WTTM 2015

7th Workshop on the Theory of Transactional Memory

On Exploring Markov Chains for Transaction Scheduling Optimization in Transactional Memory

Pierangelo Di Sanzo, Marco Sannicandro, Bruno Ciciani, Francesco Quaglia

DIAG, Sapienza University of Rome

Sapienza University of Rome

HPDCS Research Group

Effects of Concurrent Execution of Transactions



Sapienza University of Rome

HPDCS Research Group

Identifying the optimal concurrency level ...

The optimal concurrency level depends on:

- transaction/workload profile (transaction length, data access distribution, read/write ratio, ...)

- hardware architecture

The workload profile may **change** during the execution of the application

the optimal concurrency level may **change** during the execution of the application



Transaction Scheduling: Transactions are blocked or allowed to run depending on some scheduling policy

The workload profile may change during the execution of the application



Adaptive scheduling approach: the scheduler takes decisions on basis of run-time observations



- Parametric system performance models [2]
- · Machine learning-based approaches [3]
- · Interpolating functions [4]

 Hill-climbing [5]
Pro-active transaction scheduling: serializing transactions when the abort probability is (estimated to be) high [6,7,8] **Pros:** the ability of predicting the system throughput for different system configurations allows at run-time to quickly "jump" to the best scheduler configuration

Cons: these approaches require <u>a priori analysis phases (e.g.</u> <u>training phases</u>), during which various parameters have to be measured while running the application with different workload profiles and/or different scheduler configurations (in terms of, e.g., number of admitted transactions).

Pros: no a priori analysis phases of the application are required

Cons:

- \cdot the optimal solution is not guaranteed
- time to converge
- the user has to configure some parameters (e.g. conflict rate/abort probability thresholds) on the basis of which the scheduler takes decisions

Effects of Different Thresholds of Scheduling Algorithms



Hardware configuration: 16-cores HP ProLiant server, equipped with 2GHz AMD Opteron 6128 processors, 64 GB of RAM and the Linux operating system (kernel version 2.7.32-5-amd64). Application configuration: Intruder - input: -a10 -l128 -n262144 Yada - input: -a15 -i yada/inputs/ttimeu1000000.2

Sapienza University of Rome

HPDCS Research Group

A system performance model ...

 \cdot for predicting the system throughput depending on the concurrency level, ...

 \cdot that can be instantiated on-the-fly (no a priori observation phases required), \ldots

 \cdot easy to be re-configured when the workload profile changes, ...

and

 \cdot that does not require a "skilled" user for setting up the optimal scheduler configuration.

 $\cdot \mathbf{N}$ running threads

• each thread can execute both **transactions** or **non-transactional code** (*ntc*) **blocks.**

 \cdot a transaction is aborted and restarted upon conflict

• Transaction scheduling policy:

- the scheduler accepts at most m (with $m \le N$) concurrent transactions (other transactions are blocked)
- a blocked transaction is allowed to run when another running transaction commits.

Modelling the system behaviour through a set of *states* and *states transitions*

Continuous-time homogeneous Markov Chain (CTMC) with **N** states (finite state space)

 \cdot state k of the CTMC represents a state of the system when there are k threads executing transactions (both running or blocked transactions)

 \rightarrow when the system is in state **k** there are **N-k** threads executing *ntc* blocks.



A lightweight Markov Chain-based Performance Prediction Model

- A *transition* from state **k** to **k+1** occurs when a threads starts a new transaction
- A *transition* from state *k* to state *k-1* occurs upon the commit of whichever running transaction



t_{ntc:}: average execution time of ntc blocks → inter-arrival rate of transactions along any thread $\lambda = \frac{1}{t_{ntc}}$ → transition rate from state **k** to **k+1**: $\lambda_k = (N - k) \cdot \lambda$



 $t_{k:}$: average transaction execution time when there are k executing transactions →

transaction execution rate of a thread for state **k**: $\mu_k = \frac{1}{t_k}$



A lightweight Markov Chain-based Performance Prediction Model

· for any state $k \le m$, since exactly k transactions are running (i.e. none is blocked), the transition rate from state k to state k - 1 is $\gamma_k = k \cdot \mu_k$



· for any state k > m, the there are m running transactions (the other k - m transactions are blocked \rightarrow

for all states such that k > m, the *transition rate* from the state k to the state k - 1 is $\gamma_k = m \cdot \mu_k$

Transaction execution time when the system is in state *k*:



 $\{q_k: 0 \le k \le N\}$ stationary probability vector of the CTMC



System throughput thr_m when the scheduler admits at most m running transactions:

$$thr_m = \sum_{i=1}^N q_k \cdot \gamma_k$$

Model Instantiation

To instantiate the model, values of parameters $u_{t,k}$, $w_{t,k}^r$, t_{ntc} and p_k have to be known.

They can be can be measured, for each k, by observing the system running with a **fixed scheduler configuration** (i.e., admitting at most m concurrent transactions)



What-if Analysis

Once above parameters have been measured, we can use the model for predicting the system throughput for **different scheduler configurations** (i.e. for different values of m). This can be done by modifying the value of m of the CTMC and solving the model.

Sapienza University of Rome

Experimental Evaluation: model validation with STAMP and TinySTM



HPDCS Research Group

Sapienza University of Rome

What-if Analysis with STAMP and TinySTM



Sapienza University of Rome

HPDCS Research Group

Thank you

References

[1] Diego Rughetti, Pierangelo di Sanzo, Bruno Ciciani, and Francesco Quaglia. Machine learning-based selfadjusting concurrency in software transactional memory systems. In Proc. 20th Int. Symposium on 4Modeling, Analysis and Simulation of Computer and Telecommunication Systems, pages 278–285. IEEE, 2012.

[2] Aleksandar Dragojeví c and Rachid Guerraoui. Predicting the scalability of an stm a pragmatic approach.

In Proc. 5th ACM Workshop on Transactional Computing. ACM, 2010.

[3] Pierangelo Di Sanzo, Francesco Del Re, Diego Rughetti, Bruno Ciciani, and Francesco Quaglia. Regulating concurrency in software transactional memory: An effective model-based approach. In Proceedings of the Seventh IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO), 2013.

[4] Pierangelo Di Sanzo, Francesco Del Re, Diego Rughetti, Bruno Ciciani, Francesco Quaglia Regulating Concurrency in Software Transactional Memory: An Effective Model-based Approach. In Proc. Seventh IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2013)

[5] Diego Didona, Pascal Felber, Diego Harmanci, Paolo Romano, and Joerg Schenker. Identifying the optimal level of parallelism in transactional memory systems. In Proc. International Conference on Networked Systems, NETYS. Springer, 2013.

[6] Mohammad Ansari, Christos Kotselidis, Kim Jarvis, Mikel Luj´an, Chris Kirkham, and Ian Watson. Advanced concurrency control for transactional memory using transaction commit rate. In Proc. 14Th Int. Euro-Par Conference on Parallel Processing, pages 719–728. Springer-Verlag, 2008.

[7] Aleksandar Dragojevi' c, Rachid Guerraoui, Anmol V. Singh, and Vasu Singh. Preventing versus curing: avoiding conflicts in transactional memories. In Proc. 28th ACM Symposium on Principles of Distributed Computing, pages 7–16. ACM, 2009.

[8] Richard M. Yoo and Hsien-Hsin S. Lee. Adaptive transaction scheduling for transactional memory systems. In Proc. 20th Symposium on Parallelism in Algorithms and Architectures, pages 169–178. ACM, 2008.