

# HaTS: Hardware-assisted Transaction Scheduler

**Zhanhao Chen**

Palo Alto Networks Inc, USA  
zhachen@paloaltonetworks.com

**Ahmed Hassan**

Lehigh University, USA  
ahmed.hassan@lehigh.edu

**Masoomeh Javidi Kishi**

Lehigh University, USA  
maj717@lehigh.edu

**Jacob Nelson**

Lehigh University, USA  
jjn217@lehigh.edu

**Roberto Palmieri**

Lehigh University, USA  
palmieri@lehigh.edu

---

## Abstract

In this paper we present HaTS, a Hardware-assisted Transaction Scheduler. HaTS improves performance of concurrent applications by classifying the executions of their atomic blocks (or in-memory transactions) into scheduling queues, according to their so called conflict indicators. The goal is to group those transactions that are conflicting while letting non-conflicting transactions proceed in parallel. Two core innovations characterize HaTS. First, HaTS does not assume the availability of precise information associated with incoming transactions in order to proceed with the classification. It relaxes this assumption by exploiting the inherent conflict resolution provided by Hardware Transactional Memory (HTM). Second, HaTS dynamically adjusts the number of the scheduling queues in order to capture the actual application contention level. Performance results using the STAMP benchmark suite show up to 2x improvement over state-of-the-art HTM-based scheduling techniques.

**2012 ACM Subject Classification** Software and its engineering → Scheduling

**Keywords and phrases** Transactions, Scheduling, Hardware Transactional Memory

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2019.31

**Funding** This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-17-1-0367 and by the National Science Foundation under Grant No. CNS-1814974.

## 1 Introduction

Without reservation, in-memory transactions have experienced a significant growth in adoption during the last decade. Specifically, the advent of Hardware Transactional Memory (HTM) support in commodity processors [27, 7, 16, 30] has changed the way concurrent programs' execution is handled, especially in terms of performance advantages. Whether a multi-threaded application implements atomic blocks using locks or transactions, HTM can be exploited in both cases (e.g., using Hardware Lock Elision [26] in the former case or Restricted Transactional Memory [27] in the latter case) to accelerate its performance.

Hardware transactions are significantly faster than their software counterpart because they rely on the hardware cache-coherence protocol to detect conflicts, while Software



© Zhanhao Chen and Ahmed Hassan and Masoomeh Javidi Kishi and Jacob Nelson and Roberto Palmieri;

licensed under Creative Commons License CC-BY

23rd International Conference on Principles of Distributed Systems (OPODIS 2019).

Editors: Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller; Article No. 31; pp. 31:1–31:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 Transactional Memory (STM) [18] adds a significant overhead of instrumenting shared  
 46 memory operations to accomplish the same goal [9]. Relying on the cache-coherence protocol  
 47 also makes HTM appealing for mainstream adoption since it requires minimal changes in  
 48 hardware. However, this inherent characteristic of HTM represents an obstacle towards  
 49 defining contention management and scheduling policies for concurrent transactions, which  
 50 are crucial for both progress and fairness of HTM execution in the presence of conflicting  
 51 workloads. In fact, most TM implementations achieve high concurrency when the actual  
 52 contention level is low (i.e., few transactions conflict with each other). At higher contention  
 53 levels, without efficient scheduling, transactions abort each other more frequently, possibly  
 54 with a domino effect that can easily lead to performance similar to, if not worse than,  
 55 sequential execution [22, 21, 8].

56 A *Contention Manager* (CM) is the traditional, often encounter-time technique that  
 57 helps in managing concurrency. When a transaction conflicts with another one, the CM is  
 58 consulted to decide which of the two transactions can proceed. A CM collects statistics about  
 59 each transaction (e.g., start time, read/write-sets, number of retries, user-defined parameters)  
 60 and decides priorities among conflicting transactions according to the implemented policy.  
 61 Schedulers are similar to CMs except that they may proactively defer the execution of some  
 62 transactions to prevent conflicts rather than react to them. In both cases, performance is  
 63 improved by decreasing abort rate and fairness is achieved by selecting the proper transaction  
 64 to abort/defer [20, 31, 30].

65 The conflict resolution strategy of current off-the-shelf HTM implementations is provided  
 66 entirely in hardware, and can be roughly summarized as follows:

- 67 - the L1 cache of each CPU-core is used as a buffer for transactional read and write  
 68 operations<sup>1</sup>;
- 69 - the granularity of conflict detection is the cache line; and
- 70 - if a cache line is evicted or invalidated, the transaction is aborted (reproducing the idea  
 71 of read-set and write-set invalidation of STM [11]).

72 The above strategy thus implies a requester-wins contention management policy [6], which  
 73 informally means that a transaction  $T_1$  aborts another transaction  $T_2$  if  $T_2$  performed an  
 74 operation on a memory location that is physically stored in the same cache line currently  
 75 requested by  $T_1$ , excluding the case of two read operations, which never abort each other.  
 76 Due to this simple policy, classical CM policies cannot be trivially ported for scheduling HTM  
 77 transaction mainly because of two reasons. First, transactions are immediately aborted when  
 78 one of the cache lines in their footprint is invalidated, which makes it too late for CM to avoid  
 79 conflicts or manage them differently (e.g., by deciding which transaction is more convenient  
 80 to abort). Second, it is hard to embed additional metadata to monitor transactions behavior,  
 81 since all reads and writes executed within the boundaries of transactions are considered  
 82 transactional, even if the accessed locations store metadata rather than actual data.

83 In this paper we introduce HaTS (*Hardware-assisted Transaction Scheduler*), a transaction  
 84 scheduler that leverages the unique characteristics of HTM to accelerate scheduling in-memory  
 85 transactions. To overcome the aforementioned limitations of HTM, HaTS neither aims at  
 86 altering HTM's conflict resolution policy nor adds metadata instrumentation inside hardware  
 87 transactions, but instead relies on it to relax the need for the scheduler to define a non-  
 88 conflicting schedule. HaTS effectively arranges incoming transactions according to a set  
 89 of metadata collected either at compilation time (leveraging developer's annotations) or  
 90 at run time (after transactions commit/abort). HTM is then used to execute transactions

---

<sup>1</sup> In some HTM implementation [27], reads can be logged on the L2 cache to increase capacity.

91 concurrently while maintaining atomicity, isolation, and performance.

92 In a nutshell, HaTS works as follows. It uses a software classifier to queue incoming  
93 transactions with the goal of allowing only those transactions that do not conflict with  
94 each other to execute concurrently. The fundamental innovation of HaTS, which makes it  
95 practical, is that it works with incomplete or even erroneous information associated with  
96 incoming transactions. This is because even if the classifier erroneously allows conflicting  
97 transactions to run concurrently, HTM will inherently prevent them from committing (at  
98 least one of the conflicting transactions will abort). Therefore, misclassifications cannot  
99 impact the correctness of the transactional execution.

100 More in detail, HaTS offers a set of *scheduling queues* to group conflicting transactions.  
101 Membership of a queue is determined based on a single metadata object associated with  
102 each transaction, called *conflict indicator*. A conflict indicator might be provided by the  
103 programmer (e.g., the address of a contended memory location accessed transactionally) or  
104 computed by the system (e.g., transaction abort rate).

105 A queued transaction waits until the scheduler signals it when it becomes top-standing in  
106 its respective queue. When the transaction actually executes, the built-in HTM takes care of  
107 possible conflicts with transactions dispatched from other queues due to misclassification,  
108 which also includes the case where no conflict indicator is provided.

109 Another key feature of HaTS is that it adapts the number of scheduling queues based on  
110 a set of global statistics, such as the overall number of commits and aborts. This adaptation  
111 significantly improves performance in two common, and apparently dual cases. On the one  
112 hand, since conflict indicators are mainly *best effort* indicators, a single per transactions  
113 conflict indicator will not suffice when the application workload is highly conflicting. For  
114 this reason, HaTS reduces the number of queues when the overall conflict level increases,  
115 *enforcing transactions with different conflict indicators to be executed sequentially*. On the  
116 other hand, if the overall conflict level significantly decreases, HaTS increases the number  
117 of queues in order to (re-)allow transactions with different conflict indicators to execute  
118 concurrently. Additionally, when conflict level remains low, it enables dispatching multiple  
119 transactions from the same queue simultaneously, *allowing transactions with the same conflict*  
120 *indicator to execute in parallel*. Our framework aims at adaptively converging on an effective  
121 configuration of scheduling queues for the actual application workload. By leveraging HTM  
122 and its built-in atomicity guarantees, transitions between different configurations do not  
123 entail stalling transaction executions.

124 We implemented HaTS in C++ and integrated it into the software framework of SEER [13].  
125 We contrasted HaTS performance against Hardware Lock Elision (HLE) [26], Restricted  
126 Transactional Memory (RTM) [27], Software-assisted Conflict Management (SCM) [1], and  
127 SEER itself. We used the STAMP suite as our benchmark and we used a testbed of four-  
128 socket Intel platform with HTM implemented through TSX-NI. Results, including speedup  
129 over the sequential non-instrumented code and two abort rate metrics, show that HaTS  
130 outperforms competitors in both high contention and low contention scenarios. This stems  
131 from both leveraging conflict indicators and performing dynamic adjustment of scheduling  
132 queues, which leads to a notable performance improvement (e.g., 2x speedup in execution  
133 time for `Kmeans` and a 50% improvement for the `Vacation` benchmarks).

134 The rest of this paper is structured as follows. We review previous work in Section 2  
135 and the limitations of the current HTM implementations in Section 3. The design and  
136 implementation details of HaTS are presented in Section 4. We compare the performance of  
137 HaTS against state-of-art competitors in Section 5, and we conclude our paper in Section 6.

138 **2 Related Work**

139 TM schedulers can be classified based on whether they target STM systems [28, 15, 14, 5] or  
140 HTM systems [13, 25, 1, 4, 29, 31]. Although we deploy and test HaTS in an HTM-based  
141 environment, due to its performance advantages, HaTS can be deployed in STM systems as  
142 well.

143 Among the HTM-based schedulers, the closest one to HaTS is SEER [13]. SEER’s main  
144 idea is to infer the probability that two atomic blocks conflict, basing its observation on the  
145 commit/abort pattern witnessed in the past while the two were scheduled concurrently. Thus,  
146 HaTS and SEER are similar in their best-effort nature: both of them do not require precise  
147 information on the pair of conflicting transactions nor on the memory location(s) where they  
148 conflict, which is the key to coping with the limitations of current HTM implementations.

149 HaTS differs from SEER in two core points. First, SEER focuses only on managing  
150 HTM limitations, and thus it schedules transactions based on their commit/abort patterns.  
151 On the other hand, HaTS is a generic scheduler that uses HTM to preserve atomicity and  
152 consistency, and thus it defines a generic *conflict indicator* object that can embed both  
153 online and offline metadata. Second, SEER adopts a fine-grained (pairwise) locking approach  
154 to prevent transactions that are more likely to conflict from running concurrently. HaTS  
155 replaces this fine-grained locking scheme with a lightweight queueing scheme that controls the  
156 level of conflict by increasing/decreasing the number of scheduling queues. Our experimental  
157 results in Section 5 show that this lightweight scheme results in a lower overhead in different  
158 scenarios, as the case of Vacation where the percentage of committed transactions in HTM  
159 is comparable with SEER’s but overall application execution time is about 50% faster.

160 The work in [1] proposes a Software-assisted Conflict Management (SCM) extension to  
161 HTM-based lock elision (HLE) [26], where aborted transactions are serialized (using an  
162 auxiliary lock) and retried in HTM instead of falling back to a slow path with a single global  
163 lock. The main advantage of this approach is avoiding the lemming effect that causes new  
164 (non-conflicting) transactions to fall back to the slow path as well. As opposed to HaTS,  
165 SCM uses a conservative scheme where all aborted transactions are serialized without any  
166 further (even imprecise) conflict indicators, which limits concurrency. Moreover, SCM does  
167 not leverage the observed conflict pattern to proactively prevent conflicts in the future.

168 The idea of using scheduling queues to group conflicting transactions has been briefly  
169 discussed in [25], where authors introduced the concept of Octonauts. Octonauts uses  
170 statistics from transactions that committed in HTM to speculate over the characteristics of  
171 the associated transaction profile for future classification. HaTS is an evolution of Octonauts  
172 where a comprehensive software infrastructure, along with conflict indicators and a dynamic  
173 scheduling queueing techniques have been used to improve application performance.

174 A few other HTM-based schedulers were proposed prior to the release of Intel TSX  
175 extensions [27]. However, they either assume HTM implementations that are different from  
176 the currently deployed hardware [4, 29] or rely on a conservative single-lock-based serialization  
177 scheme similar to SCM [31].

178 STM-based schedulers rely on more precise information about transactions conflict.  
179 ProPS [28] uses a probabilistic approach similar to SEER but with a precise knowledge of  
180 the pair of conflicting transactions. Shrink [15] additionally uses the history of recently  
181 completed transactions’ read/write-sets to predict conflicts. CAR-STM [14] and SOA [5] use  
182 per-core queues such that an aborted transaction is placed in the queue of the transaction  
183 that causes it to abort. The necessity of precise information represents an obstacle towards  
184 adopting such techniques in HTM-based systems.

### 3 Background: Scheduling Best-effort Hardware Transactions

HTM provides a convenient concurrent programming abstraction because it guarantees safe and efficient accesses to shared memory. HTM executes atomic blocks of code optimistically, and during the execution all read and write operations to shared memory are recorded in a per-transaction log, which is maintained in a thread-local cache. Any two operations generated by two concurrent transactions accessing memory mapped to the same cache line trigger the abort of one of the transactions. HTM is known to have limited progress guarantees [3, 8, 24]. To guarantee progress, all HTM transactions are guaranteed to commit after a number of retries as HTM transactions by exploiting the traditional software lock-based fallback path [27]. To implement that, hardware transactions check if the fallback lock is acquired at the beginning of their execution. If so, the transactional execution is retried; otherwise the execution proceeds in hardware and mutual exclusion with the fallback path is implemented by leveraging the strong atomicity property of HTM, which aborts any hardware execution if the fallback lock is acquired at any moment. To reduce the well-known lemming effect [10] in HaTS, a transaction is not retried in HTM until the global lock is released.

For simplicity, in the rest of the paper we refer to HTM execution as the above process, which encompasses hardware trials followed by the fallback path, if needed. It is important to note that, since HaTS does not assume a specific methodology to provide HTM with progress, more optimized alternative solutions [22, 8, 12] can be integrated into our HTM execution to improve performance even further.

The off-the-shelf HTM implementation only provides limited information about reasons behind aborted transactions, which makes it very hard for programmer to introduce modifications that would increase the likelihood for that transaction to commit. As a result, in the presence of applications with contention, HTM might waste many CPU cycles until a transaction can succeed by either retrying multiple times, or by falling back to a single global lock where the protected HTM execution can be relaxed in favor of the mutual exclusion implemented by the lock itself.

Contention management for practical transaction processing systems is often formulated as an online problem where metadata, in the form of statistics (e.g., actual access pattern), can be collected by aborted and committed transactions in order to fine-tune scheduling activities. However, this methodology cannot be directly ported to HTM-protected concurrent executions since HTM cannot distinguish between a cache line that stores actual application data, or scheduling metadata. Because of that, conflicting accesses to shared metadata executed by two concurrent hardware transactions may cause at least one of them to abort, even if at the semantic level no conflict occurred.

The above issues motivated us to design a transaction scheduler where HTM is exploited as-is, instead of providing software innovations or hardware extensions aimed at influencing the HTM conflict resolution mechanism [2], which likely lead to degradation of HTM effectiveness.

### 4 Hardware Transaction Scheduler

In this section we overview the two core components of HaTS, namely the transaction conflict indicator (Section 4.1) and the dynamic scheduling queues (Section 4.2), along with a description of the complete transaction execution flow (Section 4.3) and the details of how threads execution passes through the scheduling queues (Section 4.4).

**Terminology.** HaTS has a set of  $N$  concurrent queues, called *scheduling queues*. Each thread that is about to start a transaction (i.e., an atomic block) is mapped to one of

231 those scheduling queues, and it starts executing only when HaTS dispatches it from that  
 232 queue. Each scheduling queue has one (or more) *dispatcher thread(s)*. As we detail later,  
 233 the mapping between each transaction and its corresponding scheduling queue is based on  
 234 the transaction's *conflict indicator* and the mapping is implemented using hashing. The  
 235 overall transaction commit/abort statistics are collected by HaTS and recorded into a shared  
 236 *knowledge base*. HaTS periodically consults the knowledge base to increase/decrease the  
 237 number of scheduling queues or the number of dispatcher threads, dynamically.

#### 238 4.1 Transaction Conflict Indicator

239 HaTS uses a so called *transaction conflict indicator*, provided as a parameter to TM-BEGIN  
 240 in our experimental study, to represent in a compact way characteristics that affect the  
 241 probability of aborting a hardware transaction due to conflict. Having this information is  
 242 indeed powerful because it allows HaTS to group transactions that access the same system's  
 243 hot spot in the same conflict queue, which saves aborts and increases throughput.

244 The transaction conflict indicator is an abstraction that can be deployed in many different  
 245 ways. A simple and effective example of conflict indicator is the address of the memory  
 246 location associated with the accessed system hot spot. As a concrete example in a real  
 247 application, let us consider a monetary application where transactions work on given bank  
 248 accounts. A transaction would use the address of the accessed bank account, which uniquely  
 249 identifies that object (or memory location) in the system, as its conflict indicator. This way,  
 250 although transactions might still (occasionally) conflict on other shared memory elements,  
 251 HaTS will be able to prevent conflicts between accesses to the same bank account, which is  
 252 one of the most contended set of objects in the system. Because of its effectiveness, in our  
 253 evaluation study we focused on system hot spots as the transaction conflict indicator.

254 Other examples of transaction conflict indicators include:

- 255 - *Abstract data types of accessed objects*: transactions accessing (possibly different) objects  
 256 of the same abstract data type will have the same conflict indicator. This represents a  
 257 more conservative approach than our adopted (per-object) conflict indicator, and it can  
 258 work better in workloads with higher contention levels.
- 259 - *HLE fallback lock(s)*: if hardware transactions are used for lock elision (HLE) [26], the  
 260 fallback paths of HTM transactions acquire the elided locks rather than a single global  
 261 lock, as in RTM. Using this fallback lock as a conflict indicator, transactions that elide  
 262 the same lock are grouped together.
- 263 - *Profile of aborted transactions*: HaTS's knowledge base can record the profile identification  
 264 of aborted transactions within a window of execution, and group incoming invocations of  
 265 those transactions using a single conflict indicator. This can significantly reduce abort  
 266 rate because those transactions are more prone to conflict in the future as well. To avoid  
 267 unnecessary serialization, the knowledge base can record only transactions aborted due to  
 268 conflict, and exclude transactions aborted for any other reason (i.e., capacity and explicit  
 269 aborts).
- 270 - *Core ID*: transactions running on the same physical core are assigned the same conflict  
 271 indicator. This can be beneficial because multiple hardware threads running on the  
 272 same physical core share the same L1 cache, which means that transactions concurrently  
 273 invoked by those threads are more prone to exceed cache capacity and abort.

274 The last two points reflect similar ideas already used in literature in different ways [13, 1,  
 275 28, 15, 14, 5]. Although we refer to Section 2 for a detailed comparison of these approaches  
 276 with HaTS, it is worth to mention here that HaTS's innovations relies on the fact that it  
 277 deploys those ideas in an abstract way, using conflict indicators, which allows for a better

278 scheduling management.

279 HaTS allows for the specification of a single conflict indicator per transaction. Although  
280 a single indicator might seem limited in terms of expressiveness, we adopt this approach  
281 because of the following reasons. First, there will always be a trade-off between the precision  
282 achieved by allowing multiple indicators and the additional cost needed to analyze them.  
283 Our decision is the consequence of an empirical study, that we excluded for space limitations,  
284 where we analyzed this trade-off. Second, the way HaTS adapts the number of scheduling  
285 queues (as detailed in the next section) is a dominant factor to manage contention that  
286 mitigates the effect of having imprecise, yet lightweight, conflict indicators. Finally, it is  
287 still possible to extend HaTS's infrastructure to support multiple indicators. For example,  
288 Bloom Filters can be used to compact multiple indicators and bit-wise operations (using  
289 either conjunction or disjunction operators) can be used to hash each bloom filter to the  
290 corresponding queue. As a future work, we plan to study the trade-off mentioned above;  
291 however, our current evaluation shows that even with a single indicator, HaTS outperforms  
292 existing approaches.

## 293 4.2 Dynamic Distribution of Scheduling Queues

294 Mapping transaction conflict indicators to scheduling queues is critical for achieving the goal  
295 of HaTS because it guarantees that transactions with the same conflict indicators are grouped  
296 in the same queue. However, using a static number of scheduling queues in such a mapping  
297 might lead to issues such as unbalancing, unfair scheduling, and poor adaptivity to some  
298 application workloads. For this reason, HaTS deploys a dynamic number of scheduling queues  
299 to cope with application workloads and effectively provide an elastic degree of parallelism.  
300 As we detailed in Section 4.3, this number is decided at run time according to the overall  
301 commit/abort statistics calculated in HaTS's knowledge base.

302 The following two examples clarify the need for having a dynamic set of scheduling queues.  
303 First, consider two incoming transactions with different conflict indicators. Since we have a  
304 finite number of scheduling queues, it is possible that those two transactions are mapped to  
305 the same queue. When the number  $N$  of queues is increased, the probability of mapping  
306 those transactions to the same queue decreases, and the level of parallelism increases. Second,  
307 consider a workload where transactions are highly conflicting so that the conflict indicator is  
308 not sufficient to capture all raised conflicts. In this case, decreasing the number of queues  
309 reduces parallelism and potentially reduces abort rate.

310 Adaptively changing the number of queues also covers more complex, yet not uncommon,  
311 cases. For example, it covers the case when transactions' data access pattern is hard to  
312 predict; therefore having a single conflict indicator per transaction may not be sufficient (e.g.,  
313 when each transaction accesses multiple system hot spots). Also, it covers the cases when no  
314 effective conflict indicator can be defined but the workload experiences high abort rates due  
315 to other reasons (e.g., aborts due to false sharing of the same cache lines). Finally, it allows  
316 schedulers to temporarily disable the usage of conflicting indicator as a medium for grouping  
317 transactions, in favor of a random policy, without hindering performance.

318 As will be clear in Section 4.3, dynamically changing the number of scheduling queues  
319 neither introduces blocking phases nor trades off correctness, thanks to the underlying HTM.

### 320 4.2.1 Multiple Dispatchers


321 An interesting example that is not covered by the aforementioned policy is when transactions  
322 with the same conflict indicator (and hence grouped in the same queue) are actually able to

323 execute concurrently. Although it may appear as an infrequent case, we recall that conflict  
 324 indicators are best-effort indicators that can be imprecise. Also, since conflicts are raised at  
 325 runtime according to the transaction execution pattern, it may happen that two conflicting  
 326 concurrent transactions succeed to commit even if they run concurrently (e.g., one of them  
 327 commits before the other one reaches the conflicting part).

328 HaTS addresses this case by allowing multiple dispatcher threads for a single scheduling  
 329 queue. Similar to the way we increase/decrease the number queues, we use abort rate  
 330 as an indicator to increase/decrease the number of dispatchers per queue. For additional  
 331 fine-tuning, we allow programmers to statically configure the number of scheduling queues.  
 332 In that sense, transactions with the same conflict indicator are executed in parallel only if  
 333 the overall contention level is low.

### 334 4.3 Transaction Execution Flow

335 Transactional operations are executed directly by application threads, without relying on  
 336 designated worker threads managed by HaTS. In fact, HaTS’s role is to dispatch thread  
 337 executions.



■ **Figure 1** HaTS software architecture and high-level threads execution flow.

338 HaTS restricts transactions mapped to the same scheduling queue to run sequentially,  
 339 while offloading the concurrency control handling to the underlying HTM engine. Figure 1  
 340 shows the execution flow of a transaction  $T$  executed by an application thread  $T_r$ .  $T_r$  first  
 341 hashes the *conflict indicator* of  $T$  (using `module N` hashing, where  $N$  is the current number  
 342 of scheduling queues) in order to find the matching scheduling queue  $Q$  (Step 1). After that,  
 343  $T_r$  effectively *suspends* its execution by enqueueing itself into  $Q$  (Step 2) and waiting until a  
 344 dispatcher thread *resumes* its execution (Step 3).

345 HaTS provides one (or more) dispatcher thread  $T_Q$  per conflict queue  $Q$ . Each dispatcher  
 346 thread resumes one waiting transaction execution at a time, in a closed-loop manner, meaning  
 347 the next queued transaction execution is resumed only after the previous one is successfully  
 348 committed. For the sake of fairness, each queue is implemented as a priority queue, where  
 349 the priority of each transaction is proportional to the number of aborts its enclosing atomic  
 350 block experienced in former executions (similar to the approach used in SEER [13] to infer  
 351 the conflict pattern between atomic blocks.).

352 When a thread execution is resumed, the corresponding transaction starts to execute  
 353 leveraging the HTM implementation (Step 4). During the hardware transaction execution,  
 354  $T_Q$  waits until  $T_r$  completes its transactional execution. After  $T_r$ ’s commit,  $T_Q$  takes control



(Step 5) and performs two operations: it updates the *knowledge base* with the needed information (namely the number and types of aborts before committing); and it dispatches the next thread execution waiting in  $Q$ .

A special background thread, called *updater*, is used to dynamically change the number of scheduling queues depending upon the effectiveness of the parallelism achieved by the current scheduling queues configuration. To do so, the updater thread queries the knowledge base and decides, according to the transaction abort rate measured so far, whether the total number  $N$  of scheduling queues should be increased, decreased, or unchanged (Step 6). In our implementation, we adopt a simple hill-climbing approach similar to the one used in [17, 12]. Briefly, if the observed abort rate is greater (less) than the last observed rate, we decrease (increase) the number of queues by one. The maximum number of queues is set to the number of physical cores and the minimum is set to one. We also allow programmers to override this policy by setting a fixed number of scheduling queues in order to eliminate the overhead of this dynamic behavior, especially when the most effective configuration is known. Interestingly, as we show later in our experimental results, this simple approach pays off in most of the tested benchmarks. As a future work, we plan to investigate more complex approaches, such as collecting more detailed information (e.g., the types of aborts) an use reinforcement learning to reach better estimates.

Changing the scheduling queues configuration does not cause stalls of the transaction execution and does not affect execution correctness. This is a great advantage of leveraging HTM conflict detection. Let us consider two incoming transactions  $T_1$  and  $T_2$  that in the current scheduling queues configuration would map to the same scheduling queue  $Q_1$ . Let us assume that the scheduling queues configuration changes after  $T_1$  is enqueued in  $Q_1$  and before  $T_2$  is enqueued. In this case, there is the possibility for  $T_2$  to be mapped to another queue  $Q_2$  in the new configuration, which ends up having  $T_1$  and  $T_2$  in two different queues (even though it might be the case they were both to be mapped to the same queue  $Q_2$  in the new configurations). Although one can consider this scenario as an example of misclassification of incoming transactions due to ongoing change of  $N$ , safety cannot be affected because of leveraging the HTM execution. Even if those transactions are conflicting, they will be correctly serialized by (likely) enforcing one of them to abort and fallback to HTM's global locking phase.

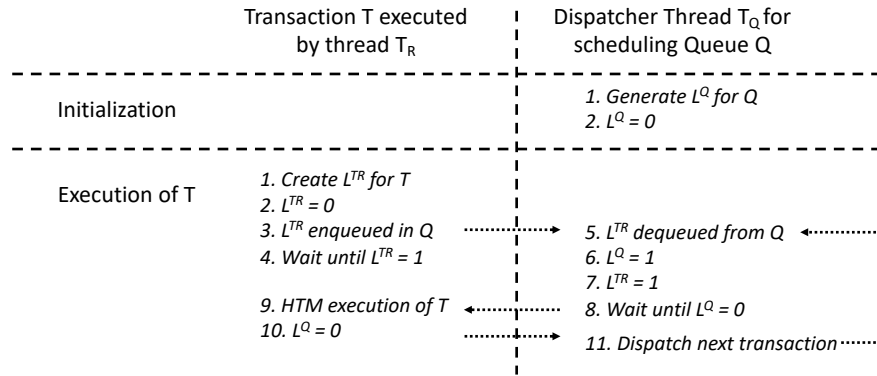
#### 4.4 Suspending/Resuming Executions with Scheduling Queues

As we mentioned in the previous section, a thread that wants to execute a transaction suspends and enqueues its execution until a dispatcher thread of the mapped queue resumes it. In order to synchronize this suspend/resume process, we use two synchronization flags<sup>2</sup>, one handled by the application thread and the other handled by the dispatcher thread.

Figure 2 pictures the synchronization scheme between an application thread  $T_r$  performing a transaction  $T$  and a dispatcher thread  $T_Q$  responsible for handling the scheduling queue  $Q$  that matches  $T$ 's conflict indicator. Numbers represent the sequence of operations completion.

When  $T_r$  wants to execute  $T$ , it creates a flag object  $L^{T_r}$  initialized to 0 and (atomically) enqueues it in  $Q$ , which effectively suspends its execution. After that,  $T_r$  spins until  $L^{T_r}$  is set to 1. When  $L^{T_r}$  becomes top standing in  $Q$ ,  $T_Q$  dequeues it. Then,  $T_Q$  sets a flag associated with  $Q$ , called  $L^Q$ , and also sets  $L^{T_r}$  to 1 (in that order). By setting  $L^{T_r}$ ,  $T_r$  will be signaled to proceed with its HTM execution. By setting  $L^Q$ ,  $T_Q$  is suspended until the

<sup>2</sup> Flags are implemented as `volatile` shared memory locations.



■ **Figure 2** Synchronization between application threads and dispatcher threads. For simplicity, the example accounts for a single dispatcher thread per scheduling queue.

399 completion of  $T$  by  $T_r$ . This suspension is implemented by spinning over the  $L^Q$  flag. When  
 400  $T$  is committed,  $T_r$  resets  $L^Q$  so that  $T_Q$  can dequeue the next thread execution waiting on  
 401  $Q$ . Note that  $T_Q$  is not notified if  $T$  is aborted and restarted for another HTM trial or if  $T$ 's  
 402 execution falls back to the software path.  $T_Q$  is resumed only after  $T$ 's successful completion.

403 In our implementation we use simple flags to synchronize two threads (i.e., application  
 404 thread and dispatcher thread) because we deploy one dispatcher thread for each scheduling  
 405 queue. As introduced earlier, HaTS allows for multiple dispatcher threads per queue in order  
 406 to cope with the case where the mapping between conflict indicators and scheduling queues  
 407 is unnecessarily unbalanced, meaning many transactions, possibly with different conflict  
 408 indicators, are grouped on the same scheduling queue. In the case where multiple dispatcher  
 409 threads are deployed per conflict queue, the same synchronization scheme illustrated before  
 410 applies, with the following differences. First,  $L^{T_r}$  flags should be atomically set (e.g., using a  
 411 Compare-And-Swap operation) to synchronize between dispatcher threads. Also, multiple  
 412  $L^Q$  flags, one per dispatcher thread, are needed to signal each dispatcher thread that it may  
 413 proceed to schedule the next transaction.

414 Scheduling queues are implemented in a lock-free manner [19] in order to speed up  
 415 the thread execution's suspension step. Also, in the above description we simplified the  
 416 presentation by saying that application threads and dispatcher threads spin over flags to  
 417 suspend/resume their execution. In the actual implementation, threads yield their execution  
 418 in order to let computing resources available so that the machine can be oversubscribed (if  
 419 needed) to maximize CPU utilization.

## 420 5 Evaluation

421 HaTS is implemented in C++ and integrated into the software framework of SEER [13]. An  
 422 advantage of using a unique software architecture for all competitors is that independent  
 423 optimization of low-level components does not bias performance towards some implementation.  
 424 In other words, the performance differences reported in our plots are due to the algorithmic  
 425 differences between competitors.

426 Our goal is to assess the practicality of HaTS in standard benchmarks for in-memory  
 427 transactional execution. Hence, we used STAMP [23], a benchmark suite of eight concurrent  
 428 applications that span several application domains with different execution patterns. Due to  
 429 space limitations, we present in detail the results of two applications, *Kmeans* and *Vacation*,

430 since they cover two different and important cases in which the characteristics of HaTS are  
431 highlighted. Then, we summarize the results with the other applications.


432 We compare HaTS against the scheduling techniques provided in the SEER framework,  
433 which are (in addition to SEER itself): Hardware Lock Elision (HLE) [26], Restricted  
434 Transactional Memory (RTM) [27], and Software-assisted Conflict Management (SCM) [1].  
435 Shortly, SEER is the state-of-the-art probability-based scheduling mechanism that uses  
436 fine-grained (pairwise) locks to prevent transactions that are likely to conflict from executing  
437 concurrently. HLE transforms atomic blocks protected by locks into hardware transactions  
438 and its software fallback path is protected by the original lock itself. In STAMP, a single lock  
439 for each atomic block is deployed. RTM supports a configurable number of retries before  
440 falling back to a single global lock shared among all hardware transactions. SCM implements  
441 a scheduling technique that serializes the aborted transactions to decrease the chance of  
442 further aborts. In all implementations, except HLE, transactions try at most five times in  
443 hardware before migrating to the software fallback path.

444 Experiments were conducted using a multi-processor platform equipped with 4 Intel Xeon  
445 Platinum 8160 processors (2.1GHz, 24 cores per CPU). The machine provides 96 physical  
446 cores and a total of 768 GB of memory divided into 4 NUMA zones. In our experiments we  
447 ran up to 64 application threads to leave resources for dispatcher threads (one per queue)  
448 and the updater thread. The maximum number of scheduling queues is statically set to 30  
449 prior execution, and we used the default operating system policy to map application threads  
450 to cores.

451 In Figure 3, we report for each application three performance metrics: (left column) speedup  
452 over sequential non-instrumented execution; (middle column) percentage of transactions  
453 committed through HTM; (right column) among those committed in HTM, percentage of  
454 transactions retried more than one time. Generally, the last two metrics are indicators of  
455 the scheduling effectiveness in reducing abort rate. The speedup metric is an indicator of  
456 whether such a reduction in abort rate is reflected in overall performance improvement or the  
457 scheduling overhead nullifies performance benefits. Performance at one application thread  
458 represent the slowdown of the sequential instrumented execution. All results are the average  
459 of 10 repeated tests.

460 The first application is *Kmeans*, a clustering algorithm that groups points into  $K$  clusters.  
461 Transactions are used by the application to synchronize the concurrent updates to the same  
462 cluster's center node. For this reason, we select the address of the center node updated by  
463 the transaction as its conflict indicator. We implemented that by passing the address of this  
464 center node as a parameter to STAMP's *TM-BEGIN* function. Our main observation is that  
465 identifying this conflict indicator allows HaTS to significantly reduce abort rate, reaching up  
466 to 1.5x reduction with respect to the closest competitor, and improve performance, reaching  
467 up to 2x better speedup over the closest competitor. SEER's probabilistic approach is the  
468 second best in terms of abort rate, which means that its approach is still able to capture  
469 conflicts, but not as effectively as using HaTS's conflict indicator. Moreover, SEER's speedup  
470 significantly decreases with higher number of threads, due to its locking overhead.

471 HLE does not perform well due to the lemming effect, which is visible as soon as few  
472 transactions fall back to locking. RTM is generally better than HLE due to its multiple  
473 retries in HTM before falling back to locking. SCM provides the worst performance. This is  
474 because the way SCM serializes conflicting transactions does not prevent new conflicts of  
475 incoming transactions, as opposed to the proactive scheduling approach, such as the one of  
476 HaTS and SEER. Also, the probability SCM executes non-conflicting transactions serially is  
477 higher than HaTS because it does not use any conflict indicators.

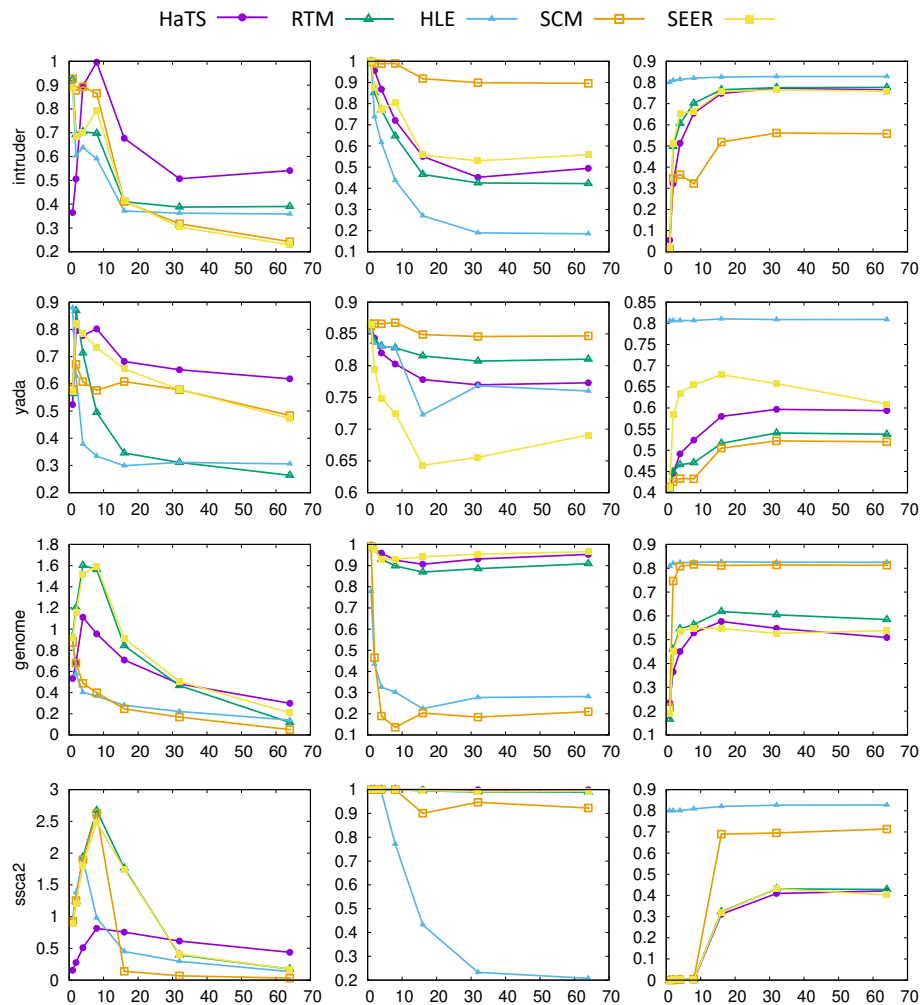


■ **Figure 3** Results with applications of STAMP benchmark. For each row, the left plot shows speedup over sequential (non-instrumented) execution; center plot shows the ratio of committed transactions in HTM; the right plot shows the ratio of transactions retried more than one time. X-axis shows number of application threads.

478 The difference between the high and low configuration of Kmeans is mainly in the  
 479 maximum achieved speedup. However, the patterns of the compared algorithms remain  
 480 the same, which shows the capability of HaTS to gain performance even in relatively low-  
 481 contention workloads.

482 The second application is *Vacation*, which simulates an online reservation system. Unlike  
 483 *Kmeans*, most transactions in *Vacation* apply operations on a set of randomly selected  
 484 objects, therefore with this pattern it is hard to identify a single conflict indicator per  
 485 transaction. For that reason, we adopt an approach where each transaction uses a unique  
 486 conflict indicator, with the exception of transactions that access a single customer object,  
 487 where we use the customer ID as a conflict indicator. Our rationale behind this decision is  
 488 that even if transactions are conflicting, HaTS’s dynamic adjustment of scheduling queues  
 489 reduces the level of parallelism and saves aborts. Indeed HaTS achieves an abort rate similar  
 490 to SEER, and moreover it scales better than SEER.

491 Summarizing our results of the other STAMP benchmarks (Figure 4) *Intruder* and *Yada*  
 492 give the same conclusions: the lightweight queuing approach in HaTS allows it to perform  
 493 better than SEER, especially for high number of threads, due to the overhead of SEER’s  
 494 locking mechanism. *SSCA* and *Genome* are low-contention benchmarks, and their abort rates  
 495 are very low even without scheduling or contention management. Hence, none of the compared



■ **Figure 4** Performance results on the remaining applications of STAMP benchmark. X-axis shows number of application threads.

496 algorithms had a significant improvement over the others. However, HaTS maintains its  
 497 performance better than others when the number of threads (and thus contention level)  
 498 increases. We excluded *Bayes* and *Labyrinth* because it is known they provide unstable,  
 499 and thus unreliable, results [13].

## 500 6 Conclusion

501 In this paper we presented HaTS, a Hardware-assisted Transaction Scheduler. HaTS groups  
 502 incoming transactions into scheduling queues depending upon the specified conflict indicators.  
 503 HaTS exploits the HTM conflict resolution to cope with the possibility of having erroneous  
 504 conflict indicators or when conflict indicators are complex to identify. Results using the  
 505 STAMP benchmark show that HaTS provides improvements in both high contention and  
 506 low contention workloads.

## 507 — References

- 508 1 Yehuda Afek, Amir Levy, and Adam Morrison. Software-improved hardware lock elision. In  
509 *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*, PODC  
510 '14, pages 212–221, New York, NY, USA, 2014. ACM. URL: <http://doi.acm.org/10.1145/2611462.2611482>, doi:10.1145/2611462.2611482.
- 512 2 Dan Alistarh, Syed Kamran Haider, Raphael Kübler, and Giorgi Nadiradze. The transactional  
513 conflict problem. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms  
514 and Architectures*, SPAA '18, pages 383–392, New York, NY, USA, 2018. ACM. URL:  
515 <http://doi.acm.org/10.1145/3210377.3210406>, doi:10.1145/3210377.3210406.
- 516 3 Dan Alistarh, Justin Kopinsky, Petr Kuznetsov, Srivatsan Ravi, and Nir Shavit. Inherent  
517 limitations of hybrid transactional memory. *Distributed Computing*, 31(3):167–185, Jun 2018.  
518 URL: <https://doi.org/10.1007/s00446-017-0305-3>, doi:10.1007/s00446-017-0305-3.
- 519 4 Mohammad Ansari, Behram Khan, Mikel Luján, Christos Kotselidis, Chris C. Kirkham, and  
520 Ian Watson. Improving performance by reducing aborts in hardware transactional memory.  
521 In *High Performance Embedded Architectures and Compilers, 5th International Conference,  
522 HiPEAC 2010, Pisa, Italy, January 25-27, 2010. Proceedings*, pages 35–49, 2010.
- 523 5 Mohammad Ansari, Mikel Luján, Christos Kotselidis, Kim Jarvis, Chris C. Kirkham, and  
524 Ian Watson. Steal-on-abort: Improving transactional memory performance through dynamic  
525 transaction reordering. In *High Performance Embedded Architectures and Compilers, Fourth  
526 International Conference, HiPEAC 2009, Paphos, Cyprus, January 25-28, 2009. Proceedings*,  
527 pages 4–18, 2009.
- 528 6 Jayaram Bobba, Kevin E. Moore, Haris Volos, Luke Yen, Mark D. Hill, Michael M. Swift,  
529 and David A. Wood. Performance pathologies in hardware transactional memory. In *34th  
530 International Symposium on Computer Architecture (ISCA 2007), June 9-13, 2007, San Diego,  
531 California, USA*, pages 81–91, 2007.
- 532 7 Harold W. Cain, Maged M. Michael, Brad Frey, Cathy May, Derek Williams, and Hung Le.  
533 Robust architectural support for transactional memory in the power architecture. In *ISCA*,  
534 pages 225–236, 2013.
- 535 8 Irina Calciu, Justin Gottschlich, Tatiana Shpeisman, Gilles Pokam, and Maurice Herlihy.  
536 Invyswell: A hybrid transactional memory for haswell's restricted transactional memory. In  
537 *Proceedings of the 23rd International Conference on Parallel Architectures and Compilation  
538 Techniques*, PACT '14, 2014.
- 539 9 Calin Cascaval, Colin Blundell, Maged Michael, Harold W. Cain, Peng Wu, Stefanie Chiras, and  
540 Siddhartha Chatterjee. Software transactional memory: Why is it only a research toy? *Queue*,  
541 6(5):40:46–40:58, September 2008. URL: <http://doi.acm.org/10.1145/1454456.1454466>,  
542 doi:10.1145/1454456.1454466.
- 543 10 Dave Dice, Yossi Lev, Mark Moir, and Daniel Nussbaum. Early experience with a commercial  
544 hardware transactional memory implementation. In *Proceedings of the 14th International  
545 Conference on Architectural Support for Programming Languages and Operating Systems*,  
546 ASPLOS XIV, pages 157–168, New York, NY, USA, 2009. ACM. URL: [http://doi.acm.org/  
547 10.1145/1508244.1508263](http://doi.acm.org/10.1145/1508244.1508263), doi:10.1145/1508244.1508263.
- 548 11 Dave Dice, Ori Shalev, and Nir Shavit. Transactional Locking II. In Shlomi Dolev, editor,  
549 *Distributed Computing*, volume 4167 of *Lecture Notes in Computer Science*, pages 194–208.  
550 Springer Berlin Heidelberg, 2006. URL: [http://dx.doi.org/10.1007/11864219\\_14](http://dx.doi.org/10.1007/11864219_14), doi:  
551 10.1007/11864219\_14.
- 552 12 Nuno Diegues and Paolo Romano. Self-tuning intel transactional synchronization extensions.  
553 In *11th International Conference on Autonomic Computing*, ICAC '14. USENIX Association,  
554 2014.
- 555 13 Nuno Diegues, Paolo Romano, and Stoyan Garbatov. Seer: Probabilistic scheduling for  
556 hardware transactional memory. In *Proceedings of the 27th ACM Symposium on Parallelism  
557 in Algorithms and Architectures*, SPAA '15, pages 224–233, New York, NY, USA, 2015. ACM.  
558 URL: <http://doi.acm.org/10.1145/2755573.2755578>, doi:10.1145/2755573.2755578.

- 559 **14** Shlomi Dolev, Danny Hendler, and Adi Suissa. CAR-STM: scheduling-based collision avoidance  
560 and resolution for software transactional memory. In *Proceedings of the Twenty-Seventh Annual*  
561 *ACM Symposium on Principles of Distributed Computing, PODC 2008, Toronto, Canada,*  
562 *August 18-21, 2008*, pages 125–134, 2008.
- 563 **15** Aleksandar Dragojevic, Rachid Guerraoui, Anmol V. Singh, and Vasu Singh. Preventing  
564 versus curing: avoiding conflicts in transactional memories. In *Proceedings of the 28th Annual*  
565 *ACM Symposium on Principles of Distributed Computing, PODC 2009, Calgary, Alberta,*  
566 *Canada, August 10-12, 2009*, pages 7–16, 2009.
- 567 **16** Ricardo Filipe, Shady Issa, Paolo Romano, and João Barreto. Stretching the capacity of  
568 hardware transactional memory in ibm power architectures. In *Proceedings of the 24th*  
569 *Symposium on Principles and Practice of Parallel Programming, PPOPP '19*, pages 107–119,  
570 New York, NY, USA, 2019. ACM. URL: <http://doi.acm.org/10.1145/3293883.3295714>,  
571 doi:10.1145/3293883.3295714.
- 572 **17** Ahmed Hassan, Roberto Palmieri, and Binoy Ravindran. Transactional interference-less  
573 balanced tree. In *Distributed Computing - 29th International Symposium, DISC 2015, Tokyo,*  
574 *Japan, October 7-9, 2015, Proceedings*, pages 325–340, 2015.
- 575 **18** Maurice Herlihy and J. Eliot B. Moss. Transactional memory: Architectural support for  
576 lock-free data structures. In *Proceedings of the 20th Annual International Symposium on*  
577 *Computer Architecture, ISCA '93*, pages 289–300, New York, NY, USA, 1993. ACM. URL:  
578 <http://doi.acm.org/10.1145/165123.165164>, doi:10.1145/165123.165164.
- 579 **19** Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming*. Morgan Kaufmann  
580 Publishers Inc., San Francisco, CA, USA, 2008.
- 581 **20** Junwhan Kim, Roberto Palmieri, and Binoy Ravindran. Enhancing concurrency in distributed  
582 transactional memory through commutativity. In Felix Wolf, Bernd Mohr, and Dieter an Mey,  
583 editors, *Euro-Par 2013 Parallel Processing*, pages 150–161, Berlin, Heidelberg, 2013. Springer  
584 Berlin Heidelberg.
- 585 **21** Alexander Matveev and Nir Shavit. Reduced hardware transactions: A new approach to  
586 hybrid transactional memory. In *Proceedings of the Twenty-fifth Annual ACM Symposium on*  
587 *Parallelism in Algorithms and Architectures, SPAA '13*, pages 11–22, New York, NY, USA,  
588 2013. ACM. URL: <http://doi.acm.org/10.1145/2486159.2486188>, doi:10.1145/2486159.  
589 2486188.
- 590 **22** Alexander Matveev and Nir Shavit. Reduced hardware norec: A safe and scalable hybrid  
591 transactional memory. In *Proceedings of the Twentieth International Conference on Architec-*  
592 *tural Support for Programming Languages and Operating Systems, ASPLOS '15*, pages 59–71,  
593 New York, NY, USA, 2015. ACM. URL: <http://doi.acm.org/10.1145/2694344.2694393>,  
594 doi:10.1145/2694344.2694393.
- 595 **23** Chi Cao Minh, Jaewoong Chung, C. Kozyrakis, and K. Olukotun. STAMP: Stanford trans-  
596 actional applications for multi-processing. In *Workload Characterization, 2008. IISWC*  
597 *2008. IEEE International Symposium on*, pages 35–46, Sept 2008. doi:10.1109/IISWC.2008.  
598 4636089.
- 599 **24** M. Mohamedin, R. Palmieri, A. Hassan, and B. Ravindran. Managing Resource Limitation of  
600 Best-Effort HTM. *IEEE Transactions on Parallel and Distributed Systems*, 28(8):2299–2313,  
601 Aug 2017. doi:10.1109/TPDS.2017.2668415.
- 602 **25** Mohamed Mohamedin, Roberto Palmieri, and Binoy Ravindran. Brief announcement: On  
603 scheduling best-effort htm transactions. In *Proceedings of the 27th ACM Symposium on*  
604 *Parallelism in Algorithms and Architectures, SPAA '15*, pages 74–76, New York, NY, USA,  
605 2015. ACM. URL: <http://doi.acm.org/10.1145/2755573.2755612>, doi:10.1145/2755573.  
606 2755612.
- 607 **26** Ravi Rajwar and James R. Goodman. Speculative lock elision: enabling highly concurrent  
608 multithreaded execution. In *Proceedings of the 34th Annual International Symposium on*  
609 *Microarchitecture, Austin, Texas, USA, December 1-5, 2001*, pages 294–305, 2001.

## 31:16 HaTS: Hardware-assisted Transaction Scheduler

- 610 **27** James Reinders. Transactional synchronization in haswell. *Intel Software Network*. URL: <http://software.intel.com/en-us/blogs/2012/02/07/transactional-synchronization-in-haswell/>, 2012.
- 611
- 612
- 613 **28** Hugo Rito and João P. Cachopo. Props: A progressively pessimistic scheduler for software transactional memory. In *Euro-Par 2014 Parallel Processing - 20th International Conference, Porto, Portugal, August 25-29, 2014. Proceedings*, pages 150–161, 2014.
- 614
- 615
- 616 **29** Christopher J. Rossbach, Owen S. Hofmann, Donald E. Porter, Hany E. Ramadan, Bhandari Aditya, and Emmett Witchel. Txlinux: using and managing hardware transactional memory in an operating system. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles 2007, SOSP 2007, Stevenson, Washington, USA, October 14-17, 2007*, pages 87–102, 2007.
- 617
- 618
- 619
- 620
- 621 **30** Lingxiang Xiang and Michael L. Scott. Conflict reduction in hardware transactions using advisory locks. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '15*, pages 234–243, New York, NY, USA, 2015. ACM. URL: <http://doi.acm.org/10.1145/2755573.2755577>, doi:10.1145/2755573.2755577.
- 622
- 623
- 624
- 625 **31** Richard M. Yoo and Hsien-Hsin S. Lee. Adaptive transaction scheduling for transactional memory systems. In *SPAA 2008: Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures, Munich, Germany, June 14-16, 2008*, pages 169–178, 2008.
- 626
- 627
- 628