

Insights into the Fallback Path of Best-Effort Hardware Transactional Memory Systems

Ricardo Quislan, Eladio Gutierrez, Emilio L. Zapata, and Oscar Plata

Department of Computer Architecture, University of Málaga, Spain, 29071
email: {quislan, eladio, zapata, oplata}@uma.es

Abstract. Current industry proposals for Hardware Transactional Memory (HTM) focus on best-effort solutions (BE-HTM) where hardware limits are imposed on transactions. These designs may show a significant performance degradation due to high contention scenarios and different hardware and operating system limitations that abort transactions, e.g. cache overflows, hardware and software exceptions, etc. To deal with these events and to ensure forward progress, BE-HTM systems usually provide a software fallback path to execute a lock-based version of the code.

In this paper, we propose a hardware implementation of an irrevocability mechanism as an alternative to the software fallback path to gain insight into the hardware improvements that could enhance the execution of such a fallback. Our mechanism anticipates the abort that causes the transaction serialization, and stalls other transactions in the system so that transactional work loss is minimized. In addition, we evaluate the main software fallback path approaches and propose the use of ticket locks that hold precise information of the number of transactions waiting to enter the fallback. Thus, the separation of transactional and fallback execution can be achieved in a precise manner.

The evaluation is carried out using the Simics/GEMS simulator and the complete range of STAMP transactional suite benchmarks. We obtain significant performance benefits of around twice the speedup and an abort reduction of 50% over the software fallback path for a number of benchmarks.

1 Introduction

Transactional Memory (TM) [8] was first presented in 1993 [9] as a non-blocking synchronization mechanism for shared memory chip multiprocessors (CMPs). TM provides the programmer with the *transaction* construct that executes the code within it atomically and in isolation. Such transactional properties are ensured by the TM system via the cache coherence protocol and dedicated hardware (hardware TM – HTM).

It is not until recently that some processor manufacturers have included HTM support in their commercial off-the-shelf CMPs [4, 10, 19, 21]. Current industry proposals focus on best-effort solutions (BE-HTM) where hardware limits are imposed on transactions. For instance, transactions cannot survive to capacity overflows, exceptions, interrupts, page faults, migrations,... To deal with these limitations, BE-HTM systems usually provide a software fallback path to execute a non-transactional version of the code, often comprising a global lock.

In this paper we propose an implementation of a hardware irrevocability mechanism as an alternative to the software fallback path to gain insight into the hardware improvements that could enhance the execution of such a fallback. Irrevocability [3, 20] is a transactional execution mode that ensures transaction forward progress since an irrevocable transaction cannot be aborted. Our mechanism anticipates the abort that causes the transaction serialization, and stalls other transactions in the system so that transactional work loss is minimized. In addition, we evaluate the main software fallback path approaches and propose the use of a ticket lock that hold precise information of the number of transactions waiting to enter the fallback. Thus, the separation of transactional and fallback execution can be achieved in a precise manner with the corresponding performance benefits. The result is an enhanced Lemming effect avoidance [6].

The evaluation is carried out using the Simics/GEMS simulator and the complete range of STAMP transactional suite benchmarks. We obtain significant performance benefits of around twice the speedup and an abort reduction of 50% over the software fallback path for a number of benchmarks.

2 Baseline Architecture

Figure 1 shows the baseline architecture used in this paper. The system relies on the L1 caches to store new transactional values of memory blocks, while old values are kept into the L2 cache. A pair of read and write transactional bits per L1 cache block marks whether the block was read or written within a transaction. Such bits can be flash-cleared on transaction commit and abort. In case of abort, the blocks whose transactional write bit is set are also invalidated. The cache coherence protocol maintains strong isolation [13] and implements an eager conflict detection policy. The conflict resolution policy is requester-wins, where the requesting transaction wins the conflict and the requested one is aborted. The baseline cache coherence protocol is modified to support the execution of transactions:

- *Backup on first transactional store*: If an L1 cache block is in M state and its write transactional bit is not set, the L1 cache has to send the data to the L2 cache before a transactional store is performed. This way the L2 cache holds the last old value for the block.
- *Abort on evictions*: The replacement of a transactional block in an L1 cache implies losing track of transactional loads and stores, which jeopardizes transaction isolation; so transactions must be aborted on these type of evictions. Beside, L2 cache block replacements may abort a transaction because of the inclusion property.
- *L2 cache serves data of aborted transactions*: The L2 cache must send the data of aborted transactions. There are two situations: (i) The requester is already the owner of the block. In such a case, the L2 cache simply responds with the data; (ii) The requester is not the owner of the block. In this case, the directory forwards the request to the owner, which receives a forward message for a block that is no longer present in its L1 cache. Then, the L1 cache informs the L2 cache and the L2 cache sends the data.

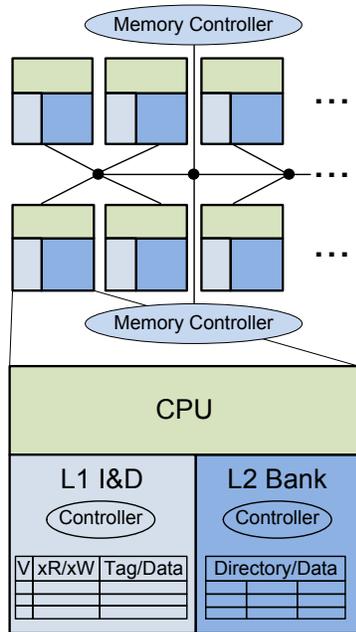


Fig. 1: Baseline architecture of the BE-HTM system.

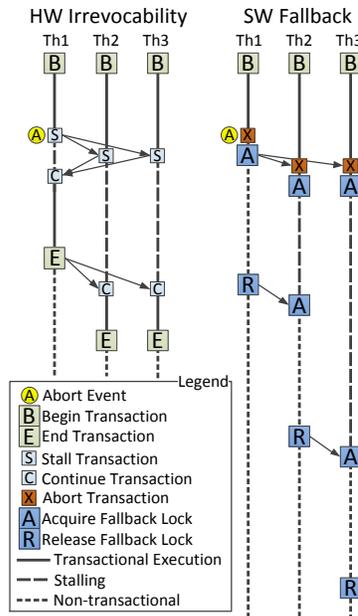


Fig. 2: Execution scenario of hardware irrevocability vs. software fallback.

3 Hardware Irrevocability Fallback Mechanism

A common way to deal with hardware capacity overflows and to ensure forward progress in commercial BE-HTM systems is a software fallback path. The code that Intel suggests as fallback path in its optimization manual [1] comprises a global lock to execute a failed transaction as a non-transactional critical section. Once a transaction aborts a given number of times, the fallback path is taken. In addition, when a transaction is successfully started, the fallback global lock is checked. If the lock was acquired, the transaction aborts. If not, the transaction goes ahead with the lock in its read set so that another transaction acquiring the lock can abort it. The clash of transactions and fallback path sections is thereby avoided.

A hardware irrevocability mechanism provides several benefits over a software fallback code of that kind:

- The programmer is not burdened with the task of writing and tuning a fallback code, which reduces the programming effort of transactional applications, one of the main goals of transactional memory.
- There is no need for a lock so it is neither cached nor added to the read set of the transaction, thus freeing limited hardware resources.

- Performance benefits: Figure 2 shows an execution scenario where a hardware irrevocability mechanism performs better than a software fallback code. The fallback path version aborts transactional execution and retries the transaction as a locked critical section. The other transactions running in the system abort as well, since they read the lock at the beginning¹. Execution is rebooted and serialized. However, the hardware irrevocability mechanism does not discard the transactional work done so far. The other transactions are stalled when a transaction gets irrevocable. Furthermore, the irrevocable one does not have to abort if it gets irrevocable just before the event that causes irrevocability, e.g. before an L1 cache replacement.

The scenario in Figure 2 is optimistic. It considers no contention between the irrevocable transaction and the stalled ones, which would cause the abort of the stalled, conflicting transactions. Additionally, the fallback code causes a chain reaction, also called as *Lemming effect* [6], by which all transactions take the fallback path even if they do not have reached the retry limit yet (Section 5 evaluates the Lemming effect problem). Nonetheless, the figure depicts the potential of hardware irrevocability and the weaknesses of a software fallback path.

3.1 Implementation

We propose a token-based implementation of the irrevocability mechanism where only the core that owns the token can run irrevocably. Each core has a flag that indicates whether there is an irrevocable transaction running in the system or not (the I bit). Another flag in the core signals whether the irrevocable transaction belongs to this core or to another core, i.e. whether the core owns the token or not (the T bit). Along with the pair of bits (I,T), each core has a counter (C) that holds the number of transaction retries. The core asks for irrevocability when C is 0.

When a transaction reaches the limit of retries, the L1 cache controller of the core checks its (I,T) bits and acts depending on their value:

- $(I,T) = (0,0)$: There are no irrevocable transactions running in the system and the token is not owned. In this case, the controller broadcasts a token request message that will be responded by the core that owns the token. Should the owner just start irrevocability, the token is not sent and the requester keeps stalling until the owner ends its transaction. If the token is received, the T bit is set to 1 and the controller broadcasts an irrevocability request message for the other cores to set the I bit to 1. The requester can safely continue its transaction in irrevocable mode, $(I,T) = (1,1)$, after acknowledgement of the other L1 cache controllers.
- $(I,T) = (0,1)$: The core owns the token, so it can request irrevocability directly.
- $(I,T) = (1,0)$: Someone else is running an irrevocable transaction. Consequently, the transaction stalls. This value for the (I,T) pair can be found on transaction beginning and after receiving an irrevocability message.

¹ The non-transactional write to the lock causes these aborts by means of strong isolation [13]. Correctness is ensured as locks and transactions are not allowed simultaneously.

Table 1: L1 cache coherence protocol modifications for irrevocability (highlighted in gray).

State	Events					
	L1 Replace $\neg(xR \vee xW)$ $\vee (1,1)$	L1 Replace $(xR \vee xW) \wedge (C > 1)$	L1 Replace $(xR \vee xW) \wedge (C \leq 1)$	L2 Replace $\neg(xR \vee xW)$ $\vee (1,1)$	L2 Replace $(xR \vee xW)$ $(1,0) \vee (C > 1)$	L2 Replace $(xR \vee xW)$ $(0,-) \wedge (C \leq 1)$
I	–	–	–	ACK	–	–
S	– / I	Abort, C-1 / I	Irre, Z	ACK / I	Abort, C-1 / I	Irre, Zz
E	PUT(no data) / I	Abort, C-1 / I	Irre, Z	ACK / I	Abort, C-1 / I	Irre, Zz
M	PUT+Data / I	Abort, C-1 / I	Irre, Z	ACK+Data / I	Abort, C-1 / I	Irre, Zz

Irre: ask for irrevocability

Z and Zz: recycle mandatory and request queue, respectively

(#, #): pair of bits (I, T)

We have modified the L1 cache controller to implement the anticipation to a block replacement. Table 1 shows the modifications made to the protocol highlighted in gray. L1 cache replacements are left untouched whenever either the block to be replaced is not transactional, $\neg(xR \vee xW)$, or the core is in irrevocable mode and owns the token, $(I, T) = (1, 1)$. However, if the block is transactional, $xR \vee xW$, the counter (C) is checked. If $C > 1$ (1 instead of 0 to anticipate the last abort) the transaction aborts and C is decremented. Conversely, if $C \leq 1$, the core asks for irrevocability and the mandatory queue is recycled² so that the event is triggered later on. Should the core manage to get irrevocable, the L1 replacement is performed safely. If irrevocability is not granted, the core stalls by continuously recycling the message that causes the eviction.

In case of L1 transactional block replacements due to L2 cache evictions (L2 Replace events in Table 1) we have different scenarios. If the core is running an irrevocable transaction, $(I, T) = (1, 1)$, the event is treated as a normal L2 cache replacement. However, if the irrevocable transaction is of another core, $(I, T) = (1, 0)$, the transaction in this core must be aborted in favour of the irrevocable one. Thus, the only situation in which a transaction asks for irrevocability on an L2 Replace event is when $C \leq 1$ and there is no other irrevocable transaction in the system, $(I, T) = (0, -)$.

The special case in which several transactions ask for the token at the same time is arbitrated by the controller queue of the core that owns the token. The owner of the first token request message found in such a queue is the one that gets the token. The rest of the token request messages are ignored and the requesters stalled. They will ask for irrevocability again after receiving a message of end of irrevocability.

4 Simulation Environment

The simulation environment comprises the full system simulator Simics [12], and the Wisconsin GEMS [14] toolkit that includes Ruby. Ruby is a multiprocessor memory system timing simulator, which we have modified to simulate the best-effort HTM system outlined in Section 2, and the proposals described in this paper.

² The cache controller comprises queues where coherence messages are buffered until they are served by the controller [18]. In this case, there are a mandatory queue that holds the messages from the CPU to the L1 cache, a request queue that holds request messages from/to the L1 cache and a response queue with response messages from/to the L1 cache.

Table 2: Workloads: Input parameters and transactional characteristics.

Bench	Input	# Xact	% Time in Xact	avg RS	avg WS
Bayes	-v32 -r1024 -n2 -p20 -i2 -e2 -s1	654	94%	87.64	48.91
Genome	-g512 -s32 -n32768	19496	85%	23.34	3.58
Intruder	-a10 -i16 -n4096 -s1	54933	92%	9.87	3.06
Kmeans	-m15 -n15 -t0.05 -i random-n2048-d16-c16	8235	46%	6.23	1.75
Labyrinth	-i random-x32-y32-z3-n96	222	100%	139.34	95.12
SSCA2	-s14 -i1.0 -u1.0 -l9 -p9	93721	13%	3.00	2.00
Vacation	-n4 -q60 -u90 -r16384 -t4096	4095	95%	63.20	10.16
Yada	-a20 -i633.2	5447	100%	62.45	38.21

The target system is organized as shown in Fig. 1. It comprises 16 in-order single-issue cores, with a private 32KB split 4-way L1 cache where the data cache holds two read and write transactional bits per 64B block. The L2 cache is unified, shared and divided into 16 banks of 512KB each. L2’s associativity is 8-way and it does not hold transactional information. The directory keeps a full bit-vector of sharers. Each thread is bound to a core, and so it is the operating system, so that there are not interferences such as migrations and context changes. Consequently, there is a maximum of 15 threads for the use of benchmarks.

The whole Stanford STAMP suite [16] was used for the evaluation. Table 2 shows the parameters and characteristics of the benchmarks. Namely, the number of transactions that successfully commits (# Xact), the percentage of time running transactions (% Time in Xact), and the average RS/WS (read set/write set) cardinality of the transactions, in cache blocks.

5 Software Fallback Path Evaluation

Figure 3 shows the fallback path code we have evaluated, which includes a variable to specify the number of transaction retries and the Lemming effect³ avoidance code [6, 11]. The code defines a thread’s local retry variable that is initialized to 0 (line 1). The retry limit is defined globally (RETRY_LIMIT). We define two primitives to begin a transaction: (i) TAKE_XACT_CHECKPOINT takes a register checkpoint where we want to resume the transaction on abort, but it does not start transactional bookkeeping; (ii) BEGIN_XACT begins transactional bookkeeping. Then, we can have non-transactional code between the two primitives to check whether we have to take the fallback path or not. The code to begin a transaction (lines 2-13) first takes a checkpoint and then increments the thread’s local retry variable. Next, if the number of retries is greater than the retry limit (line 5), the fallback path is taken by acquiring a single spin lock (line 6). If the retry limit is not reached, the code executes transactionally and adds the lock to the read set (line 10). The transaction is explicitly aborted if the lock is taken (line 11). It should be noted that the thread waits for the lock to be released just before beginning the transaction to avoid the Lemming effect (line 8). The code to end a transaction

³ If one transaction takes the fallback path, the others abort and wait for the fallback path lock to be released, i.e. a complete serialization of the ongoing transactions is carried out.

```

1 localRetries = 0;
2 void beginTransaction(&localRetries) {
3   TAKE_XACT_CHECKPOINT; //Return point on abort
4   localRetries++; //Increment xact retries
5   if (localRetries > RETRY_LIMIT) { //Fallback?
6     while(!lockAcquire(globalLock)) ; //Acquire lock
7   } else { // Execute transactionally
8     while(lock != 0) ; //Avoid Lemming effect
9     BEGIN_XACT;
10    if(lock != 0) //Add lock to the read-set
11      ABORT_XACT;
12  }
13 }
14 void endTransaction(localRetries) {
15   //If not retry limit, assume lock's elided
16   if (localRetries <= RETRY_LIMIT)
17     COMMIT_XACT; //Commit
18   else lockRelease(globalLock);
19 }

```

Fig. 3: Fallback code with retry limit and Lemming avoidance. Ticket lock alternative on the right.

(lines 14-19) checks the number of retries to execute either a transaction commit or a lock release.

On the right hand side of Figure 3 we show an alternative implementation of the fallback path which replaces the single spin lock by a two-variable *ticket lock* [15]. Each thread takes its own ticket before entering the critical section by atomically incrementing and reading the global ticket variable (line 6). Then, the thread waits for his turn by checking it against the global turn variable. The global turn is atomically incremented to release the lock (line 18). The implementation of the Lemming effect avoidance loop (line 8) is more accurate with the ticket lock as the thread waits not only when the lock is taken (`lock != 0`) but also when there is a queue of threads waiting to acquire the lock (`globalTicket >= globalTurn`).

Figure 4 depicts the speedup results obtained for those STAMP benchmarks that scale to some extent. The fallback code used is that of Figure 3, with or without Lemming avoidance (\pm Lemm) and with single or ticket lock. The lazy single lock approach [5] is also shown, which is the same as the single lock without Lemming avoidance but the lock is checked lazily at the end of transactions.⁴ The retry limit has been set to 5, which is a frequently used value [10, 21]. We have evaluated 3, 8 and 10 retries as well. An increased number of retries (8 or 10) seems to perform better when the number of threads, and therefore the contention, is high. For a low number of threads, a low number of retries suffices (3 retries up to 4 threads).

The results show that the fallback path versions with Lemming effect avoidance always beat the ones without it, due to the reduction in unnecessary serializations. As far as the type of lock is concerned, the ticket lock reveals itself as a good option since it reduces lock contention and ensures fairness in lock acquisition. But more importantly, the ticket lock provides the information of how many threads are waiting to enter the critical section and therefore, the Lemming loop waits for them to finish. Conversely, the single lock does not provide such information. Thus, the threads waiting at the

⁴ In this manner, multiple transactions are allowed to execute in parallel with the one in the fallback path, as long as such transactions commit after the lock release and the fallback code does not conflict with them.

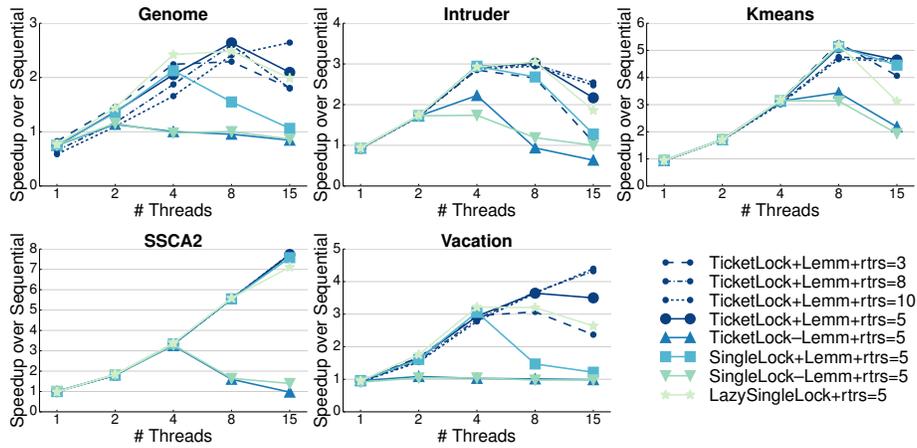


Fig. 4: Speedup over the sequential application for different fallbacks and parameters (Lemm: lemming effect avoidance, rtrs: number of retries).

Lemming loop may begin a transaction while other threads are contending for acquiring the lock. Those transaction will be aborted by the eventual lock acquisition. This fact is more probable in those benchmarks that spend a lot of time in transactions such as Genome, Intruder and Vacation (see Table 2), which take advantage of the ticket lock Lemming loop enhancement to avoid unnecessary aborts. SSCA2 and Kmeans are most of the time out of transactions and they are not affected by the type of lock. The lazy single lock yields good results since it encourages parallelism. However, the performance is worse than the ticket lock with Lemming effect avoidance as the number of threads increases, thus increasing the contention (e.g. Intruder, Kmeans and Vacation with 15 threads). The fallback conflicts with the concurrent transactions.

6 Hardware Irrevocability Mechanism Results

Figure 5 shows the speedup of the baseline BE-HTM system with the hardware irrevocability mechanism (Irre) and the software fallback path (Fback) with ticket lock and enhanced Lemming effect avoidance. The hardware irrevocability mechanism counter has been set to 5, as well as the retry counter of the fallback code. From these results we can classify the STAMP benchmarks in the following groups.

Bayes, Labyrinth and Yada

The speedup obtained for these benchmarks is barely that of the sequential version. And when there is only one thread the results are even worse than the sequential. The problem with performing worse than the sequential when we have only one thread running in the system is the number of retries before getting irrevocable or taking the fallback (set to 5 in this evaluation). With only one thread there is no abort due to conflicts, so all

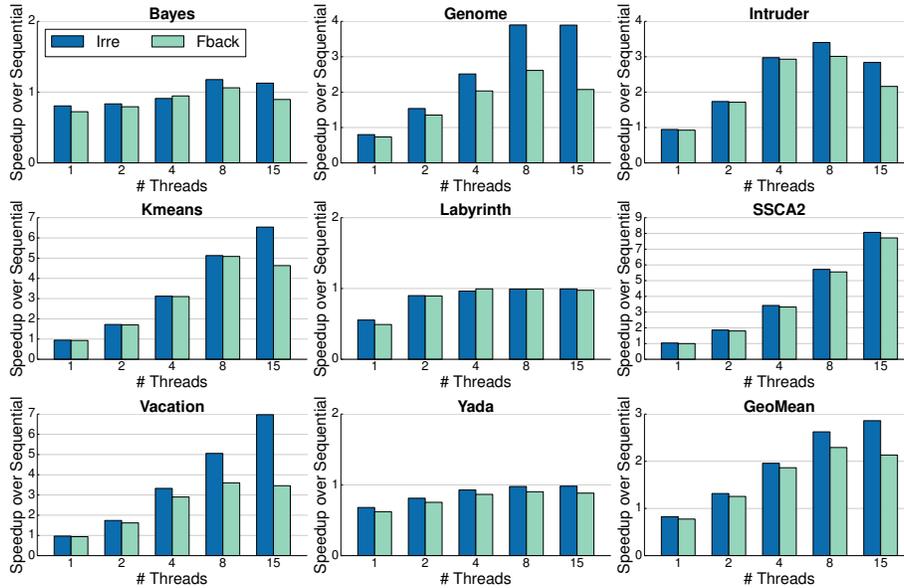


Fig. 5: Speedup of the hardware irrevocability mechanism (Irre) and the software fallback path (Fback) over the sequential application. The geometric mean is also shown (GeoMean).

aborts are because of capacity overflows, that are usually persistent. This can be avoided by maintaining different retry counters as stated in Nakaïke et al. [17], where they adapt the number of retries depending on the cause of abort. Three counters are used: one for aborts due to the fallback lock, a second for persistent aborts such as capacity aborts, and a third for transient aborts. In any case, the hardware irrevocability mechanism can implement different counters as well and it performs slightly better than the fallback path due to the last abort anticipation.

Although the irrevocable mechanism is better than the fallback one, these benchmarks do not scale because they exhibit large transactions in average, as shown by Table 2. In addition, Table 3 shows the number of irrevocable transactions and its cause, and the majority of them are due to L1 replacements. We can also see how the number of irrevocable transactions increases with the number of threads because of conflict aborts and capacity overflows due to L2 evictions (the latter primarily in Yada).

Kmeans and SSCA2

These two benchmarks scale well and behave similarly either by using hardware irrevocability or software fallback. This is due to the short time spent in transactions that amounts to 46% for Kmeans and only 13% for SSCA2, which reduces contention.

The size of transactions in Kmeans and SSCA2 is also a factor to consider. Their small transactions make that the fallback path or hardware irrevocability are barely taken. Actually, Table 3 shows 0 irrevocable transactions due to L1 and L2 replacements. However, contention makes some transactions to abort and take the fallback or

Table 3: Average number of irrevocable transactions, broken down into those due to L1 or L2 replacements, and those due to conflicts. Average number of aborts of irrevocability and fallback.

Bench	# Irrevocable Xacts (L1/L2/Conflicts)			Aborts(IRRE/FBACK)		
	4 th's	8 th's	15 th's	4 th's	8 th's	15 th's
Bayes	135(86/0/49)	148(68/0/80)	179(37/3/139)	653/728	788/1212	1040/1585
Genome	1203(1120/0/83)	1195(1017/0/178)	1679(1358/21/301)	5022/7942	5582/10821	8217/16321
Intruder	1055(22/0/1033)	3794(36/0/3759)	10562(110/1/10450)	11455/10428	35861/39628	77299/96499
Kmeans	400(0/0/400)	970(0/0/970)	1815(0/0/1815)	2193/2074	5425/6537	10296/19185
Labyrinth	97(76/0/21)	122(57/0/64)	160(44/0/116)	435/631	617/783	797/931
SSCA2	127(0/0/127)	283(0/0/283)	515(0/0/515)	657/575	1583/2140	3208/5486
Vacation	249(217/0/33)	347(280/0/68)	433(301/0/132)	1272/2924	1773/6221	2357/9874
Yada	1021(702/0/319)	1245(710/0/535)	1557(628/57/872)	4651/9128	5895/12561	7600/13643

the irrevocability mechanism when we have more threads. For this configurations we can see a slight benefit of irrevocability over the fallback version, or not so slight for Kmeans and 15 threads, because the irrevocability mechanism stalls the transactions instead of aborting them. Table 3 shows such an abort reduction that is up to 9000 transactions for Kmeans and 15 threads, which supposes an abort rate of 1.2 with irrevocability in contrast to the 2.32 of the fallback path.

Genome, Intruder and Vacation

For this group of benchmarks we obtain considerable benefits by using the BE-HTM system with hardware irrevocability over the fallback configuration. They are benchmarks with medium and small-sized transactions (Genome and Vacation) or that are more contended (Intruder). These characteristics can be noted in the number of irrevocable transactions that are due to replacements or conflicts in Table 3.

The hardware irrevocability mechanism not only performs better due to the anticipation of the last abort but also reduces the number of aborts by stalling non-irrevocable transactions instead of aborting them. Table 3 shows that the number of transaction aborts for the system with irrevocability is usually lower than that of its software fallback counterpart. The amount of wasted work is larger for the fallback path, specially for Genome and Vacation with 15 threads, where the abort reduction is more than 50%.

Summarizing, the BE-HTM system with irrevocability speeds up the execution about 2x with respect to the fallback path counterpart for Genome and Vacation, and it is around 20% better for Intruder and Kmeans, for 15 threads. The rest of the benchmarks yields similar or slightly better speedup by using irrevocability.

7 Related Work

Irrevocability in the context of HTM was first proposed in TCC [7] to deal with overflowed transactions. Blundell et al. [3] introduces OneTM-Serialized as a system where overflowed transactions gets irrevocable and serializes the system to ensure forward progress. Their implementation comprises a log-based HTM where the irrevocable transaction can be aborted as old data can be recovered from the log. They use a shared transaction status word residing in a fixed virtual location that acts as a mutex lock

to implement the irrevocability mechanism. We implement irrevocability with a token-based mechanism distributed through the cache controllers, in the context of a best-effort HTM system, comparing its performance with a software fallback path to gain insight into the hardware that could enhance the fallback.

IBM Blue Gene/Q HTM [19] ensures forward progress on capacity overflows and contention scenarios by means of an irrevocable mode. The irrevocability mechanism is implemented in a runtime system, thus freeing the programmer from the task of providing a fallback code. The runtime decides if a transaction gets irrevocable in an adaptive way. However, it has to abort a transaction to run it in irrevocable mode, whereas our hardware irrevocable mechanism anticipates the abort and initiates the irrevocable mode without wasting the work done so far by the transaction.

Afek et al. [2] propose a ticket-lock-based technique to improve the performance of Haswell's hardware lock elision (HLE). It is a different approach to our use of the ticket lock. In this case, the ticket lock guards the HLE lock and is acquired by those transactions that abort due to conflicts. Thus, the conflicting transactions are executed speculatively in turn, in parallel with the non-conflicting ones. After a given number of aborts, the transaction holding the ticket lock acquires the HLE lock and aborts all other transactions in the system. In fact, it is a contention management approach.

8 Conclusions

In this paper we propose a hardware implementation of an irrevocability mechanism to gain insight into the hardware enhancements that may speedup the execution of a fallback path in BE-HTM systems. We find that anticipating the abort that causes the execution of the fallback path and stalling the other transactions running in the system yields a significant improvement over the *abort-all* fallback solution.

On the other hand, we propose an enhanced Lemming effect avoidance loop by means of a ticket lock. A ticket lock provides precise information of how many threads are waiting to acquire the lock, so the separation of transactional and non-transactional execution can be performed more precisely.

We suggest having a hardware accelerated fallback path to retain both hardware benefits and software versatility. However, the possibility of having a hardware alternative to the software fallback path can be interesting for the user due to its simplicity.

Acknowledgement This work has been supported by the Government of Spain under project TIN2013-42253-P, Junta de Andalucía under project P12-TIC-1470, and the University of Málaga, Campus de Excelencia Internacional Andalucía Tech.

References

1. Intel 64 and IA-32 Architectures Optimization Reference Manual. Chapter 12.3: Developing an Intel TSX Enabled Synchronization Library (Sep 2014)
2. Afek, Y., Levy, A., Morrison, A.: Programming with hardware lock elision. ACM SIGPLAN Notices 48(8), 295–296 (aug 2013)

3. Blundell, C., Devietti, J., Lewis, E.C., Martin, M.M.K.: Making the Fast Case Common and the Uncommon Case Simple in Unbounded Transactional Memory. In: 34th Ann. Int'l. Symp. on Computer Architecture (ISCA'07). pp. 24–34. ISCA '07 (2007)
4. Cain, H.W., Michael, M.M., Frey, B., May, C., Williams, D., Le, H.: Robust Architectural Support for Transactional Memory in the Power Architecture. In: 40th Ann. Int'l. Symp. on Computer Architecture (ISCA'13). pp. 225–236 (2013)
5. Calciu, I., Shpeisman, T., Pokam, G., Herlihy, M.: Improved Single Global Lock Fallback for Best-effort Hardware Transactional Memory. In: 9th Workshop on Transactional Computing (TRANSACT'14) (2014)
6. Dice, D., Herlihy, M., Lea, D., Lev, Y., Luchangco, V., Mesard, W., Moir, M., Moore, K., Nussbaum, D.: Applications of the Adaptive Transactional Memory Test Platform. In: 3rd Workshop on Transactional Computing (TRANSACT'08) (2008)
7. Hammond, L., Wong, V., Chen, M., Carlstrom, B., Davis, J., Hertzberg, B., Prabhu, M., Wijaya, H., Kozyrakis, C., Olukotun, K.: Transactional memory coherence and consistency. In: 31th Ann. Int'l. Symp. on Computer Architecture (ISCA'04). pp. 102–113 (2004)
8. Harris, T., Larus, J., Rajwar, R.: Transactional Memory, 2nd edition. Morgan & Claypool Publishers (2010)
9. Herlihy, M., Moss, J.: Transactional memory: Architectural support for lock-free data structures. In: 20th Ann. Int'l. Symp. on Computer Architecture (ISCA'93). pp. 289–300 (1993)
10. Jacobi, C., Slegel, T., Greiner, D.: Transactional Memory Architecture and Implementation for IBM System z. In: 45th Ann. Int'l. Symp. on Microarchitecture (MICRO'12). pp. 25–36 (Dec 2012)
11. Liu, Y., Spear, M.: Toxic transactions. In: 6th Workshop on Transactional Computing (TRANSACT'11). ACM (2011)
12. Magnusson, P., Christensson, M., Eskilson, J., Forsgren, D., Hallberg, G., Hogberg, J., Larsson, F., Moestedt, A., Werner, B., Werner, B.: Simics: A full system simulation platform. *IEEE Computer* 35(2), 50–58 (2002)
13. Martin, M.M.K., Blundell, C., Lewis, E.: Subtleties of Transactional Memory Atomicity Semantics. *IEEE Computer Architecture Letters* 5(2), 17 (2006)
14. Martin, M., Sorin, D., Beckmann, B., Marty, M., Xu, M., Alameldeen, A., Moore, K., Hill, M., Wood, D.: Multifacet's general execution-driven multiprocessor simulator GEMS toolset. *ACM SIGARCH Computer Architecture News* 33(4), 92–99 (2005)
15. Mellor-Crummey, J.M., Scott, M.L.: Algorithms for scalable synchronization on shared-memory multiprocessors. *ACM Trans. on Computer Systems* 9(1), 21–65 (Feb 1991)
16. Minh, C., Chung, J., Kozyrakis, C., Olukotun, K.: STAMP: Stanford Transactional Applications for Multi-Processing. In: *IEEE Int'l Symp. on Workload Characterization (IISWC'08)*. pp. 35–46 (2008)
17. Nakaike, T., Odaira, R., Gaudet, M., Michael, M.M., Tomari, H.: Quantitative comparison of hardware transactional memory for Blue Gene/Q, zEnterprise EC12, Intel Core, and POWER8. In: 42nd Ann. Int'l. Symp. on Computer Architecture (ISCA'15). pp. 144–157 (2015)
18. Sorin, D.J., Hill, M.D., Wood, D.A.: *A Primer on Memory Consistency and Cache Coherence*. Morgan & Claypool Publishers, 1st edn. (2011)
19. Wang, A., Gaudet, M., Wu, P., Amaral, J.N., Ohmacht, M., Barton, C., Silvera, R., Michael, M.: Evaluation of Blue Gene/Q hardware support for transactional memories. In: 21st Int'l Conf. on Parallel Architectures and Compilation Techniques (PACT'12). pp. 127–136 (2012)
20. Welc, A., Bratin, S., Adl-Tabatabai, A.R.: Irrevocable transactions and their applications. In: 20th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA'08). pp. 285–296 (june 2008)

21. Yoo, R.M., Hughes, C.J., Lai, K., Rajwar, R.: Performance Evaluation of Intel Transactional Synchronization Extensions for High-performance Computing. In: Int'l Conf. on High Performance Computing, Networking, Storage and Analysis (SC'13). pp. 19:1–19:11 (2013)