Noname manuscript No.
(will be inserted by the editor)

Restricted Admission Control in View-Oriented Transactional Memory

Kai-Cheung Leung · Yawen Chen · Zhiyi Huang

the date of receipt and acceptance should be inserted later

Abstract This paper proposes a Restricted Admission Control (RAC) scheme for View-Oriented Transactional Memory. The scheme can control the number of threads concurrently accessing a view in order to reduce the number of aborts of transactions. The RAC scheme has the merits of both the locking mechanism and the transactional memory. A theoretical model is proposed to analyze the performance of the RAC scheme and to provide guidance for dynamic adjustment of the number of concurrent threads accessing the same view. Experimental results demonstrate that theoretical RAC model can mostly provide correct guidance to transactional concurrency control. Our RAC implementation shows that RAC can optimize concurrency control of transactions and performs much better than conventional transactional memory systems such as TinySTM that have no dynamic admission control.

Keywords View-Oriented Transactional Memory (VOTM) · transactional memory · deadlock · concurrency control · Restricted Admission Control (RAC)

1 Introduction

Parallel programming is becoming mainstream since multicore CPUs have become pervasive. There is a pressing need for parallel programming models to facilitate both performance and convenience. Traditional lock-based programming models can be made efficient but have tedious programmability and are prone to errors such as deadlock. New programming models based on transactional memory are more convenient, but may suffer from low performance [6, 21].

Traditionally locking [20, 24] is used for concurrency control, where multiple threads/processes[1] have to access a shared data object in an exclusive way. Atomic

Kai-Cheung Leung, Yawen Chen and Zhiyi Huang
Department of Computer Science
University of Otago
E-mail: {kcleung, yawen, hzy}@cs.otago.ac.nz

[1] In the rest of the paper, we use "thread" to mean both process and thread for simplicity since they are identical in terms of concurrency control.

access to a shared object is achieved through a locking mechanism. This lock-based concurrency control is generally regarded as pessimistic approach [29] where conflicts are resolved before they are allowed to happen. Even though locking is an effective mechanism to resolve conflicts, it could result in the deadlock problem if multiple objects are locked in different orders by multiple threads. Moreover, apart from the deadlock problem, fine-grained locks are tedious for programming, while coarse-grained locks often suffer from poor performance due to lack of concurrency.

To avoid the deadlock problem as well as to increase concurrency, Transactional Memory (TM) [15,23] was proposed for shared-memory programming models. In TM, atomic access to shared objects is achieved through transactions. All threads can freely enter a transaction, access the shared objects, and commit the accesses at the end of the transaction. If there are access conflicts among threads, one or more transactions will be aborted and rolled back. TM will undo the effects of the rolled-back transactions and restart them from the beginning. This transaction based concurrency control is labelled as an optimistic approach [4,19] where it is assumed nothing will go wrong and if it does go wrong deal with it later.

In terms of performance, both lock-based and TM-based approaches have their own merits in different situations. When access conflicts are rare (i.e., the contention is low), the TM-based approach has little roll-back overhead and encourages high concurrency since multiple threads can access different parts of the shared data simultaneously. In this situation, however, the lock-based approach has little concurrency due to the sequential access to the shared data, which results in low performance. To increase concurrency and performance, the programmer has to break the shared data into finer parts and use a different lock for each part. This solution using fine-grained locks often complicates the already-complex parallel programs and could incur deadlocks.

On the other hand, when access conflicts are frequent (i.e., the contention is high), the TM-based approach could have staggering roll-back overheads and is not scalable due to a large number of aborts of transactions. In such a situation, it is more effective to use the pessimistic lock-based approach to avoid the excessive operational overheads of transactions.

In order to adaptively improve performance of applications under various contention situations, we proposed a View-Oriented Transactional Memory (VOTM) paradigm [22] that seamlessly integrates the locking mechanism and transactional memory into the same programming model. VOTM is an extension of our previous work on View-Oriented Parallel Programming (VOPP) model [17,18,32] in the area of transactional memory. In VOTM, shared data objects are partitioned into "views" by the programmer according to the memory access pattern of a program. The grain (size) and content of a view are decided by the programmer as part of the programming task, which is as easy as declaring a shared data structure or allocating a block of memory space. Each view can be dynamically created, merged, and destroyed. The most important property for views is that they do not intersect with each other. Before a view is accessed (read or written), it must be acquired; after the access of a view, it must be released. This data-centric model bundles concurrency control and data access together and therefore relieves the programmer from controlling concurrent data access directly with either locks or transactions. When a shared data (i.e. a view) is to be accessed, the programmer just simply uses *acquire_view* to inform the system that the corresponding view is going to be accessed. It is up to the system to decide whether the locking mech-

anism should be adopted or a transaction should be started for the concurrent access of the shared data.

In VOTM, we adopt a Restricted Admission Control (RAC) scheme that can dynamically adjust the number of threads allowed to access the same view. With the RAC scheme, a view in VOTM is restricted to be accessed by a limited number of threads Q (called admission quota) whose value ranges from 1 to the maximum number of threads (N). If Q is 1, the threads access the set of data objects sequentially as in the lock-based approach. If Q equals N, the RAC scheme behaves like the conventional TM systems where any thread is allowed to start a transaction to access the data objects of the view. However, if Q is greater than 1 but smaller than N, only Q threads are allowed to access the data objects concurrently through transactions. If there are already Q threads accessing the data objects inside uncommitted transactions, other threads are excluded from accessing the set of data objects and have to wait until some existing transactions commit. In addition, RAC can flexibly adjust Q at runtime according to the contention situation, e.g., the number of transactional aborts, to achieve optimal performance, which will be described in details in Section 2.

This paper has the following contributions. First, we propose the novel Restricted Admission Control (RAC) scheme that adapts flexibly to runtime contention situations in order to achieve optimal performance for concurrent accesses to shared data objects in transactions.

Second, we propose a theoretical model for RAC to measure the contention levels and decide when Q should be adjusted to achieve optimal performance for TM applications. As far as we know, this is the first time that a theoretical analysis is applied to model admission control of transactions.

Third, we evaluate the RAC model with microbenchmarks and show the model can correctly decide if Q should be adjusted at various contention levels. Our experiment shows that this theoretical model is general enough to help measure the contentions for various TM systems.

The rest of the paper is organized as follows: Section 2 will present the RAC scheme and its theoretical model; Section 3 will evaluate the RAC model with microbenchmarks; Section 4 will discuss related work and Section 5 concludes the paper.

2 Restricted Admission Control

The RAC scheme is implemented for every view in VOTM. Each view consists of memory blocks that may store an entire linked list, tree or graph. Each view has an admission quota Q that restricts the maximum number of threads accessing the view concurrently. Before a view is accessed, the primitive *acquire_view* is used. If Q equals 1, *acquire_view* is equivalent to a lock acquisition. In this case, lock mechanism is used instead of the transaction mechanism to avoid transactional overheads. If Q is greater than 1, *acquire_view* will either start a new transaction or wait according to the following RAC scheme.

Suppose a view has an admission quota Q. We assume the current number of threads concurrently accessing the view is P. When the view is acquired through *acquire_view*, RAC follows the steps below:

1. Compare P with Q. If P is smaller than Q, increase P by 1, start a new transaction, and return with success.
2. If P equals Q, the calling thread is blocked until P becomes smaller than Q, and then goes to Step 1.

When the view is released through *release_view*, RAC executes the following steps:

1. Try to commit the transaction. If the commit fails, abort and roll back the transaction, decrease P by 1, and reacquire the view as shown above.
2. If the commit succeeds, decrease P by 1, and then return with success.

Furthermore, RAC can dynamically adjust the admission quota Q in the following way according to the contention situation. The admission quota Q of each view is initialized as the maximum number of threads (N). RAC regularly checks the contention situation. If the contention is high, RAC will relieve the contention of the view by halving the admission quota Q. This process can be repeated periodically until Q reaches 1, in which case the concurrency control is switched to the lock-based approach and the transaction mechanism is no longer used to access the view. Conversely, when the contention is low, RAC will increase concurrency by doubling Q. This process will repeat periodically until Q reaches N.

Obviously, to find out when to adjust Q is crucial to the performance of RAC. The following theoretical analysis helps understand when RAC can outperform conventional TM systems and when Q should be adjusted to achieve optimal performance.

2.1 RAC vs. conventional TM

Consider a set of *transactions* $S_T = \{T_1, ..., T_n\}$, which access the same view and are executed by N threads. The *duration* of transaction $T_i (1 \leq i \leq n)$ is denoted by t_i and refers to the time period that T_i is executed from start to commit without conflicts and interruptions. For simplicity of the analysis, we assume that, during the execution of T_i, the expected number of aborts is c_i and the average time spent by an aborted transaction is d_i, where $c_i, d_i \geq 0$. Therefore, the expected execution time for T_i is $c_i d_i + t_i$ in conventional TM that has no admission control of transactions.

Makespan is defined as the total time needed to perform all transactions [3]. Suppose that N threads are continuously executing the transactions, then the best possible *makespan* for S_T in conventional TM, denoted by $makespan_{TM}(S_T)$, can be calculated as

$$makespan_{TM}(S_T) = \frac{\sum_{i=1}^{n} c_i d_i + t_i}{N} \tag{1}$$

In RAC, Q transactions are allowed to be executed at any given time, where $1 \leq Q \leq N$. The expected execution time for T_i is $\frac{Q-1}{N-1} \times c_i d_i + t_i$, which can be proven as follows.

Suppose T_i aborts due to the conflict of shared memory location s accessed by $T_{i'}$ in conventional TM. However, in RAC, if T_i is allowed to access s at a given time, the probability that $T_{i'}$ is also allowed to access s is $\frac{Q-1}{N-1}$, because RAC allows only Q threads accessing s at any given time. So, the probability that T_i has 1 abort due to the conflict with $T_{i'}$ is $\frac{Q-1}{N-1}$. According to the binomial

distribution, the probability that T_i has k aborts ($k \in \{0, 1, ..., c_i\}$) is $p(k) = \binom{c_i}{k}(\frac{Q-1}{N-1})^k(\frac{N-Q}{N-1})^{c_i-k}$. Therefore, the expected execution time for T_i in RAC is $\sum_{k=1}^{c_i}(kd_i + t_i)p(k) = \sum_{k=1}^{c_i} kp(k)d_i + \sum_{k=1}^{c_i} p(k)t_i = \frac{Q-1}{N-1} \times c_i d_i + t_i$. (By the binomial distribution, $\sum_{k=1}^{c_i} kp(k) = \frac{Q-1}{N-1} \times c_i$ and $\sum_{k=1}^{c_i} p(k) = 1$)

Suppose the Q threads are continuously executing the transactions in RAC, then the *makespan* for S_T in RAC, denoted by $makespan_{RAC}(S_T)$, is

$$makespan_{RAC}(S_T) = \frac{\sum_{i=1}^{n} \frac{Q-1}{N-1} \times c_i d_i + t_i}{Q} \tag{2}$$

Therefore, the difference of $makespan_{RAC}(S_T)$ and $makespan_{TM}(S_T)$, denoted by Δ, can be obtained by Equation (1) and (2) as follows.

$$\begin{aligned} \Delta &= makespan_{RAC}(S_T) - makespan_{TM}(S_T) \\ &= \frac{\sum_{i=1}^{n} \frac{Q-1}{N-1} \times c_i d_i + t_i}{Q} - \frac{\sum_{i=1}^{n} c_i d_i + t_i}{N} \\ &= \frac{1}{N-1}(\frac{1}{N} - \frac{1}{Q})(\sum_{i=1}^{n} c_i d_i - \sum_{i=1}^{n} t_i(N-1)) \end{aligned} \tag{3}$$

Let $\delta = \frac{\sum_{i=1}^{n} c_i d_i}{\sum_{i=1}^{n} t_i(N-1)}$. It can be derived from Equation (3) that

(a) if $\delta > 1$, then $\Delta < 0$ and $makespan_{RAC}(S_T) < makespan_{TM}(S_T)$. That is, RAC outperforms conventional TM and the performance improvement is $|\Delta|$ when $\delta > 1$ (i.e., $\sum_{i=1}^{n} c_i d_i > \sum_{i=1}^{n} t_i(N-1)$). From this condition, it can be seen that RAC works especially well for transactions with high contention (c_i can be considered as the number of conflicts experienced by T_i), which will be verified in our experimental results.

(b) If $\delta \leq 1$, then $\Delta \geq 0$ and $makespan_{RAC}(S_T) \geq makespan_{TM}(S_T)$. That is, when $\delta \leq 1$, we should set Q to N in RAC. When Q equals to N, $\Delta = 0$ and RAC works the same as the conventional TM.

2.2 RAC with Q' threads vs. Q threads

Similar to the deduction of Equation (3), the difference between makespans of RAC using Q' (new) threads ($makespan_{RAC}(S_T, Q')$) and Q (previous) threads ($makespan_{RAC}(S_T, Q)$) is

$$\begin{aligned} &makespan_{RAC}(S_T, Q') - makespan_{RAC}(S_T, Q) \\ &= \frac{1}{Q-1}(\frac{1}{Q} - \frac{1}{Q'})(\sum_{i=1}^{n} c_i(Q) \times d_i(Q) - \sum_{i=1}^{n} t_i \times (Q-1)) \end{aligned} \tag{4}$$

where $c_i(Q)$ and $d_i(Q)$ are the expected number of aborts and the average time spent by an abort of T_i when using Q threads in RAC.

Let $\delta(Q) = \frac{\sum_{i=1}^{n} c_i(Q) \times d_i(Q)}{\sum_{i=1}^{n} t_i \times (Q-1)}$. It can be derived from Equation (4) that

(a) if $\delta(Q) > 1$ and $Q' < Q$, then $makespan_{RAC}(S_T, Q) > makespan_{RAC}(S_T, Q')$. That is, if $\delta(Q) > 1$, RAC should decrease Q to reduce the execution time of the concurrent transactions.

(b) if $\delta(Q) < 1$ and $Q' > Q$, then $makespan_{RAC}(S_T, Q) > makespan_{RAC}(S_T, Q')$. Therefore, to reduce the execution time of the concurrent transactions, RAC should increase Q.

In summary, the following theorem can be derived:

Theorem 1 *If $\delta(Q)$ is larger than 1, Q should be decreased; if $\delta(Q)$ is smaller than 1, Q should be increased, in order to reduce the makespan of S_T in RAC.*

In our implementation of RAC, $\sum_{i=1}^{n} c_i(Q) \times d_i(Q)$ is estimated with the total CPU cycles spent in aborted transactions, and $\sum_{i=1}^{n} t_i$ is estimated with the total CPU cycles spent in successful transactions. Therefore, $\delta(Q)$ is estimated with Equation (5) in RAC:

$$\delta(Q) = \frac{CPUcycles_{aborted\_tx}}{CPUcycles_{successful\_tx} \times (Q-1)} \tag{5}$$

3 Experimental evaluation

This experiment aims at verifying the above Theorem 1 with microbenchmarks. The experiment also shows that the RAC scheme works well in terms of dynamic adjustment of Q. We use the microbenchmark suite Eigenbench [16], and the real TM application Lee-TM [1] in our experiment.

3.1 Eigenbench

Eigenbench models transactions using orthogonal parameters, and allows a better understanding of which parameter contributes to a particular behaviour of the TM system.

For example, contention in Eigenbench is controlled by adjusting the size of the shared *hot_array* (A1) and number of read and write accesses to the *hot_array* in a transaction (R1 and W1 respectively). *hot_array* can be accessed by all transactions. High contention can be modelled by large number of accesses to the *hot_array* and/or small *hot_array* size (thus each element in the *hot_array* is more likely to be accessed by multiple concurrent transactions and have conflicts). The shared *mild_array* is also accessed by transactions, but each thread has its own subarray in the *mild_array* and each transaction can only access elements in the subarray owned by its thread, so access of the array will not cause conflict, but will increase transaction size.

Moreover, long transactions can be modelled with one or more of the following features:

- reading/writing to a large range of locations in shared memory;
- many repeated accesses to the same location(s) in shared memory;
- frequent access to local memory and/or high computation load inside transactions.

In Eigenbench, a transaction is modelled by a sequence of reads/writes to the shared memory with accesses to local memory and computation (represented by a number of NOPs) in between. A microbenchmark can also have computation and accesses to local memory outside transactions.

Below is the pseudocode outlining the Eigenbench model:

```
shared word hot_array[A1];    /* shared array where conflict occurs,
                                 accessed in tx */

shared word mild_array[A2];   /* shared array where each thread accesses
                                 its own subarray, so does not cause
                                 conflict, but still needs rollback
                                 should tx be aborted */

thread_local word cold_array[A3]   /* private to each thread, can be accessed
                                      either inside or outside tx.
                                      if accessed inside tx and tx aborted
                                      then needs to roll back changed */


each thread:

for loops:
do
        tx_start()
        perform r1 reads and w1 writes to the shared hot_array, and
                r2 reads and r2 writes to the shared mild_array
                in *random order*
                each access touches a random element (word) in
                the shared hot_array, or in the dedicated subarray within
                the shared mild_array

                between two accesses to shared arrays, there will
                also be r3i reads and w3i writes to the thread-local
                cold array, and NOPi instructions
        tx_end()

        /* activities outside transactions:
        perform r3o reads and w3o writes to the thread-local array
        perform NOPo instructions
done
```

In Eigenbench, each thread executes *loops* of iterations, and each iteration consists of a transaction, and then activities outside transactions.

In summary, Eigenbench allows us to model contentions quantitatively with various orthogonal parameters, which is ideal for us to evaluate the RAC model.

In order to verify the theoretical RAC model, we use four Eigenbench microbenchmarks: Highcon, FutileStall, StarvingElder and StarvingWriter. Highcon is configured by ourselves, but FutileStall, StarvingElder and StarvingWriter are taken from [5]. Parameters of each microbenchmark are listed in Table 1.

Table 1 Eigenbench parameters for the microbenchmarks

Application	N	loops	A1	A2	A3	R1	W1	R2	W2	R3i	W3i	R3o	W3o	NOPi	NOPo
Highcon	16	400	4k	1m	8k	100	10	20	20	10	20	0	0	20k	10m
FutileStall	16	10k	256	16k	8k	80	20	10	10	0	0	0	0	0	0
StarvingElder	1	100k	1k	1m	8k	128	32	20	20	0	0	100	100	0	0
	15	1m	1k	1m	8k	2	2	20	20	0	0	100	100	0	0
StarvingWriter	1	10k	32	1m	1m	0	30	0	0	500	0	0	0	0	0
	15	100k	32	1m	1m	30	0	0	0	0	0	0	0	0	0

Highcon has high contention long transactions that have computation loads both inside and outside transactions. It mimics realistic TM applications with transactions that are computationally intensive apart from having many high-contentious accesses to the *hot_array*. Each thread executes 400 transactions. In addition, this microbenchmark also has substantial computational workload *outside* transactions.

In FutileStall, transactions read from and then write to highly-contended data. Some transactions are stalled by other transactions that eventually will abort, and thus have futile stalls. In this microbenchmark, each thread executes 10000 transactions.

In StarvingElder, thread 0 executes 100,000 long transactions that read/write many locations in the *hot_array*, while each of threads 1-15 execute 1,000,000 short transactions that can conflict with and abort long transactions. As a result, these long transactions will make little progress.

In StarvingWriter, thread 0 executes 10,000 long writer transactions that write to the small *hot_array*, while each of threads 1-15 executes 100,000 read-only transactions that have frequent access to the *hot_array*, thus conflict with the writer transactions executed by thread 0. These read-only transactions will impede progress of the long writer transactions, since the multiple readers may continuously lock the same location for a long time.

3.2 Lee-TM

Lee-TM [1] is a printed circuit board autorouter application that has long transactions with high contention. A shared grid is allocated as the search space. Each routing transaction attempts to find and commit a route. If a transaction attempts to access a grid cell that is currently occupied by another transaction, the transaction will be aborted and restarted. The data file "mainboard.txt" is used in this experiment.

3.3 Verification of Theorem 1

In this experiment, we are going to verify if Theorem 1 can correctly suggest that the admission quota Q should be adjusted.

In the experiment, RAC is implemented over the TinySTM-1.0.0 [11] with both its default encounter-time locking algorithm (TinySTM-ETL) and its alternative commit-time locking algorithm (TinySTM-CTL). These two RAC implementations are denoted as "TinySTM-ETL+RAC" and "TinySTM-CTL+RAC" respectively. We use the two different implementations to show if the RAC model can work with different TM algorithms. To examine the correctness of Theorem 1, Q in RAC is fixed to 1, 2, 4, 8, 12 and 16 respectively, without dynamic adjustment during runtime. In all tests, the number of total threads (N) is fixed to 16. Thus the $Q = 16$ case is equivalent to the conventional TM that has no restriction of admissions. In Lee-TM and all microbenchmarks, we use only one view, which is sufficient to verify Theorem 1. All tests are carried out on a Dell PowerEdge R905 server with four AMD Opteron 8380 quad-core processors running with 2.5GHz and 16GB DDR2 memory. Linux kernel 2.6.32 and the compiler gcc-4.4 are used during benchmarking. All programs are compiled with the optimization flag -O2.

As mentioned previously, $\delta(Q)$ in Theorem 1 is estimated with Equation (5). Perfctr-2.6.42 [25] is used to measure the CPU cycles spent in aborted transactions and successful transactions.

For each Q, the runtime of each benchmark and its $\delta(Q)$ are presented in the following tables. To verify if Theorem 1 is correct for a benchmark, we compare

the runtimes of Q and Q', where $Q' < Q$. If $\delta(Q) > 1$ and the runtime of Q is larger than the runtime of Q', or alternatively if $\delta(Q) < 1$ and the runtime of Q is smaller than the runtime of Q', then Theorem 1 is correct; otherwise it is not quite correct.

For example, Table 2 shows the results of Highcon (on TinySTM-ETL+RAC) at different values of Q. When Q is 16, $\delta(Q)$ is 1.25, larger than 1, which suggests we should decrease Q according to Theorem 1. Compared with $Q = 12$, we find the runtime at $Q = 12$ is smaller than the runtime at $Q = 16$. This shows Theorem 1 has correctly suggested a smaller Q should relieve the contention situation and improve performance. Likewise, Theorem 1 is correct for $Q = 12$ and $Q = 2$.

The exceptions are $Q = 4$ and $Q = 8$. For $Q = 8$, $\delta(Q)$ is 0.95, slightly smaller than 1, but the experimental result suggests that a smaller Q (=4) can improve performance. Similarly, for case $Q = 4$, $\delta(Q)$ is smaller than 1, but a larger Q (=8) cannot improve performance. We attribute this inaccuracy to the approximation error of $\delta(Q)$, which will be discussed shortly in Section 3.3.1. To accommodate the error, we should use a critical zone with minimum and maximum thresholds instead of a critical point value (=1). If $\delta(Q)$ is in the critical zone, Q should not be adjusted. Our RAC implementation in Section 3.4 will suggest a suitable critical zone for $\delta(Q)$ that makes the RAC model work correctly for all cases.

Table 2 Highcon with ETL

Q	1	2	4	8	12	16
Runtime(s)	16.9	11.2	8.4	17.2	22.5	24.3
$\#abort$	0	1.55k	6.80k	130k	453k	643k
$\#tx$	6.4k	6.4k	6.4k	6.4k	6.4k	6.4k
$CPUcycles_{aborted\_tx}$	0	6.85G	27.9G	274G	611G	774G
$CPUcycles_{successful\_tx}$	41.1G	41.2G	41.2G	41.3G	41.4G	41.3G
$\delta(Q)$	N/A	0.17	0.23	0.95	1.34	1.25

Table 3 Highcon with CTL

Q	1	2	4	8	12	16
Runtime(s)	16.8	10.5	5.99	4.79	4.79	4.78
$\#abort$	0	660	2.03k	3.01k	3.02k	3.04k
$\#tx$	6.4k	6.4k	6.4k	6.4k	6.4k	6.4k
$CPUcycles_{aborted\_tx}$	0	4.16G	12.6G	19.2G	19.3G	19.3G
$CPUcycles_{successful\_tx}$	41.2G	41.2G	41.2G	41.2G	41.2G	41.2G
$\delta(Q)$	N/A	0.10	0.10	0.07	0.04	0.03

In Table 3 (TinySTM-CTL+RAC), the actual optimal value of Q is 16 with the smallest runtime of 4.78s. In the table, $\delta(Q)$ is around 0.1 or smaller for all values of Q. Therefore, RAC should keep increasing Q and use the maximum possible value (=16), which is consistent with Theorem 1.

It is worth noting that TinySTM-CTL does not cause as much false aborts and its actual contention is small. Therefore, there is no need for RAC to restrict admission.

Table 4 FutileStall with ETL

Q	1	2	4	8	12	16
Runtime(s)	1.37	2.84	5.31	14.2	63.3	237
$\#abort$	0	866k	4.52m	25.2m	185m	1.06G
$\#tx$	160k	160k	160k	160k	160k	160k
$CPUcycles_{aborted\_tx}$	0	7.28G	37.2G	211G	1.54T	8.05T
$CPUcycles_{successful\_tx}$	2.95G	5.06G	6.39G	8.82G	10.6G	11.9G
$\delta(Q)$	N/A	1.44	1.94	3.41	13.2	45.2

Table 5 FutileStall with CTL

Q	1	2	4	8	12	16
Runtime(s)	1.60	1.94	2.23	2.68	3.12	3.55
$\#abort$	0	163k	497k	1.19m	1.94m	2.69m
$\#tx$	160k	160k	160k	160k	160k	160k
$CPUcycles_{aborted\_tx}$	0	4.44G	14.9G	41.7G	80.4G	129G
$CPUcycles_{successful\_tx}$	3.54G	4.36G	5.31G	6.91G	8.38G	9.69G
$\delta(Q)$	N/A	1.01	0.94	0.86	0.87	0.89

From Table 4 and 5, we can see the actual optimal Q for FutileStall on TinySTM-ETL+RAC and TinySTM-CTL+RAC is 1. When Q is between 2 and 16 in TinySTM-ETL+RAC, $\delta(Q)$ is larger than 1. Thus Q should be kept decreasing to the smallest value (=1) which performs the best. Therefore, in these cases, Theorem 1 is correct.

However, for FutileStall on TinySTM-CTL+RAC in Table 5, when Q is between 4 and 16, $\delta(Q)$ is slightly smaller than 1, but decreasing Q can still further reduce runtime. This inaccuracy can be again attributed to the approximation error of $\delta(Q)$.

Table 6 StarvingElder with ETL

Q	1	2	4	8	12	16
Runtime(s)	173.2	88.1	49.6	27.2	21.9	21.8
$\#abort$	0	134k	515k	2.03m	3.74m	3.84m
$\#tx$	15.1m	15.1m	15.1m	15.1m	15.1m	15.1m
$CPUcycles_{aborted\_tx}$	0	3.30G	13.2G	49.0G	88.0G	92.1
$CPUcycles_{successful\_tx}$	381G	386G	399G	424G	451G	464G
$\delta(Q)$	N/A	0.009	0.011	0.016	0.018	0.013

Table 7 StarvingElder with CTL

Q	1	2	4	8	12	16
Runtime(s)	176.0	88.7	49.9	27.9	22.7	22.5
$\#abort$	0	67.1k	269k	1.11m	1.11m	1.11m
$\#tx$	15.1m	15.1m	15.1m	15.1m	15.1m	15.1m
$CPUcycles_{aborted\_tx}$	0	3.13G	13.5G	56.3G	56.3G	56.0G
$CPUcycles_{successful\_tx}$	388G	388G	402G	432G	453G	449G
$\delta(Q)$	N/A	0.008	0.011	0.019	0.011	0.008

Table 8 StarvingWriter with ETL

Q	1	2	4	8	12	16
Runtime(s)	14.6	11.4	9.99	9.37	9.22	9.18
$\#abort$	0	5.32m	15.7m	36.3m	56.4m	61.5m
$\#tx$	15.1m	15.1m	15.1m	15.1m	15.1m	15.1m
$CPUcycles_{aborted\_tx}$	0	3.43G	10.0G	23.2G	36.5G	40.1G
$CPUcycles_{successful\_tx}$	27.8G	27.8G	27.9G	28.1G	28.3G	28.3G
$\delta(Q)$	N/A	0.12	0.12	0.12	0.12	0.010

Table 9 StarvingWriter with CTL

Q	1	2	4	8	12	16
Runtime(s)	14.8	10.0	9.08	9.04	9.04	9.03
$\#abort$	0	1.34k	2.32k	2.69k	3.02k	3.22k
$\#tx$	15.1m	15.1m	15.1m	15.1m	15.1m	15.1m
$CPUcycles_{aborted\_tx}$	0	7.38m	16.4m	29.5m	45.6m	58.4m
$CPUcycles_{successful\_tx}$	27.8G	27.8G	27.9G	27.9G	27.9G	27.9G
$\delta(Q)$	N/A	0.0003	0.0002	0.0002	0.0001	0.0001

Long transactions in StarvingElder and writers in StarvingWriter can have poor progress. However, as shown in Table 6, 7, 8 and 9, the overall contention in both cases on TinySTM-ETL+RAC and TinySTM-CTL+RAC is not sufficiently high to justify admission control, as other short transactions still make good progress. Since $\delta(Q)$ is much smaller than 1 with all values of Q, Theorem 1 correctly predicts that it is inappropriate to restrict admission of transactions.

Table 10 Lee-TM with ETL

Q	1	2	4	8	12	16
Runtime(s)	82.6	69.7	162.2	livelock	livelock	livelock
$\#abort$	0	669	14.3k			
$\#tx$	1507	1507	1507	1507	1507	1507
$CPUcycles_{aborted\_tx}$	0	127G	1.23T			
$CPUcycles_{successful\_tx}$	205G	221G	243G			
$\delta(Q)$	N/A	0.56	1.69			

Table 11 Lee-TM with CTL

Q	1	2	4	8	12	16
Runtime(s)	82.3	69.7	62.9	56.4	52.0	50.6
$\#abort$	0	263	710	1.48k	2.12k	2.76k
$\#tx$	1507	1507	1507	1507	1507	1507
$CPUcycles_{aborted\_tx}$	0	134G	390G	822G	1.16T	1.52T
$CPUcycles_{successful\_tx}$	196G	208G	224G	255G	291G	318G
$\delta(Q)$	N/A	0.64	0.58	0.46	0.36	0.32

From Table 10, we can see the actual optimal Q for Lee-TM on TinySTM-ETL+RAC is 2. When Q is set to 4, $\delta(Q)$ is larger than 1 and RAC has correctly predicted that Q should be decreased. Lee-TM livelocks when Q is between 8 and 16, but if RAC is run in the adaptive mode (in Section 3.4), livelock will cause a very large $\delta(Q)$ and will drive Q down. When Q is 2, $\delta(Q)$ is smaller than 1 and RAC has correctly predicted that Q should not be further decreased.

In Table 11, the actual optimal Q for Lee-TM on TinySTM-CTL+RAC is 16, and $\delta(Q)$ is much smaller than 1 for all Q. Thus according to Theorem 1, Q should stay at 16, and is consistent with the actual optimal Q. Therefore in both cases, Theorem 1 is correct.

In conclusion, Theorem 1 based on the theoretical RAC model can correctly predicts whether Q should be adjusted in most cases for both TinySTM-ETL+RAC and TinySTM-CTL+RAC. However, for a few cases like Highcon on TinySTM-ETL+RAC and FutileStall on TinySTM-CTL+RAC, when $\delta(Q)$ is slightly smaller than 1, decreasing Q can still improve performance. This discrepancy between our

theoretical model and the experimental results will be discussed in the following section.

3.3.1 Approximation error of $\delta(Q)$

The RAC theoretical model assumes that the total number of transactions n is very large (towards infinity). Under this assumption, the number of aborts and the time spent in aborted transactions are linear to Q (the number of concurrent transactions), as proved in our theoretical analysis.

However, if the total number of transactions is finite and small, the more accurate formula for the number of aborts in the RAC model is $\frac{Q(Q-1)}{2} + (n - Q)(Q - 1)$, assuming the worst case where any transaction has conflict with any other transaction. Also the above formula assumes the TM system can manage conflicts well enough so that, every time a conflict occurs, at least one transaction can commit successfully. In this formula, if n moves towards infinity, the linear item $(n - Q)(Q - 1)$ becomes dominant and the quadratic item $\frac{Q(Q-1)}{2}$ becomes negligible, as our original model assumes.

On the other hand, if n is not large enough, the quadratic item of Q in the formula has important impact on the number of aborts. It can cause the discrepancy between our theoretical model and the experimental results when n, since our model has omitted it. The following is a further explanation how the item $\frac{Q(Q-1)}{2}$ can affect $\delta(Q)$.

Let $I(Q)$ represent the impact of the quadratic item $\frac{Q(Q-1)}{2}$ on the wasted time of the aborted transactions. If this impact is considered in $\delta(Q) = \frac{\sum_{i=1}^{n} c_i(Q) \times d_i(Q)}{\sum_{i=1}^{n} t_i \times (Q-1)}$, then we have $\delta'(Q) = \frac{\sum_{i=1}^{n} c_i(Q) \times d_i(Q) + I(Q)}{\sum_{i=1}^{n} t_i \times (Q-1)}$, where $\delta'(Q)$ is more accurate than $\delta(Q)$. When $\delta'(Q)$ is moving beyond 1, it means Q should be decreased according to Theorem 1. However, in such a situation, $\delta(Q)$ is still below 1 due to the missing item of $I(Q)$. This explains why Q should be decreased even when $\delta(Q)$ is slightly smaller than 1 in our experimental results.

In real applications, even though the total number of transactions is high, it is still not infinite. Moreover, the transactions do not arrive at the same time as our model assumed, since they can be very bursty in some applications. In these situations, the quadratic item of Q has its non-negligible impact on the experimental results. Therefore, in some of our benchmarks, decreasing Q can improve performance even when $\delta(Q)$ is slightly smaller than 1.

Though we know the reason of this approximation error, to include the quadratic item of Q will make the theoretical model intractable, as the mathematical deduction for $I(Q)$ is very complicated if not impossible. Thus, it is also challenging to produce the range (or interval) of the approximation error with mathematical deduction. Therefore, we choose to tolerate this error as long as the model plus our heuristics proposed in our implementation works reasonably well. A more complicated and accurate model could be an interesting but challenging attempt for future research.

3.4 Performance of the RAC scheme

We use Theorem 1 to guide the dynamic adjustment of Q in RAC. However, due to approximation error, the previous results indicate that even when $\delta(Q)$ is slightly smaller than 1, contention can still be high, and performance can still be improved by decreasing Q. For the same reason, conversely $\delta(Q)$ would need to be much smaller than 1 to indicate low contention to justify increasing Q. To accommodate the approximation error of $\delta(Q)$, we need to use a critical zone instead of a critical point for $\delta(Q)$.

In our implementation of RAC, we use a critical zone with two values MAX and MIN. Q will be decreased when $\delta(Q) > MAX$; and Q will be increased when $\delta(Q) < MIN$. MAX and MIN are set to 0.8 and 0.2 respectively based on our experimental results.

Additionally, in our implementation, to eliminate cache flushing overheads on incrementing the shared counter P (number of threads holding the view) in the RAC metadata, when it is clear that admission control to the view is unnecessary because of the following low contention condition, the RAC mechanism will be disabled.

- 20000 transactions are executed since Q is set to N, and
- $\delta(Q) < MIN$

After the RAC mechanism of the view is disabled, access to the view will no longer be restricted until the contention situation of the view changes.

Table 12 and 13 show the performance of the RAC scheme for the previous benchmarks. We compare the following three implementations: RAC with the above dynamic adjustment of Q (denoted as RAC in Table 12 and 13), no admission control ($Q = 16$, denoted as Q16), and RAC with static Q whose value is set to the optimal value of Q (denoted as OPT). We compare RAC with OPT to verify whether the RAC scheme can find the optimal Q in practice. Also we compare RAC with Q16 to show the performance improvement of RAC over conventional TM. Note that, in the tables, "$Q(RAC)$" is the Q picked up by the RAC with dynamic adjustment, and "$Q(OPT)$" represents the Q with which RAC works the best.

In all benchmarks, we use only one view, though multiple views would make RAC perform even better, as demonstrated in [22].

Table 12 Performance of Adaptive RAC in TinySTM-ETL

Benchmark	$time$(s) (RAC)	Q (RAC)	$\#aborts$ (RAC)	$time$(s) (Q16)	$\#aborts$ (Q16)	$time$(s) (OPT)	Q (OPT)	$\#aborts$ (OPT)
Highcon	8.52	4	8.86k	24.3	643k	8.43	4	6.80k
FutileStall	0.98	1	2.39k	237	1.06G	1.37	1	0
StarvingElder	22.4	16	3.85m	21.8	3.84m	21.8	16	3.84m
StarvingWriter	9.24	16	60.9m	9.18	61.5m	9.18	16	61.5m
Lee-TM	72.5	2	2.84k	livelock		69.7	2	1.51k

Table 13 Performance of Adaptive RAC in TinySTM-CTL

Benchmark	*time*(s) (RAC)	*Q* (RAC)	*#aborts* (RAC)	*time*(s) (Q16)	*#aborts* (Q16)	*time*(s) (OPT)	*Q* (OPT)	*#aborts* (OPT)
Highcon	4.80	16	3.02k	4.78	3.04k	4.78	16	3.04k
FutileStall	1.01	1	34.1k	3.55	2.69m	1.60	16	0
StarvingElder	22.4	16	1.12m	22.5	1.11m	22.5	16	1.11m
StarvingWriter	9.06	16	3.50k	9.03	3.22k	9.03	16	3.22k
Lee-TM	51.7	16	2.80k	50.6	2.76k	50.6	16	2.76k

From Table 12 and 13, it can be seen that in all benchmarks on both TinySTM-ETL and TinySTM-CTL, RAC has correctly settled to the optimal Q values. For example, in Highcon on TinySTM-ETL, RAC correctly settled to $Q = 4$ and has a 200% improvement over conventional TM (Q16). The RAC scheme settles to the correct Q value very quickly, as the runtime of RAC (8.52s) is similar to the runtime of OPT (8.43s). Also the number of aborts in RAC (8.86k) is only slightly more than OPT (6.80k).

In FutileStall, RAC has correctly settled to $Q = 1$ to in both TinySTM-ETL and TinySTM-CTL, and the runtimes in both cases (0.98s and 1.01s respectively) are shorter than OPT (1.37s and 1.60s respectively). This improvement is attributed to the turn-off of the transactional mechanism in RAC when Q reaches 1.

In StarvingElder, StarvingWriter and the TinySTM-CTL version of Highcon, the contention is not high enough to justify restricting admission and RAC has correctly stayed at $Q = 16$. The runtime of RAC is similar to OPT in all cases.

In the real TM application Lee-TM, RAC has correctly settled to $Q = 2$ and $Q = 16$ in TinySTM-ETL and TinySTM-CTL respectively. In the TinySTM-ETL version, the RAC scheme prevents livelock in Q16 by quickly adjusting Q to 2, and the runtime of RAC (72.5s) is close to the runtime of OPT (69.7s). In the TinySTM-CTL version, RAC has correctly stayed at $Q = 16$, and its runtime (51.7s) is close to the runtime of OPT (50.6s), indicating that RAC has low overheads.

From the above results, it can be seen that RAC based on Theorem 1 can quickly choose the optimal Q in different TM algorithms, and can therefore improve performance of TM applications with high contention. Lee-TM and our previous work [22] showed RAC was also able to improve performance of real TM applications.

4 Related work

Approaches to concurrency control in TM can be classified into three types: in-transaction conflict resolution, transactional scheduling, and adaptive locks.

4.1 In-transactional conflict resolution

In-transaction conflict resolution aims to resolve conflicts effectively to reduced wasted work of aborted transactions. All in-transaction conflict resolution algorithms, including both encounter-time locking (DSTM [14,28], SXM [12] and

McRT-STM [27]) and commit-time locking (TL-2 [8] and NOrec [7]) algorithms, resolve conflicts *within* a transaction only *after* these conflicts have been detected, but threads are still freely admitted into transactions. Therefore, the aborts cannot be stemmed in high contention and work is still wasted by transactions that eventually aborts, as shown in our experimental results.

4.2 Transactional scheduling

Transactional scheduling can control the admission of transactions when contention is high. It can prevent conflicts before they occur and therefore reducing wasted work on aborted transactions. For example, transaction scheduling algorithms such as [31] use a thread-local contention score. When a thread experiences high contention, it queues the starting transaction to a central scheduler, which will execute queued transactions serially. A similar approach is adopted in [9], except when a thread experiences high contention it uses a heuristic approach that predicts read and write sets of the starting transaction using read and write sets of previous transactions of the threads. If any address in the predicted read and write sets is being written by any other currently executing transactions, then the starting transaction will be queued to be executed serially. Otherwise, the transaction executes immediately. This algorithm relies on heuristic prediction of what will be read/written in the starting transactions. The admission control algorithm in [2] also adopts a similar approach. This family of transactional scheduling algorithms works orthogonally with the in-transaction conflict resolution algorithms mentioned above.

However, the above transactional scheduling algorithms use empirical thresholds to decide contention level. There is no theoretical model to guide the selection of those threshold values, as we discussed in this paper. As far as we know, the theoretical RAC model proposed in this paper is the first model that provides theoretical guidance to transactional concurrency control.

Furthermore, all transaction scheduling algorithms described above treat the entire TM with the same scheduling decision. Therefore, access to a low-contention shared object can be unreasonably restricted due to other high-contention shared objects in TM. Also the statistics collected for the entire TM are not as accurate as those collected per view basis, which will enlarge the estimation error of $\delta(Q)$ in the RAC model. However, RAC in VOTM treats each view individually and the estimation of $\delta(Q)$ is thus more accurate.

4.3 Adaptive locks

The speculative lock elision (SLE)-based model [26] was proposed to avoid unnecessary exclusive accesses in lock-based programs. An elidable lock can be acquired "speculatively" (using TM) or "non-speculatively" (using mutex). At any time, an elidable lock can be acquired speculatively by multiple threads, but only one thread can hold an elidable lock non-speculatively at any time.

The adaptive lock model in [30] has a similar approach, except a thread trying to acquire the lock in mutex mode must wait until all existing threads holding the lock in transaction mode to finish.

Like VOTM, both SLE and adaptive lock models have separate access control on each elidable lock, to ensure restrictions placed by the system on locks with high contention will not unnecessarily affect concurrency of accessing other elidable locks with low contention. However these models either allows all threads to hold the elidable lock in speculative mode, or exclusive access to one thread during non-speculative (mutex) mode, yet as shown in Sections 3 and 2, there are some cases where the optimal admission quota of a lock/view is actually between 1 and N. Therefore, the RAC scheme can achieve a superior performance by finding out the optimal admission quota to achieve maximal concurrency rather than only choosing between the two extremes – exclusive access to one thread or admitting all threads.

5 Conclusions and future work

As shown in the results in Eigenbench microbenchmarks and the real TM benchmark Lee-TM, our theoretical RAC model can mostly provide correct guidance to transactional concurrency control. The RAC scheme based on the model can improve performance of TM regardless of which underlying TM algorithm is used. In any TM algorithms, there will be situations where the contention becomes very high (the number of aborts becomes much larger than the number of transactions). In these situations, RAC will dynamically adjust the number of threads admitted to a view to control contention, thereby reducing works wasted by aborted transactions and improving progress. Experimental results show that RAC has superior performance to conventional TM because of the ability of RAC controlling admission and switching between TM and locking, whereas conventional TM has a performance issue when the contention is high and lock-based approach only works well in fine-grained locking but poorly in coarse-grained locking. Therefore, the RAC scheme enables VOTM to have better performance than the conventional TM and better convenience (and sometimes better performance) than lock-based programming.

One issue with RAC is blocking of threads by RAC when Q is smaller than N. This blocking seems to violate the lock-free or obstruction-free feature of TM systems [13]. Even though this feature is arguably necessary [10], RAC can quickly resolve this kind of blocking when the contention becomes low and thus Q is increased up to N, as long as Q does not become 1. If necessary, RAC can completely avoid blocking by using transactions even when Q equals 1, though it will lose some performance gain. In this way, if the system discovers that blocking is too long, the blocking can be easily lifted by increasing Q. Actually, in normal situations, the blocking in RAC is not worse than the live-locking in TM when transactions abort each other without progress under high contention.

As a future work, we will further investigate how to reduce the approximation error of $\delta(Q)$ in the RAC theoretical model. We will also investigate the performance of the RAC scheme for multiple views, which we believe is another strength of VOTM.

References

1. Ansari, M., Kotselidis, C., Jarvis, K., Luján, M., Kirkham, C., Watson, I.: Lee-TM: A non-trivial benchmark for transactional memory. In: Proceedings of the 7th International Conference on Algorithms and Architectures for Parallel Processing. LNCS, Springer (2008)
2. Ansari, M., Kotselidis, C., Jarvis, K., Luján, M., Watson, I.: Adaptive concurrency control for transactional memory. Tech. rep., University of Manchester (2007)
3. Attiya, H., Epstein, L., Shachnai, H., Tamir, T.: Transactional contention management as a non-clairvoyant scheduling problem. Algorithmica **57**, 44–61 (2010). URL http://dx.doi.org/10.1007/s00453-008-9195-x. 10.1007/s00453-008-9195-x
4. Bernstein, P., Goodman, N.: Concurrency control in distributed database systems. ACM Computer Survey **13**(2), 185–221 (1981)
5. Bobba, J., Moore, K.E., Volos, H., Yen, L., Hill, M.D., Swift, M.M., Wood, D.A.: Performance pathologies in hardware transactional memory. In: Proceedings of the 34th annual international symposium on Computer architecture, pp. 81–91. ACM, New York, NY, USA (2007). DOI http://doi.acm.org/10.1145/1250662.1250674. URL http://doi.acm.org/10.1145/1250662.1250674
6. Cascaval, C., Blundell, C., Michael, M., Cain, H.W., Wu, P., Chiras, S., Chatterjee, S.: Software transactional memory: Why is it only a research toy? Queue **6**, 46–58 (2008). DOI http://doi.acm.org/10.1145/1454456.1454466
7. Dalessandro, L., Spear, M.F., Scott, M.L.: NOrec: streamlining STM by abolishing ownership records. In: Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pp. 67–78. ACM, New York, NY, USA (2010). DOI http://doi.acm.org/10.1145/1693453.1693464
8. Dice, D., Shalev, O., Shavit, N.: Transactional locking II. In: Proceedings of the 20th International Symposium on Distributed Computing (2006)
9. Dragojević, A., Guerraoui, R., Singh, A.V., Singh, V.: Preventing versus curing: avoiding conflicts in transactional memories. In: Proceedings of the 28th ACM Symposium on Principles of Distributed Computing, pp. 7–16. ACM, New York, NY, USA (2009). DOI http://doi.acm.org/10.1145/1582716.1582725
10. Ennals, R.: Software transactional memory should not be obstruction-free. Tech. rep., Intel Corporation (2006)
11. Felber, P., Fetzer, C., Riegel, T.: Dynamic performance tuning of word-based software transactional memory. In: Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pp. 237–246. ACM, New York, NY, USA (2008). DOI http://doi.acm.org/10.1145/1345206.1345241
12. Guerraoui, R., Herlihy, M., Pochon, B.: Polymorphic contention management. In: Proceedings of the 19th International Symposium on Distributed Computing, pp. 26–29. LNCS, Springer (2005)
13. Guerraoui, R., Kapalka, M.: On obstruction-free transactions. In: 20th ACM Symposium on Parallelism in Algorithms and Architectures (2008)
14. Herlihy, M., Luchangco, V., Moir, M., Scherer III, W.N.: Software transactional memory for dynamic-sized data structures. In: Proceedings of the 22nd annual symposium on Principles of Distributed Computing, pp. 92–101. ACM, New York, NY, USA (2003). DOI http://doi.acm.org/10.1145/872035.872048
15. Herlihy, M., Moss, J.E.B.: Transactional memory: architectural support for lock-free data structures. SIGARCH Computer Architecture News **21**, 289–300 (1993). DOI http://doi.acm.org/10.1145/173682.165164
16. Hong, S., Oguntebi, T., Casper, J., Bronson, N., Kozyrakis, C., Olukotun, K.: Eigenbench: A simple exploration tool for orthogonal tm characteristics. In: Proceedings of the IEEE International Symposium on Workload Characterization (IISWC'10), pp. 1–11. IEEE Computer Society, Washington, DC, USA (2010). DOI http://dx.doi.org/10.1109/IISWC.2010.5648812. URL http://dx.doi.org/10.1109/IISWC.2010.5648812
17. Huang, Z., Cranefield, S., Purvis, M.K., Sun, C.: View-based consistency and its implementation. In: First IEEE International Symposium on Cluster Computing and the Grid, pp. 74–81 (2001)
18. Huang, Z., Purvis, M., Werstein, P.: Performance evaluation of view-oriented parallel programming. In: Proceedings of the 34th International Conference on Parallel Processing, pp. 251–258. IEEE Computer Society, Oslo (2005)

19. Kung, H., Robinson, J.: On the optimistic methods for concurrency control. ACM Transactions on Database Systems **6**(2), 213–226 (1981)
20. Lamport, L.: A new solution of Dijkstra's concurrent programming problem. Commun. ACM **17**(8), 453–455 (1974). DOI http://doi.acm.org/10.1145/361082.361093
21. Larus, J.R., Rajwar, R.: Transactional Memory. Synthesis Lectures on Computer Architecture. Morgan and Claypool (2007)
22. Leung, K., Huang, Z.: View-oriented transactional memory. In: The Fourth International Workshop on Parallel Programming Models and Systems Software for High-end Computing, in Proceedings of the 40th International Conference on Parallel Processing (2011)
23. Lomet, D.B.: Process structuring, synchronization, and recovery using atomic actions. In: ACM Conference on Language Design for Reliable Software, pp. 128–137 (1977)
24. Peterson, G.: Myths about the mutual exclusion problem. Information Processing Letters **12**(3), 115–116 (1981)
25. Pettersson, M.: The Perfctr Linux Performance Monitoring Counters Driver. Uppsala University (2004)
26. Roy, A., Hand, S., Harris, T.: A runtime system for software lock elision. In: Proceedings of the 4th ACM European Conference on Computer Systems, pp. 261–274. ACM, New York, NY, USA (2009). DOI http://doi.acm.org/10.1145/1519065.1519094. URL http://doi.acm.org/10.1145/1519065.1519094
27. Saha, B., Adl-Tabatabai, A.R., Hudson, R.L., Minh, C.C., Hertzberg, B.: McRT-STM: a high performance software transactional memory system for a multi-core runtime. In: Proceedings of the eleventh ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pp. 187–197. ACM, New York, NY, USA (2006). DOI http://doi.acm.org/10.1145/1122971.1123001
28. Scherer III, W.N., Scott, M.L.: Advanced contention management for dynamic software transactional memory. In: M.K. Aguilera, J. Aspnes (eds.) Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, pp. 240–248. ACM (2005). DOI http://doi.acm.org/10.1145/1073814.1073861
29. Tanenbaum, A., Steen, M.: Distributed Systems: Principles and Paradigms, Chapter 5. Prentice Hall (2002)
30. Usui, T., Behrends, R., Evans, J., Smaragdakis, Y.: Adaptive locks: Combining transactions and locks for efficient concurrency. In: Proceedings of the 18th International Conference on Parallel Architecture and Compilation Techniques. IEEE Computer Society, Washington, DC, USA (2009)
31. Yoo, R.M., Lee, H.H.S.: Adaptive transaction scheduling for transactional memory systems. In: Proceedings of the Twentieth Annual Symposium on Parallelism in Algorithms and Architectures, pp. 169–178. ACM, New York, NY, USA (2008). DOI http://doi.acm.org/10.1145/1378533.1378564. URL http://doi.acm.org/10.1145/1378533.1378564
32. Zhang, J., Huang, Z., Chen, W., Huang, Q., Zheng, W.: Maotai: View-oriented parallel programming on CMT processors. In: Proceedings of the 37th International Conference on Parallel Processing, pp. 636–643 (2008)