current RTM-like HTM designs which already limit the write-set size. To compound matters, applications that demand ACID tend to have relatively large transaction sizes. Therefore, one of our goals is to support transactions with a larger write-set size relative to those supported by current commercial HTMs. But in the quest for larger transactions, we do not want to introduce significant hardware complexity; in particular, we do not want to introduce changes to the shared LLC like

adding transaction tracking hardware or searching the LLC for cache lines belonging to the write set – something that current HTM designs avoid.

Our approach is to integrate hardware based redo logging to an RTM-like HTM. For atomic visibility, DHTM leverages the RTM-like HTM and for atomic durability, it employs hardware redo logging. Since logging is performed transparently, DHTM’s programming interface is similar to that of volatile transactions (Figure 1a). DHTM’s redo logging mechanism leverages the L1 cache write-back interface to dynamically write redo log entries to persistent memory for cache lines being modified within a transaction. Furthermore, DHTM allows dirty cache lines to overflow from the L1 cache into the LLC without causing an abort. This increases the transaction size with minor changes to the coherence protocol and without adding significant design complexity (in particular, without adding transaction tracking hardware to the LLC). Below, we first describe DHTM’s hardware logging mechanism. Then, we describe how logging integrates with the HTM, followed by the description on how DHTM manages overflow.

A. Logging for Durability

We ensure atomic durability using *write-ahead logging*. The idea is to maintain a persistent copy of the old and new versions at all times during the transaction, so that the state can be recovered to either of the versions. This persistent copy is maintained in the form of *log entries* which consist of the address and the old or new version of data. In this section we provide a design for a redo-log based implementation to work in conjunction with HTM.

Why Redo Logging? We choose a redo-log based design as it allows us to have both fast commits as well as fast aborts. In volatile transactions, undo logging supports faster commits because, on transaction completion all the in-place updates would have already taken place (in the cache); commit therefore only requires two simple steps: discarding the undo log and flash-clearing the speculative write bits to make the write set visible to other threads. Durable transactions, however, impose additional constraints. Both the undo log and the write set (data) have to be written to persistent memory – only then, can the transaction be committed. While techniques have been proposed for minimizing the fine grained ordering overheads while writing log entries [20], flushing the write set can significantly increase commit time.

Redo logging, in contrast, requires only the redo log to be written to persistent memory at commit time. This is because the redo log, in addition to serving as a recovery log in case of a failure, can also provide the up-to-date values on commit. This allows for the data updates to be written to persistent memory in the background, and out of the commit critical path. One traditional drawback of redo logging is that, because writes are not allowed to overwrite

previous values, subsequent reads to those addresses need to be redirected to the redo log. Our proposed hardware based redo-logging mechanism overcomes this limitation by allowing writes to overwrite previous values in the cache. A subsequent read can therefore directly read the updated value from the cache. It is worth noting, however, that the writes do not overwrite the old values in memory but are written to a separate log area. Lastly, aborts are also faster with redo logging and only require two simple steps: discarding the redo log and invalidating the modified lines in the cache.

Log management. In the DHTM design, the transaction log space is thread private and is allocated by the operating system (OS) when the thread is spawned. The OS keeps track of all the logs it has allocated so that it can recover transactions from logs in case of a system crash. This per thread transaction log is organized as a circular log buffer similar to Mnemosyne [14]. On a log overflow, DHTM aborts the transaction with an indication that the abort is because of log overflow. The OS in this case allocates a larger log space for the thread and the transaction is retried.

Hardware Support. One of the design goals of DHTM is to write log entries to the transaction log in persistent memory without adding them to the write set. To this end, logging is performed in hardware in DHTM, allowing DHTM to differentiate between log writes and data writes. The L1 cache controller is modified to enable it to write log entries to persistent memory by bypassing the LLC as shown in the Figure 2a. The L1 cache controller creates these log entries on the fly at a word granularity for every store within a transaction. Figure 2b shows the log writes that the L1 cache controller performs.

Log coalescing. Writing a word-granular redo-log entry for every store can generate a large number of log entries which can consume significant amounts of memory write bandwidth. Figure 2b highlights this with an example. Let us assume that each cache line consists of two words (all words belonging to cache lines *A* and *B* are initially 0); the subscript for each cache line refers to the word in the cache line that is being modified. Performing word-granular logging generates 5 log writes across the memory bus for 5 store requests to different words in cache lines *A* and *B*. The bandwidth consumed can be mitigated to some extent by coalescing multiple log entries into one cache line before writing them to memory. Nonetheless, creating a log entry for every store request is problematic. Recall that each log entry is composed of the data and the address (metadata). The finer the granularity of logging, the greater the amount of metadata, which in turn translates into higher bandwidth consumption. Second, logging for every store request might miss opportunities for coalescing multiple stores to the same word via a single log entry. For example, in Figure 2b the word A_0 gets written to twice which leads to the creation of 2 log entries, however only the second log entry would have sufficed.

(b) Hardware redo log at word granularity. Each redo-log entry consists of (address, new value) pair.

(c) Hardware redo log at cache line granularity using a log buffer.

Figure 2: Redo logging in hardware.

An alternative is to perform logging at cache line granularity. But naively creating a log entry for every store request will only worsen the memory bandwidth consumption. At the same time, the final state of a cache line (at the end of a transaction) must be logged for correctness. If we can predict the final store to a cache line, that would be an opportune moment to log that cache line, since that would minimize the number of entries logged for that cache line. It is important to note that the prediction must be conservative, in that, it must not miss the last store under any circumstance.

We conservatively predict the final store to a cache line via a simple structure called log buffer that is added to the L1 cache. The log buffer is a fully associative structure with a small number of entries that keeps track of cache lines with their cache line addresses. When a store is performed, the corresponding cache line address is added to the log buffer (if not already present). A log entry is written to persistent memory only when an entry is evicted from the log buffer. An entry is evicted from the log buffer under two situations: (a) when the log buffer is full, an eviction has to happen in order to make space for a new cache line address; (b) when

an L1 cache line is replaced and the log buffer holds the corresponding address, the address is evicted from the log buffer. Thus, we use eviction from the log buffer as a proxy for predicting the last store to a cache line; in practice, this simple policy works well because write reuse distance (when there is reuse) is typically low for transactional workloads. When an entry is evicted from the log buffer, the redo-log entry for that cache line is created as usual by composing the address with the contents of that cache line from the L1 cache. Then, the redo-log entry is written to persistent memory – in doing so, the stores to one cache line are temporally coalesced, such that all these coalesced stores get only one log write. Finally, at the end of the transaction, all of the cache lines being tracked in the log buffer are logged to persistent memory. It is important to note that this log buffer is different from the log buffer used in LogTM [37]. LogTM uses a buffer to reduce the contention for the L1 cache port and to hide L1 cache miss latency whereas the buffer in DHTM is to coalesce log writes to the same cache line and to predict the last write to a cache line.

Figure 2c shows the previous example in the presence of a single entry log buffer. Initially the buffer holds cache line *A* while it is being modified. When cache line *B* has to be modified, the updated value of cache line *A* is written to the log area and the buffer now holds cache line *B*. Eventually when the transaction ends, a redo log entry for *B* is also written to the log area. In this example, 5 store requests require only 2 log writes over the memory bus.

B. Integrating Logging with HTM

In this section, we will describe how to integrate our logging mechanism with an RTM-like HTM. This section will assume that the transaction will abort on a write-set overflow from the L1; we will handle write-set overflows in the next section.

Overview. Committing a volatile transaction requires that the read-/write-set tracking structures be cleared and that the speculative state be made visible to other threads. In addition to the above steps, in order to commit a durable transaction (with redo logging), the redo-log entries must be written to persistent memory. Recall that the data updates (write set) can be written to persistent memory lazily and out of the commit critical path. Conflict detection works identically to a volatile transaction. Non-transactional accesses also safely integrate with DHTM, similar to RTM, by aborting an ongoing transaction if it conflicts with a non-transactional access. Aborts also are largely identical, with an added step of (logically) clearing the redo log for the transaction. Thus, a durable transaction can be expressed in the form of a state diagram as shown in Figure 3a, with the following states: *Active*, *Commit*, *Commit Complete* and *Abort Complete*. Below, we discuss these in more detail.

Commit. Upon reaching the end of the transaction, and having written all redo-log entries to persistent memory, the

Figure 3: Transaction States. A core can start executing subsequent non-transactional instructions after reaching *Commit/Abort* and can start a new transaction after reaching *Commit Complete/Abort Complete*.

transaction effectively *commits*. To mark that the transaction has committed, DHTM writes a *commit log record* to the log area. The L1 cache controller then starts writing back the cache lines belonging to the write set of the committed transaction via the cache write-back interface. DHTM does not flash clear the write bit associated with cache lines on a commit, instead it clears those bits once a write-back is issued for the corresponding cache line. After writing back all the modified cache lines to persistent memory, DHTM marks the transaction as *completed* by writing a *complete log record* to the log area. Writing a complete log record is not a correctness requirement but reduces recovery time on a failure (as we shall see in the section on recovery). Once a transaction has committed, DHTM can start executing non-transactional code following the transaction. But since DHTM has only one set of write bits per cache line, it cannot start executing a new transaction until the previous transaction has completed. This is because, in order to complete a transaction, DHTM relies on these write bits to identify the modified cache lines that need to be written back to persistent memory.

In the DHTM design, there is a window between the commit point of a transaction and its completion point (when the cache lines modified in the transaction are being written back to persistent memory and are being marked as non-speculative) during which a conflict might be detected

incorrectly. For example, consider that a transaction T_A tries to modify a cache line X . But X has already been modified by a committed but not yet complete transaction T_B , and has not yet been marked as non-speculative. In such a scenario a conflict will be detected incorrectly. DHTM sidesteps this problem by also consulting the state of the transaction during conflict detection; as the transaction status of T_B indicates that it has committed, DHTM does not raise a conflict in this situation. Additionally, DHTM inserts a sentinel log entry in the transaction log of both T_A and T_B indicating that transaction T_A is dependent on the updates of transaction T_B . This sentinel log entry enables the recovery manager to decide the correct order of replay for transactions with conflicting updates.

Abort. Aborting a volatile transaction requires that the read-/write-set tracking structures be cleared and that the speculative state be invalidated. To abort a durable transaction, in addition to the above steps, the log entries need to be cleared. DHTM logically clears the log entries by writing an abort log record, effectively marking the log entries as being part of an aborted transaction.

Recovery. At the time of failure, a durable transaction can be in one of the following states: *Active*, *Commit*, *Commit Complete*, or *Abort Complete*. The recovery manager does not have to do anything for transactions in *Active* or *Abort Complete* state as none of the updates of the transactions would have been written back in-place in persistent memory. In other words, persistent memory has the pre-transaction state for those transactions. For committed but not completed transactions (transactions in *Commit* state), the recovery manager reads the log entries and writes the updated values in-place in persistent memory, thus recovering the updates of the transaction. Finally, for completed transactions (in *Commit Complete* state) the recovery manager does not have to do anything, since all of their updates would have already been written back in-place in persistent memory. In the absence of a complete log record, the recovery manager would have had to copy all the updates from the log area to in-place in persistent memory. Thus the writing a complete log record helps reduce the recovery time.

The replay order of committed but not complete transactions does not matter as long as they do not have conflicting updates. For transactions with conflicting updates, the recovery manager infers the required replay order by looking at the sentinel log entries in the relevant transaction logs.

The recovery manager is implemented as an operating system service which is invoked upon system re-start. As described earlier, the OS keeps track of all the logs it has allocated which it also registers with the recovery manager on creation and de-registers when the the log is deallocated. When the recovery manager is invoked, it scans all the registered logs and restores all the committed but not completed transactions.

C. Handling Overflow

The design described above continues to suffer from the transaction size limitation that is typical of an RTM design. We now describe an extension that allows the write set of a transaction to overflow the L1 cache without aborting the transaction. Consistent with current commercial HTM designs, our proposed extension also does not require expensive operations at the LLC (e.g. searching for cache lines belonging to a transaction), making it amenable to commercial adoption. We first summarize the challenges in supporting overflow efficiently and then describe our approach.

Challenges. Commercial HTM designs like RTM allow the read set to overflow the L1 cache. Conflicts are detected with the help of the overflow signature maintained in the L1 cache, which tracks the addresses of cache lines that have overflowed. On a transaction commit or an abort, the overflow signature is cleared; importantly, nothing needs to be updated in the LLC. In contrast, RTM-like designs do not support write-set overflows from the L1 cache. This is because, aborting a transaction requires that the HTM invalidate all the cache lines belonging to the write set. Whereas this can be done in private L1 caches by flash invalidating the cache lines, doing this for a shared structure as large as the LLC is expensive and involves non-trivial changes (indexing and searching operations). With durable transactions, a commit would also require a similar operation at the LLC: all the cache lines that have overflowed must be identified and written back to persistent memory.

Overview. Our DHTM design allows for the write set to overflow from the L1, with minor changes to the coherence protocol and without requiring any structural changes to the LLC. Our key idea is to leverage the redo log (which holds the speculative state of the transaction) for handling write-set overflows, thus obviating the need for expensive changes to the shared LLC.

DHTM handles write-set overflow by allowing for cache lines belonging to the write set to be replaced from the L1 cache to the LLC; in order to enable conflict detection the coherence state of the cache line in the LLC is kept unchanged, however. This ensures that the LLC continues to show the cache line as being owned by the core executing the transaction. Therefore, any coherence message will continue to be forwarded to the owner's L1, wherein a potential conflict can be detected. It is worth noting that our idea of using stale coherence state for conflict detection is similar to the *sticky state* solution used in LogTM [37]. Therefore, the resulting coherence protocol extensions in DHTM are similar to the sticky state extensions of LogTM.

While maintaining stale state in the LLC helps in conflict detection, we also need a mechanism to identify all the cache lines that have overflowed from the L1 cache to the LLC for versioning. To this end, DHTM maintains an *overflow list*

along with the redo log in memory. When a dirty cache line overflows from the L1 cache to the LLC, DHTM writes the address of the overflowed cache line to the overflow list. On a commit or an abort, DHTM uses the overflow list to identify cache lines belonging to the write set that have overflowed, and writes them back to persistent memory (in case of a commit), or invalidates the corresponding LLC cache line (in case of an abort).

The recovery procedure remains the same with overflows, since the cache lines belonging to the write set that overflowed the L1 cache are already present in the redo log. In summary, a durable transaction with write-set overflows can be expressed in the form of a state diagram as shown in Figure 3b with the following states: *Active*, *Commit*, *Commit Complete*, *Abort* and *Abort Complete*. One key difference with overflows is that like commits, aborts now require a completion phase. Below, we discuss this in more detail.

Commit. After writing the commit log record to the log area and issuing write-backs for all the cache lines belonging to the write set in the L1 cache, DHTM reads the overflow list corresponding to the committing transaction (recall that the overflow list contains addresses of all the dirty cache lines that have overflowed from the L1 cache). It then sends write-back messages for those cache lines to the LLC. On receiving a write-back request, the LLC writes back the relevant cache line in-place in persistent memory and also transitions the cache line to a clean state and clears its sharer vector. After writing back all the cache lines in-place in persistent memory, DHTM writes a *complete* log record to the transaction log and then transitions the transaction status to *Commit Complete*. This completion operation ensures that the LLC correctly reflects the status of overflowed write-set cache lines belonging to the transaction and eliminates any need for LLC modifications to clear such state.

Conflict Detection. Recall that conflicts are detected at the L1 controller, by checking coherence requests against the read/write bits associated with cache lines or the read overflow signature. In order to enable conflict detection in the presence of write-set overflows, we need to ensure that coherence requests for the overflowing cache lines continue to reach the L1 cache. To this end, when a dirty block overflows from the L1, the coherence state of the LLC is kept unchanged. Specifically, when an L1 cache receives a Fwd-GetM request or a Fwd-GetS request for a cache line that is not present in the cache, DHTM infers that the request corresponds to a cache line that has overflowed from the L1. Therefore, a conflict is detected and one of the transactions is aborted based on the conflict resolution policy.

Abort. To abort a transaction, DHTM first invalidates the cache lines belonging to the write set in L1 as described in Section III-B. But additionally, the cache lines in the LLC that overflowed from the L1 will also need to be invalidated. Therefore, in the presence of overflows, abort also has a completion phase as shown in Figure 3b.

In the completion phase, DHTM reads the cache line addresses from the overflow list in the transaction log and issues invalidate requests for those cache lines to the LLC. The LLC invalidates the cache lines on receiving an invalidate message. Similar to the completion phase of commit (Section III-B), the completion phase for abort can continue in parallel with the execution of other non-transactional instructions, but a subsequent transaction cannot begin until this phase completes. However, differently from a commit, DHTM does not (need to) write an abort complete log record for an aborted transaction as the state that needs to be invalidated is in volatile caches and will anyway be lost on a system crash.

One corner case concerns cache lines that have been reread back into the L1 during the transaction, after overflowing from the L1 to the LLC. When such a transaction aborts, these reread cache lines must be identified as belonging to the write set and invalidated by the L1. Such reread cache lines are correctly identified by DHTM as follows. On an LLC read, DHTM will look at the state of the cache line and the sharer vector; if the cache line is dirty and its state is in *modified* state with the requester marked as its owner, DHTM will identify the cache line as belonging to the write set and will set the write bit in the L1. This ensures that such reread cache lines are invalidated on an abort.

It is worth noting that, as opposed to existing proposals for supporting overflow (such as LogTM) DHTM does not stall requests from other transactions. Consider the case where transaction T_A has modified cache line X which then overflowed to the LLC. T_A is subsequently aborted because of a conflict. While T_A is in the process of aborting, another transaction T_B issues a read for cache line X . Because of the eager version management of LogTM, the read for X cannot be completed until X has been reverted to a non-speculative state from the undo log. Therefore, LogTM would NACK the read request for X while waiting for the abort process of T_A to end and T_B will have to subsequently re-issue the read. This adds significant complexity to the coherence protocol. DHTM on the other hand has non-speculative data in memory because it maintains a redo log for atomic durability. Therefore, it can immediately complete the read for X by fetching it from memory. In summary, DHTM maintains the simplicity of an RTM like design while allowing for overflows from the L1 to the LLC.

IV. PUTTING IT TOGETHER

In this section we first explain through detailed examples, the life cycle of a transaction. We also quantify the overall hardware overhead and finally describe a software fallback mechanism for transactions that do not fit in DHTM.

Transaction Lifecycle. Figure 4 shows the life cycle of a transaction for both commits (e and f) and aborts (g and h). For this example, let us assume a dual-core processor. Each figure shows three views: (i) L1 view from the perspective

of core 1, showing L1 cache lines with their read/write bits, a single entry log buffer, the read overflow signature and transaction status register; (ii) LLC view, showing for each cache line, its coherence state, sharer vector and dirty bit; and (iii) Persistent memory view, showing the overflow list, log area and in-place values of cache lines in memory.

(a) Initial state. The transaction is in *Active* state and has already modified cache line A with a value of 15 and has already read cache line B . No transactional data has overflowed from the L1. The LLC has cache lines A and B owned by the core 1 in *modified* state.

(b) Write B. The transaction modifies the value of cache line B to 25. Therefore, DHTM sets the write bit for cache line B . Also, cache line A needs to be evicted from log buffer (to make space for cache line B), so its updated value of 15 is written to the redo log area.

(c) Read C. The transaction reads cache line C , because of which cache line A gets replaced from the L1 cache. Therefore, cache line A is written back to the LLC and its dirty bit is set, but the coherence state of cache line A is not changed. Also, the address of cache line A is written to the overflow list in memory. Cache line C is present in the LLC in shared state, its sharer vector is updated to add core 1 as a sharer and the cache line is brought to the L1 with its read-bit set.

(d) Write E. The transaction writes 55 to cache line E . This leads to cache line C being replaced from the L1. Before being replaced, the address of cache line C is added to the read overflow signature (but because of inherent imprecision, let us assume that the signature conservatively shows both C and D as its members). Since cache line E is not present in the cache hierarchy, it is brought to the LLC in *modified* state with the sharer vector showing core 1 as the owner. The cache line is also added to the L1 cache where it is updated and its write bit is set. Since cache line E also needs to be added to the log buffer, cache line B is removed from the buffer and its updated value is written to the transaction log.

(e) Commit. When the transaction commits, cache line E is written from the L1 to the transaction log and a commit log record is also written to the log. Simultaneously, the read bits and the read overflow signature in the L1 are cleared and the transaction state is updated to *Commit*. The transaction commits at this point and core 1 may continue executing non-transactional instructions.

(f) Commit Complete. In the commit completion stage, the L1 writes back cache lines B and E to the LLC and clears their respective write bits. On receiving the write-backs, the LLC updates them and also writes them back to persistent memory. Then the memory controller reads the overflow list and issues a write-back request to the LLC for cache line A . The LLC on receiving the request, clears the sharer vector and dirty bit for cache line A and writes it back in-place in persistent memory. Finally, a complete log record is written to the log area, the overflow list is cleared and

Figure 4: Flow of a transaction

Register	Description
Log Buffer	Tracks cache lines pending log writes
Transaction State	Identify the state of a transaction
Log Area	
Start Pointer	The start address of the log space
Next Pointer	Address to write the next log entry
Size	Size of the log space
Overflow List	
Start Pointer	The start address of the overflow list
Next Pointer	Address to write the next entry
Size	Size of the overflow list

Table II: Hardware Overhead

the transaction state is updated to *Commit Complete*. At this point the transaction has completed and core 1 may begin a new transaction.

(g) Abort. Shows the state of the system if the transaction were to abort after (d). An abort log record is written to the log area, the read bits and the read overflow signature in the L1 are cleared, and the transaction status is updated to *Abort*. The transaction has aborted at this point and core 1 may continue executing subsequent non-transactional instructions.

(h) Abort Complete. In the abort completion phase, the L1 invalidates cache lines *B* and *E* belonging to the write set and sends invalidate messages to the LLC which then clears the sharer vector for cache lines *B* and *E*. Then the memory controller reads the overflow list and issues invalidate message for cache line *A* to the LLC. On receiving the invalidate message, the LLC invalidates cache line *A* and clears its sharer vector and dirty bit. Finally, the transaction status is updated to *Abort Complete* and at this point, core 1 may begin a new transaction.

Hardware Overhead. Table II shows the hardware overhead that DHTM adds on top of an RTM-like HTM design. DHTM adds to the L1 cache a fully-associative structure called the *log-buffer*, for keeping track of cache lines for which redo log entries need to be written to persistent memory. It also adds a *transaction state* register to identify the current state of the transaction. DHTM also adds two sets of registers to keep track of the log area and the overflow list. The registers in each set consist of a *start pointer* to identify the start address of the corresponding area, a *next pointer* to identify the address where the next entry can be written to and finally a *size* register to keep track of the size of each area so that an overflow can be detected.

Fallback Path. DHTM increases the limit for the transaction size from the L1 cache to the LLC. However, if a transaction aborts continually because of a overflow from the LLC (or due to any other reason) then it might not be able to make forward progress. Therefore, a fallback path must be provided. In principle integrating a software fallback path to DHTM is no different from the ones proposed for RTM [44]

Cores	8 In-order cores @ 2GHz
L1 I/D Cache	32KB 64B lines, 4-way
L1 Access Latency	3 cycles
L2 Cache	1MB×8 tiles, 64B lines, 16-way
L2 Access Latency	30 cycles
MSHRs	32
NVM Access Latency	360 (240) cycles write (read)

Table III: System Parameters

Workload	Description	Write Set
TPC-C	Online transaction processing	590
TATP	Mobile carrier database	167
Queue	Insert/delete entries in a queue	52
Hash	Insert/delete entries in a hash table	58
SDG	Insert/delete edges in a scalable graph	56
SPS	Random swaps between entries in an array	63
BTree	Insert/delete nodes in a b-tree	61
RBTree	Insert/delete nodes in a red-black tree	53

Table IV: Benchmarks used in our experiments along with their descriptions and write-set sizes (# cache lines).

because both employ a similar mechanism for atomic visibility. In particular, this software fallback path does not interact with hardware logging because, before initiating the fallback path DHTM would abort the transaction (taking it to abort complete state) which clears the log. The only difference is that our fallback will provide both atomic visibility and durability similar to Mnemosyne [14].

V. EXPERIMENTAL SETUP

We now describe our simulation infrastructure, system configuration, benchmarks and designs that we evaluate. We implemented DHTM on the gem5 [45] simulator with Ruby. We extend the Ruby memory model to implement DHTM functionality, with a log buffer size of 64 entries. We evaluate DHTM on an 8-core multicore (one thread per core) with each core containing a 32 KB private L1 and a multi-banked LLC. The local L1s are kept coherent using a MESI based directory protocol. DHTM is built on top of HTM that is based on an RTM-like implementation. Conflicts are detected by piggybacking on top of the coherence protocol and the HTM employs a first-writer wins conflict resolution policy similar to IBM POWER8 [29]. However, it is important to note that the choice of conflict resolution policy is not fundamental to DHTM design and it can be implemented with other policies (like requester wins) as well. Table III shows the main parameters of our system. The peak memory bandwidth in our setup is 5.3 GB/s.

Workloads and their Characteristics. We considered two classes of workloads for our study. TPC-C and TATP (the first two rows from Table. IV) are traditional online transaction processing (OLTP) workloads that require ACID guarantees. We use in-memory implementations of these

workloads [12]. It is worth noting that the OLTP workloads have write working-set sizes exceeding or comparable to the size of the L1 cache. Indeed, the write-set size of TPC-C (37 KB) exceeds the L1 cache size (32 KB) and can cause both capacity and conflict L1 misses which in turn can cause aborts when run on an HTM. Although TATP has a write-set size of around 10 KB, we find that there are significant conflict misses, which can lead to aborts.

The second class of workloads (the last six rows from Table. IV) are micro-benchmarks that perform atomic search, insert and delete operations on the corresponding data structure. The micro-benchmarks are similar to those in the benchmark suite used by NVHeaps [24]. We evaluate each of these micro-benchmarks with a data set size of 3 KB, similar to ATOM [20]. It is worth noting that the write-set sizes of the micro-benchmarks are significantly smaller than the L1 size, in contrast to the OLTP workloads.

Evaluated Designs:

- **SO:** This *software only* design uses locks for atomic visibility and software logging for atomic durability. For the OLTP workloads, we use the default software concurrency control mechanism which uses fine-grained locking. For the micro-benchmarks, we partition the data-structure into coarse-grained partitions with a lock associated with each partition, to allow for concurrency across the different partitions. We use a software logging mechanism similar to Mnemosyne [14], wherein log entries are flushed synchronously as soon as their values are finalized (thus benefiting from coalescing as well).
- **sdTM:** This design (software durability + hardware transactional memory) is based on PHyTM [32], which uses HTM similar to Intel’s RTM for atomic visibility. We disable the software concurrency control mechanism in the benchmarks and instead enclose each transaction within a hardware transaction. We use software logging similar to Mnemosyne for atomic durability.
- **ATOM:** This design uses locks (similar to the **SO** design) for atomic visibility. It uses the state-of-the-art hardware undo logging mechanism for atomic durability [20].
- **LogTM-ATOM:** This design uses LogTM [37] like HTM for atomic visibility and integrates it with ATOM [20] for atomic durability. It is worth noting that this represents a new design that has not been studied previously.
- **DHTM:** This is our proposed design which supports atomic visibility by using HTM similar to Intel’s RTM and atomic durability by using hardware based redo logging. It also allows the write set to overflow from the L1 to the LLC.

VI. RESULTS

In this section, we quantitatively compare the performance of the evaluated designs on the micro-benchmarks. We then present studies to better understand where the DHTM gains are coming from. One important parameter that could affect

Figure 5: Transaction throughput normalized to SO.

	queue	hash	sdg	sps	btree	rbtree	Ave.
sdTM	68	19	23	27	37	46	37
DHTM	46	5	13	16	18	26	21

Table V: Abort rates for sdTM and DHTM designs.

the efficacy of DHTM is the size of the log-buffer. Therefore, we quantify its impact. We also evaluate the efficacy of the designs on TPC-C and TATP workloads. Finally, we analyze the overheads of persistence by comparing our design with a non-persistent design.

A. Transaction Throughput

Figure 5 shows the transaction throughput of all the evaluated designs normalized to the software-only (SO) design, on the micro-benchmarks. As we can see, sdTM provides an average throughput improvement of 20% over SO. Recall that sdTM uses HTM for concurrency control which can potentially uncover more concurrency, especially in workloads where locking is coarse-grained. On the other hand, sdTM can suffer the negative effects of rollbacks in situation where the HTM aborts frequently. Table V shows the abort rates for the workloads for the sdTM design. In general, we can observe a correlation between the abort rates experienced by various workloads and the throughput improvement over SO. In particular, for the *rbtree* workload which experiences a significant 46% abort rate, sdTM provides only a minimal 5% improvement over SO.

We can also observe that ATOM provides a more robust average improvement of 35% over SO. Recall that ATOM uses the same concurrency control mechanism as SO (locks), but provides faster atomic durability by performing undo logging in hardware. Interestingly, we can see that ATOM has a comfortable 15% advantage over sdTM because of the aborts experienced by the latter.

Our DHTM design provides the best throughput improvement amongst all competing designs. On average, DHTM improves transaction throughput by 61% compared to SO. In comparison to SO, it can achieve faster durability by way of hardware logging and can also uncover more concurrency. In comparison with sdTM, DHTM improves transaction

Figure 6: Normalized transaction throughput sensitivity to log-buffer size for hash benchmark.

throughput by 41%. It not only benefits from faster durability, but also benefits from fewer aborts because DHTM supports write-set overflows from the L1. This is evidenced in Table V, where we can see that DHTM suffers from a relatively lower 21% abort ratio in comparison with the 37% abort ratio of sdTM. In comparison with ATOM, DHTM provides a 26% higher improvement in transaction throughput. DHTM not only benefits from better concurrency (because of HTM), but also because logging is faster in DHTM. Indeed, because ATOM uses undo logging, it suffers from the overhead of persisting data in-place in the commit critical path. In contrast, because DHTM uses redo logging, data can be persisted out of the critical path.

Finally, we look at LogTM-ATOM which implements eager version management for both atomic visibility and atomic durability. On average, DHTM provides 17% higher improvement in throughput compared to LogTM-ATOM. Since both designs leverage HTM for atomic visibility, the difference in performance between the two designs is primarily because of difference in atomic durability mechanisms. In other words, the low transaction commit latency in DHTM enabled by the redo log leads to performance improvement over LogTM-ATOM. From this we can also infer that more than half of the DHTM’s 26% improvement over ATOM is because of faster durability (and the rest owing to higher concurrency). In summary, DHTM provides a significant performance improvement, because of faster logging and/or integration with HTM, over the state-of-the-art design (ATOM [20]) and over a (novel) design combining LogTM [37] with ATOM.

B. Sensitivity to the size of the log-buffer

Figure 6 shows the impact of the size of log buffer on the performance of DHTM for the hash benchmark (other benchmarks show similar trends). We run the benchmark with buffer sizes ranging from 4 through 128 entries. As the number of entries are increased, the throughput increases, saturates at the size of 64 entries (default configuration in DHTM) and then marginally reduces upon further increase in the size. A small persist buffer size leads to creation of multiple redo log entries which consumes higher amount of memory bandwidth and adversely impacts other memory

	SO	ATOM	DHTM
TPC-C	1	1.67	1.88
TATP	1	1.27	1.53

Table VI: Transaction throughput for ATOM and DHTM normalized to SO for TPC-C and TATP benchmarks.

	1×	2×	10×
NP	2.9	3.0	3.3
DHTM	1.9	2.4	3

Table VII: Transaction throughput for NP and DHTM normalized to SO for hash benchmark with varying memory bandwidth.

requests. On the other hand, a larger buffer delays log writes which results in those log writes happening in the critical path of commit. Recall that a transaction cannot commit until all the redo log entries have been made persistent. In summary, we find that a 64-entry log-buffer provides the best coalescing effect.

C. TPC-C and TATP Throughput

Table VI shows the transaction throughput of ATOM and DHTM normalized to the throughput of SO for the TPC-C and TATP workloads. We have not shown sdTM results because it performs quite poorly. Because of the significant number of HTM aborts (owing to the large write working-set size of the OLTP workloads), these conventional HTM designs revert to the software concurrency mode often.

As we can see, DHTM continues to provide impressive speedups, not only over the SO baseline, but also over ATOM. Specifically, for the TPC-C workload, DHTM provides an 88% improvement over SO and 21% higher improvement compared to ATOM. For the TATP workload, DHTM provides a 53% improvement over SO and 26% higher improvement compared to ATOM.

D. The Cost of Atomic Durability

In this section, we wanted to see how close DHTM is compared to a non-persistent (volatile) HTM design that we call NP. For micro-benchmarks, we find that NP provides 2.2× higher transaction throughput compared to SO which is 59% better than DHTM.

Next, we wanted to understand better the reason for the performance gap. Are their inefficiencies in DHTM or is it fundamentally limited by the cost of durability? There are two primary sources of overheads in DHTM compared to NP. First, the overhead of log writes and data writes that are in the critical path. These include log writes that are pending when a transaction execution completes and is waiting to commit and data writes from the committed but yet to complete transaction pending when a core encounters the next transaction. To evaluate the performance impact of

these overheads we implemented a DHTM design where these writes happen instantaneously. This design is able to improve performance over DHTM by 16% for micro-benchmarks. Therefore log/data writes in the critical path of execution appear not to be the major source of overhead.

The second source of overhead corresponds to a fundamental difference in memory write bandwidth requirements for DHTM and NP. Recall that, in comparison with NP, DHTM needs to flush cache lines for atomic durability. To analyze the impact of this overhead, we performed experiments by varying the available memory bandwidth. Table VII shows the transaction throughput for NP and DHTM designs normalized to SO design for the hash micro-benchmark with varying memory bandwidth. With the baseline bandwidth (5.3 GB/s) the difference between NP and DHTM designs is 100% whereas with $10\times$ the baseline bandwidth the difference is only 30%. Thus in a system with higher memory bandwidth DHTM can achieve performance similar to that of a volatile only (NP) design.

VII. CONCLUSION

ACID transactions are a well-understood and widely adopted programming model. How fast can we achieve ACID in the presence of fast persistent memory? We have proposed DHTM, a HTM design in which durability is treated as a first class design constraint. It extends a commercial HTM like RTM with hardware support for atomic durability. It supports atomic visibility by employing an RTM like HTM and atomic durability by employing a hardware logging infrastructure which transparently and efficiently writes redo log entries to persistent memory. A redo-log based design allows us to commit a transaction as soon as all the log entries have been made persistent, without waiting for the data to persist.

One of our design goals was to support larger transactions, since ACID transactions tend to be considerably larger than those supported by current L1-limited RTM like HTM designs. But in supporting larger transactions we did not want to introduce significant hardware complexity. In particular, we did not want to introduce changes to the shared LLC – something that current HTM designs avoid. Our key insight here is to reuse the logging infrastructure that is necessary for durability for also supporting L1 overflows. Our experimental results showed that our proposal outperforms the state-of-the-art ACID design by an average of 21% to 25% on TATP, TPC-C and a set of micro-benchmarks.

ACKNOWLEDGMENT

We would like to thank our shepherd, Daniel Sanchez, and the anonymous reviewers for their helpful comments. This work is supported by the Intel University Research Office and by EPSRC grants EP/M001202/1 and EP/M027317/1 to the University of Edinburgh.

REFERENCES

- [1] H. Akinaga and H. Shima, “Resistive Random Access Memory (ReRAM) Based on Metal Oxides,” *Proceedings of the IEEE*, vol. 98, no. 12, 2010.
- [2] Intel Corporation and Micron, “Intel and Micron Produce Breakthrough Memory Technology,” 2015, http://newsroom.intel.com/community/intel_newsroom/blog/2015/07/28/intel-and-micron-produce-breakthrough-memory-technology.
- [3] T. Kawahara, R. Takemura, K. Miura, J. Hayakawa, S. Ikeda, Y. M. Lee, R. Sasaki, Y. Goto, K. Ito, T. Meguro, F. Matsukura, H. Takahashi, H. Matsuoka, and H. Ohno, “2Mb Spin-Transfer Torque RAM (SPRAM) with Bit-by-Bit Bidirectional Current Write and Parallelizing-Direction Current Read,” in *Proceedings of the International Solid-State Circuits Conference*, 2007.
- [4] S. Raoux, G. W. Burr, M. J. Breitwisch, C. T. Rettner, Y. Chen, R. M. Shelby, M. Salinga, D. Krebs, S. Chen, H. Lung, and C. H. Lam, “Phase-change random access memory: A scalable technology,” *IBM Journal of Research and Development*, vol. 52, no. 4-5, 2008.
- [5] V. J. Marathe, M. Seltzer, S. Byan, and T. Harris, “Persistent Memcached: Bringing Legacy Code to Byte-Addressable Persistent Memory,” in *USENIX Workshop on Hot Topics in Storage and File Systems*, 2017.
- [6] S. Harizopoulos, D. J. Abadi, S. Madden, and M. Stonebraker, “OLTP Through the Looking Glass, and What We Found There,” in *Proceedings of the International Conference on Management of Data*, 2008.
- [7] T. Karnagel, R. Dementiev, R. Rajwar, K. Lai, T. Legler, B. Schlegel, and W. Lehner, “Improving in-memory database index performance with Intel[®] Transactional Synchronization Extensions,” in *Proceedings of the International Symposium on High Performance Computer Architecture*, 2014.
- [8] V. Leis, A. Kemper, and T. Neumann, “Scaling HTM-Supported Database Transactions to Many Cores,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 2, 2016.
- [9] J. Huang, K. Schwan, and M. K. Qureshi, “NVRAM-aware Logging in Transaction Systems,” *Proc. VLDB Endow.*, 2014.
- [10] T. Wang and R. Johnson, “Scalable Logging Through Emerging Non-volatile Memory,” *Proc. VLDB Endow.*, 2014.
- [11] D. R. Chakrabarti, H.-J. Boehm, and K. Bhandari, “Atlas: Leveraging Locks for Non-volatile Memory Consistency,” in *Proceedings of the International Conference on Object Oriented Programming Systems Languages & Applications*, 2014.
- [12] A. Chatzistergiou, M. Cintra, and S. D. Viglas, “REWIND: Recovery Write-Ahead System for In-memory Non-volatile Data-Structures,” *Proc. of VLDB Endow.*, vol. 8, no. 5, 2015.
- [13] M. Liu, M. Zhang, K. Chen, X. Qian, Y. Wu, W. Zheng, and J. Ren, “DUDETM: Building Durable Transactions with Decoupling for Persistent Memory,” in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 2017.
- [14] H. Volos, A. J. Tack, and M. M. Swift, “Mnemosyne: Lightweight Persistent Memory,” in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 2011.
- [15] K. Doshi, E. Giles, and P. Varman, “Atomic Persistence for SCM with a Non-intrusive Backend Controller,” in *Proceed-*

ings of the International Conference on High Performance Computer Architecture, 2016.

- [16] A. Kolli, J. Rosen, S. Diestelhorst, A. Saidi, S. Pelley, S. Liu, P. M. Chen, and T. F. Wenisch, "Delegated Persist Ordering," in *Proceedings of the International Symposium on Microarchitecture*, 2016.
- [17] Y. Lu, J. Shu, L. Sun, and O. Mutlu, "Loose-ordering consistency for persistent memory," in *Proceedings of the International Conference on Computer Design*, 2014.
- [18] S. Nalli, S. Haria, M. D. Hill, M. M. Swift, H. Volos, and K. Keeton, "An Analysis of Persistent Memory Use with WHISPER," in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 2017.
- [19] J. Zhao, O. Mutlu, and Y. Xie, "FIRM: Fair and High-Performance Memory Control for Persistent Memory Systems," in *Proceedings of the International Symposium on Microarchitecture*, 2014.
- [20] A. Joshi, V. Nagarajan, S. Viglas, and M. Cintra, "ATOM: Atomic Durability in Non-volatile Memory through Hardware Logging," in *Proceedings of the International Symposium on High Performance Computer Architecture*, 2017.
- [21] E. L. M. Matheus A. Ogleari and J. Zhao, "Relaxing persistent memory constraints with hardware-driven undo+redo logging," University of California, Santa Cruz, Tech. Rep., 2016. [Online]. Available: <https://users.soe.ucsc.edu/~jzhao/files/HardwareLogging.pdf>
- [22] S. Shin, S. K. Tirukkovalluri, J. Tuck, and Y. Solihin, "Proteus: A Flexible and Fast Software Supported Hardware Logging Approach for NVM," in *Proceedings of the International Symposium on Microarchitecture*, 2017.
- [23] J. Zhao, S. Li, D. H. Yoon, Y. Xie, and N. P. Jouppi, "Kiln: Closing the Performance Gap Between Systems With and Without Persistence Support," in *Proceedings of the International Symposium on Microarchitecture*, 2013.
- [24] 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 International Conference on Architectural Support for Programming Languages and Operating Systems*, 2011.
- [25] A. Kolli, S. Pelley, A. Saidi, P. M. Chen, and T. F. Wenisch, "High-Performance Transactions for Persistent Memories," in *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*, 2016.
- [26] J. Chung, L. Yen, S. Diestelhorst, M. Pohlack, M. Hohmuth, D. Christie, and D. Grossman, "ASF: AMD64 Extension for Lock-Free Data Structures and Transactional Memory," in *Proceedings of the International Symposium on Microarchitecture*, 2010.
- [27] R. Cypher, A. Landin, H. Zeffer, S. Yip, M. Karlsson, M. Ekman, S. Chaudhry, and M. Tremblay, "Rock: A High-Performance Sparc CMT Processor," *IEEE Micro*, 2009.
- [28] Intel Corporation, *Intel® Architecture Instruction Set Extensions Programming Reference*, 2014.
- [29] H. Q. Le, G. Guthrie, D. Williams, M. M. Michael, B. Frey, W. J. Starke, C. May, R. Odaira, and T. Nakaike, "Transactional memory support in the IBM POWER8 processor," *IBM Journal of Research and Development*, 2015.
- [30] A. Wang, M. Gaudet, P. Wu, J. N. Amaral, M. Ohmacht, C. Barton, R. Silvera, and M. Michael, "Evaluation of blue gene/q hardware support for transactional memories," in *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, 2012.
- [31] V. Leis, A. Kemper, and T. Neumann, "Exploiting hardware transactional memory in main-memory databases," in *International Conference on Data Engineering*, 2014.
- [32] H. Avni and T. Brown, "PhyTM: Persistent Hybrid Transactional Memory," *Proceedings of VLDB Endowment*, 2016.
- [33] E. Giles, K. Doshi, and P. Varman, "Continuous Checkpointing of HTM Transactions in NVM," in *Proceedings of the International Symposium on Memory Management*, 2017.
- [34] Z. Wang, H. Yi, R. Liu, M. Dong, and H. Chen, "Persistent Transactional Memory," *IEEE Computer Architecture Letters*, 2015.
- [35] A. Joshi, V. Nagarajan, M. Cintra, and S. Viglas, "Efficient Persist Barriers for Multicores," in *Proceedings of the International Symposium on Microarchitecture*, 2015.
- [36] E. Giles, K. Doshi, and P. Varman, "Brief Announcement: Hardware Transactional Storage Class Memory," in *Proceedings of the Symposium on Parallelism in Algorithms and Architectures*, 2017.
- [37] K. E. Moore, J. Bobba, M. J. Moravan, M. D. Hill, and D. A. Wood, "LogTM: log-based transactional memory," in *Proceedings of the International Symposium on High Performance Computer Architecture*, 2006.
- [38] C. S. Ananian, K. Asanovic, B. C. Kuszmaul, C. E. Leiserson, and S. Lie, "Unbounded transactional memory," in *Proceedings of the International Symposium on High-Performance Computer Architecture*, 2005.
- [39] R. Rajwar, M. Herlihy, and K. Lai, "Virtualizing transactional memory," in *Proceedings of the International Symposium on Computer Architecture*, 2005.
- [40] M. Herlihy and J. E. B. Moss, "Transactional Memory: Architectural Support for Lock-Free Data Structures," in *Proceedings of the International Symposium on Computer Architecture*, 1993.
- [41] T. Harris, J. Larus, and R. Rajwar, *Transactional Memory*, 2nd Edition, 2nd ed. Morgan and Claypool Publishers, 2010.
- [42] 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 Trans. Database Syst.*, vol. 17, no. 1, 1992.
- [43] D. J. Sorin, M. D. Hill, and D. A. Wood, *A Primer on Memory Consistency and Cache Coherence*, 1st ed. Morgan & Claypool Publishers, 2011.
- [44] I. Calciu, T. Shpeisman, G. Pokam, and M. Herlihy, "Improved single global lock fallback for best-effort hardware transactional memory," in *Workshop on Transactional Computing*, 2014.
- [45] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, and D. A. Wood, "The Gem5 Simulator," *SIGARCH Comput. Archit. News*, vol. 39, no. 2, 2011.