Architecture Support and Comparison of Three Memory Consistency Models in NoC based Systems

Abdul Naeem, Axel Jantsch and Zhonghai Lu

Department of Electronic Systems, Royal Institute of Technology (KTH), Sweden

E-mail: {abduln, axel, zhonghai}@kth.se

Abstract—We propose a novel hardware support for three relaxed memory models, Release Consistency (RC), Partial Store Ordering (PSO) and Total Store Ordering (TSO) in Network-on-Chip (NoC) based distributed shared memory multicore systems. The RC model is realized by using a *Transaction Counter* and an *Address Stack* based approach while the PSO and TSO models are realized by using a *Write Transaction Counter* and a *Write Address Stack* based approach. In the experiments, we use a configurable platform based on a 2D mesh NoC using deflection routing policy. The results show that under synthetic workloads, the average execution time for the RC, PSO and TSO models in 8x8 network (64 cores) is reduced by 35.8%, 22.7% and 16.5% respectively, over the Sequential Consistency (SC) model. The average speedup for the RC, PSO and TSO models in the 8x8 network under different application workloads is increased by 34.3%, 10.6% and 8.9%, respectively, over the SC model. The area cost for the TSO, PSO and RC models is increased by less than 2% over the SC model at the interface to the processor.

Keywords— *Memory consistency; Release consistency; Scalability; Distributed shared memory; Network-on-Chip*

I. INTRODUCTION

The parallelization of computation, communication and memory architecture has to be matched [1]. The full potential can be harvested with Distributed Shared Memory (DSM) on-chip by exploiting the distributed nature of Network-on-Chip (NoC) based systems. Since shared memory operations can be reordered in the network, the DSM systems may show unexpected behavior. Memory consistency models define the execution order for the shared memory operations [2]. The strict *Sequential Consistency* (SC) model [3] does not take advantage of potential performance benefits in the DSM systems. As a result, several *relaxed* consistency models [2][7-8] emerged to exploit the system optimizations by relaxing the ordering constraints on the shared memory operations.

Memory consistency and cache coherence are two distinct problems. Both aim to achieve consistent view of the memory system but at different levels. The cache coherence problem arises due to different cached copies of the same shared data. Memory consistency in contrast is related to the ordering constraints on the shared memory operations for the expected behavior of the DSM systems. In some situations, where these two problems have very different requirements (e.g. on the size of the cache block and the consistency object), or when a cache is not used (e.g. for hard real time applications) an independent implementation of the memory consistency and cache coherence is preferred [16][17].

Furthermore, heterogeneous and customized systems have different design constraints and requirements than general multiprocessor systems. The former have tighter power constraints, require less but more heterogeneous memory, make

less or no use of caches, and have often soft or hard real-time constraints. In the context of customized NoC based multicore (McNoC) systems, this paper studies the memory consistency issue with the following contributions:

- A novel realization scheme of the Release Consistency (RC) model which is independent of the cache coherence protocol is proposed. In contrast to [17], reordering among the outstanding shared memory operations issued by a processor to the *same* location in the memory is avoided. To that end, a hardware structure *A-Stack* (address stack) is used in each node of the network (refer to Section V).
- The Total and Partial Store Ordering (TSO, PSO) models are realized using Write Transaction Counter (*WTC*) and Write Address Stack (*WA-Stack*) in each node of the network (more detail is in Section V).
- Performance of the RC, PSO, TSO and SC models is evaluated in the NoC based systems with 1 to 64 cores.

For the experiments, a configurable McNoC platform is used with DSM, distributed locks and on-chip 2D mesh *Nostrum* network [4] using deflection routing policy. The performance of memory models are compared in the McNoC systems. The experimental results show the performance gain of the three relaxed memory models over the SC model.

The paper organization is as follows. The next section overviews the related work. In section III, the TSO, PSO and RC models are discussed. In section IV, the DSM based McNoC platform is introduced. The realization schemes of the three consistency models are presented in section V. In section VI, simulation results and performance analysis of these consistency models are described in the McNoC systems and finally, section VII summarizes our contribution.

II. RELATED WORK

In general multiprocessor systems, several memory consistency models are discussed in the literature [2-9]. Adve et al. [2] discussed the memory consistency models from the system optimizations point of view. The SC model [3] enforces a total order on the memory operations. The TSO model [5][6] relaxes the ordering constraint in the case of a *write followed by a read* operation. The PSO model [5] further provides relaxation among the *write* operations. The ordering constraints on the memory operations under both the TSO and PSO models are enforced using different kind of *fences* (non-memory references). For instance, the SPARC architectures [5] use (MEMBARs, SBAR), and the x86 architectures [6] use (MFENCE, SFENCE, LFENCE) fences. The weak consistency model [7] classifies the shared memory operations as *data* and *synchronization* operations. The RC model [8] further classifies the synchronization operations as *acquire* and *release* operations. The DASH project [9] implements the RC model

using tracking mechanism via several counters and is dependent on the cache coherence protocol. The directory based coherence protocol is used to maintain the status informations of cache blocks. Recent work [10] on the directory based coherence protocols aims to reduce the directory overheads in terms of area, energy and power consumption. Of late, address translation aware memory consistency models at physical and virtual address levels have been proposed [11]. They focus on the detection of the design and runtime faults due to the address translation.

In NoC based systems, the proposed mechanism in [12] is very strict and allows one outstanding transaction of an initiator at a time in the network. Streaming consistency [13] is based on the software cache coherence protocol, where synchronization sections can overtake each other. However, polling and updating the circular buffer administration at each request level may not be feasible in the larger systems. The AXI [14] and OCP [15] protocols enforce the *ordering models* by using transactions IDs and thread IDs, respectively. In [14], transactions of the same master with *different* IDs can be reordered, but, transactions with the *same* ID are not allowed to be reordered. In [15], *tagged* transactions of the same master using thread IDs are allowed to be reordered, but, *non-tagged* transactions are strictly ordered. The SC model is realized in the McNoC systems by stalling the processor on issuance of each memory operation till its completion [16]. In [17], a Transaction Counter (TC) based approach is used to realize the RC model in the McNoC systems. The TC is used to keep track of the outstanding data operations issued between two consecutive release points. In this paper, an additional hardware structure *A-Stack* is used to avoid reordering of transactions to the same address to ensure parallel program correctness in [17], and the TSO and PSO models are realized by using the *WTC* and *WA-Stack* hardware structures.

III. TSO, PSO AND RC MODELS

The ordering constraints under the SC, TSO and PSO models are compared in Figure 1. The variables (A, B, C, D, E, F) are ordinary shared memory variables and the variable S is a special synchronization (lock) variable. The variables to the left side of the assignment operators are written and those to the right side are read. An *arrow* between the two variables indicates an ordering constraint between the operations on those variables. For instance, $A \rightarrow B$ indicates that an operation on A is followed by an operation on B in the program, and an operation on A is completed before the issuance of an operation on B. The two operations are not allowed to be reordered with each other.

Figure 1. Comparison of SC, TSO and PSO models

According to the SC model (Figure 1(a)), the shared memory operations are completed in the *program order*. The *sequential order* is maintained by interleaving operations on lock S among processors in the system.

A. TSO Model

The TSO model (Figure 1(b)) allows the write operation on C to be *reordered and overlapped* with respect to the following read operation on D. The TSO model compared to the SC model allows reordering and relaxation in the case of a *write followed by read* operation. It enforces the ordering constraints in the cases of a *read followed by a write* ($A \rightarrow B$), a *write followed by a write* ($E \rightarrow F$), or a *read followed by a read* operation. In addition, ordering constraints with respect to the synchronization (Sync) operations must also be enforced. The global orders to be enforced under TSO model are given in Figure 2(a). We refer to these global orders in section V.

Figure 2. Global orders under: a) TSO model b) PSO model

B. PSO Model

The PSO model [5] is a refinement of the TSO model. As demonstrated in Figure 1(c), the PSO model further eliminates the ordering constraint in the case of a *write operation on E followed by a write operation on F*. It allows additional reordering among the *write* operations compared to the TSO model. The PSO model enforces the global orders on shared memory operations as shown in Figure 2(b).

C. RC Model

The RC model [8][17] is a refinement of the weak consistency model [7][16]. It classifies the *synchronization* operations as *acquire* and *release* operations. Acquire and release operations are related to the special synchronization variables (locks, semaphores) maintained in the shared address space. The *data* operations are the read and write operations related to the ordinary shared variables.

Figure 3. a) RC model b) Global orders under RC model

As illustrated in Figure 3(a), according to the RC model, the independent data operations on (A, B) are allowed to be reordered with each other, with the acquire operation on lock S and with the data operations on (C, D) in the critical section. They are not permitted to be reordered with respect to the

release operation on lock S. The data operations (C, D) can be reordered and overlapped with respect to each other, but, they are not allowed to be reordered with the acquire and release operations on lock S. The data operations on (E, F) are allowed to be reordered with each other, with the prior outstanding release operation on lock S and with the prior outstanding data operations on (C, D). However, they are not permitted to be reordered with respect to the prior acquire operation on lock S. The global orders to be enforced on the shared memory operations under the RC model are given in Figure 3(b).

A homogenous McNoC system is shown in Figure 4(a). All the nodes are interconnected via a packet-switched network. As demonstrated in Figure 4(b), each processor-memory (PM) node has a processor, transaction controller, synchronization handler, network interface and the local memory.

Figure 4. a) Homogeneous McNoC b) PM node

The platform uses 2D mesh packet switched *NoStrum* NoC [4] with an adaptive routing algorithm. The Network Interface (NI) connects a PM node to the network. It deals with the transactions from the processor and performs packetization, queuing, arbitration and message passing over the network. It also receives the packet from the network, de-packetizes it and hands it over to the processor or memory system. The shared memory is distributed across the network. The local memory is connected to the local processor and NI, respectively. All the shared parts in the local memories form the DSM in a single global address space. The platform also uses the distributed locks in the network. The Synchronization Handler (SH) controls N locks maintained in the global address space. Every lock is accessed in a sequential order by multiple processors in the system. The synchronization (acquire, release) requests to the SH either come from the local processor or from the remote processor via the network. If the requested lock is available then it is acquired. Otherwise, a negative acknowledgement is sent back and the source node sends again the same request until the lock is gained. A release request makes the lock available for the next acquire on it. The local data and synchronization operations are accomplished within the node. For the remote accesses, message passing is carried out to the remote node via the network. The transaction controller (TCTRL) is a customized interface which integrates the processor with the rest of the system like any standard interface [14][15]. It also implements the *memory consistency protocols* using the hardware structures like (*TC*, *Address-Stack*). The TCTRL is developed specifically for the LEON3 IP core [18] which is used in each node of the network.

A. TSO Model

The TSO model is realized (Figure 5) in the McNoC platform by enforcing the required global orders in Figure 2(a). **Write** $\rightarrow$ **Write**: A Write Transaction Counter (*WTC*) is used in each node of the network to keep track of the outstanding write operations issued by a processor. The *WTC* is incremented by the issuance of a write operation (1). It is decremented by the completion of a write operation (2). The *WTC* is checked at the issuance of each write operation and the write operation is delayed by stalling the processor till the completion of previously issued write operation (*WTC*=0). The processor is stalled by issuing an active low *keep-transaction-going* signal by the TCTRL in Figure 4.

Read $\rightarrow$ Read/Write: These global orders are enforced by *stalling* the processor on the issuance of a shared memory read operation (3) till its completion by returned data (4). The issuances of the subsequent *read* and *write* operations are delayed till the completion of previously issued read operation.

Figure 5. Realization scheme of TSO model

In order to ensure the parallel program correctness, outstanding memory operations to the same location in the memory under the TSO model are constrained to accomplish as per *program order*. A hardware structure *WA-Stack* (Write Address Stack) is used in each node of the network to this end. The *WA-Stack* keeps track of the *addresses* to be accessed by the previously issued outstanding *write* operations. At the issuance of a write operation (1) the address to be accessed by the write operation is pushed on the *WA-Stack*. On the completion of the write operation (2) the address is popped from the *WA-Stack*. The issuance of a *read* operation (3) checks the *WA-Stack*. If the address is on the *WA-Stack*, then the issuance of the read operation is delayed until the same address is popped from the *WA-Stack*. The address is popped from the *WA-Stack* on the completion of previously issued write operation to the same location in the shared memory.

The ordering constraints with respect to the Sync operations are also enforced. The shared memory operations are completed before the issuance of a Sync operation and vice versa. The issuance of a Sync operation checks the *WTC* and the Sync operation is delayed till the completion of previously issued outstanding write operations ($WTC=0$). On the issuance

of a Sync operation (5) the following memory operations are delayed by stalling the processor till its completion (6).

B. PSO Model

The PSO model is realized by enforcing the required global orders as shown in Figure 2(b). The ordering requirements under the PSO and TSO models are mostly similar except the PSO model further relaxes the ordering constraint in the case of a write followed by a write operation.

Figure 6. Realization scheme of PSO model

As illustrated in Figure 6, the issuance of a write operation (1) is not delayed till the completion of previously issued outstanding write operation ($WTC=0$). It allows multiple outstanding write operations on the network, which are not allowed under the TSO model (Figure 5). The independent shared memory write operations can be reordered with respect to each other. The PSO model also uses *WA-Stack* to prohibit the reordering among the shared memory write operations to the same location in the shared memory. There is an additional check on the *WA-Stack* under the PSO model. The *WA-Stack* is checked on the issuance of each write operation (1). If the address to be accessed by the write operation is on the *WA-Stack*, then the issuance of the write operation is delayed until the same address is popped from the *WA-Stack* by the completion of a previously issued outstanding write operation to the same location. The rest of the realization scheme is similar to that described under the TSO model.

C. RC Model

The RC model is realized (Figure 7) in the McNoC platform by enforcing the required global orders as given in Figure 3(b).

Data → Release: A Transaction Counter (*TC*) in each node of the network keeps track of the outstanding data operations issued before a release operation. The *TC* is incremented by the issuance (1) of a data operation. It is decremented by the completion (2) of a data operation. The *TC* is checked at the issuance of a release operation (5) and the release operation is delayed by stalling the processor till the completion of previously issued outstanding data operations ($TC=0$).

Acquire → Data/Release: These global orders are enforced by stalling the processor upon the issuance of an acquire operation (3) till the acquisition of a lock (4). The lock must be gained by a processor before entering to the critical section and also before trying to release it.

Release → Acquire: This global order is enforced by sequential order on a lock in the multiprocessors system as discussed in section IV. A lock must be released by a processor before the next acquire on it.

The *A-Stack* is used in each node of the network to constrain the data (read, write) operations issued by a processor to the same location in the shared memory. The *A-Stack* also keeps track of the addresses of the outstanding read operations (in addition to the write operations) which were not tracked by the *WA-Stack* under the TSO and PSO models.

Figure 7. Realization scheme of RC model

VI. EXPERIMENTS AND RESULTS

A. Hardware implementation cost

The designs are synthesized by using Synopsis Design Compiler with SMIC 90nm technology and optimized for the area. The synthesis results in term of nand-gate equivalent and maximum frequency are given in Table I. The difference in the area costs of the designs is mainly in the transaction controller (TCTRL), which uses *transaction counter* and *address stack* to implement the memory models. The TCTRL consumes 26.85%, 27.3%, 26.97% and 27.04% of the total area under the SC, TSO, PSO and RC models, respectively. The area cost for the TSO, PSO and RC models are increased by 1.8% (463 gates), 0.57% (153 gates) and 0.93% (243 gates) over the SC model in the TCTRL. The switch, synchronization handler (SH) and network interface (NI) together consume 73.14%, 72.78%, 73.03% and 72.95% of the total area under the SC, TSO, PSO and RC models, respectively. In all cases, the maximum clock frequency is 500 MHz or above.

TABLE I. SYNTHESIS RESULTS WITH 90 NM SMIC TECHNOLOGY

	SC Model		TSO Model		PSO Model		RC Model	
	A	F	A	F	A	F	A	F
Switch	13.24	0.5	13.24	0.5	13.24	0.5	13.24	0.5
SH	3.76	1.25	3.76	1.25	3.76	1.25	3.76	1.25
NI	49.99	1.25	49.99	1.25	49.99	1.25	49.99	1.25
TCTRL	24.6	0.5	25.05	0.5	24.74	0.5	24.83	0.5
Total	91.59		92.04		91.73		91.83	

A: AREA (KILO NAND GATES), F: FREQUENCY (GHZ)

B. Experimental Setup

We experimented on a configurable multicore NoC based cycle true simulation platform constructed in VHDL (Figure 4). The LEON3 processor is used in each node. The platform

uses a DSM system and the size of the shared, local memory of each node is 16 MB. The SH in each node maintains 256 locks in the shared address space. The *TC* and *WTC* each are 32 bits. The *A-Stack* and *WA-Stack* can stack up to 64 addresses each with 24 bits. The sizes of the Stacks are kept small and they are utilized efficiently. The addresses are popped from the *Stack* continuously by the completion of operations in a pipelined style. The packet formation in the NI uses 7 fields (96 bits). The buffering capacity at the NI is 64 packets. The *Nostrum* NoC [4] uses 2D regular mesh topology and deflection routing policy. The caches are disabled from the LEON3 processors in the experiments, since they are neutral for our evaluation of the memory models. In the experiments, the effects on execution time of network size and different traffic patterns are investigated and speedup, overhead and efficiency under memory models are reported.

C. Experiments with Synthetic Workloads

The performance of the RC, PSO, TSO and SC models are evaluated with three different synthetic workloads (WL1-WL3) as shown in Figure 8. These workloads are manually mapped on the LEON3 processors in the network. WL1 contains data and synchronization operations and has both *write followed by read* and *read followed by read* cases. WL2 contains *write followed by write*, *write followed by read* and *read followed by read* sequences. WL3 has, in addition, *read followed by write* operations and uses two non-overlapped critical sections. For the lock and protected (critical section) data operations hotspot traffic pattern is generated.

Figure 8. Sequences of transactions generated a) WL1 b) WL2 c) WL3

D. Results and Discussion

The Synthetic workload Execution Times (SETs) are compared for the RC, PSO, TSO and SC models (Figure 9(a)). The SC model is used as the baseline model. As the system scales up, the SETs are quickly increased under all the consistency models. It is due to the increasing network traffic and increasing waiting time for acquiring the locks. The SET is decreased under the RC model by allowing further reordering and relaxation compared to the other memory models. The average SETs of the RC, PSO and TSO models in the 8x8 network are reduced by 35.8%, 22.7% and 16.5% over the SC model, respectively. For the three synthetic workloads the SETs under all memory models in the 8x8 network are shown in Figure 9(b). The SET reduction is more under the WL2

compared to the WL1/WL3, due to the issuance of more data operations before the release operation.

Figure 9. SETs for WL1, WL2, and WL3.

E. Application Workloads

1) Matrix Multiplication

A matrix multiplication application workload calculates the product of two integer matrices $X[64 \times 1]$ and $Y[1 \times 64]$, resulting in a $Z[64 \times 64]$. Three matrices are decomposed into sub-matrices and stored in the distributed shared memory.

Figure 10. Matrix Multipl: a) AETs b) Speedup c) Overhead d) Efficiency

The Application workload Execution Times (AETs) under all memory models are decreased due to the division of the computation cost in the larger networks (Figure 10(a)). The AETs under the RC, PSO and TSO models in the 8x8 network are reduced by 24.1%, 3.2% and 1.8% over the SC model, respectively. The *speedup* is defined as the ratio of the single core execution time (T_s) and the execution time of the multicore (T_m). The speedup (Figure 10(b)) grows faster under the RC model compared to the other memory models as the system size scales up. The AET reduction and speedup under the RC model is higher compared to the other memory models due to additional reordering and relaxation in the shared memory operations. We define the *overhead* of the multicore system as: $N_c * T_m - T_s$, where, N_c is the number of cores in the system. The overhead increases as the system size grows due to the increasing communication cost (Figure 10(c)). The *efficiency* of the multicore system is defined as the ratio $speedup/N_c$. The RC model maintains high efficiency throughout compared to the stricter memory models (Figure 10(d)) when the system size is increased.

2) Pattern Search

The application searches P[64] data patterns against the M[64] data elements, which are initialized in the distributed shared memory across the network. Each node operates on the data patterns in its local shared memory and the data elements from remote nodes, and writes the output results into its local shared memory. The output values N[64] are the number of times the data patterns that appear in the data elements.

Figure 11. Pattern Search: a) AETs b) Speedup c) Overhead d) Efficiency

The AET reduction and speedup (Figure 11(a) & (b)) are higher compared to the matrix multiplication application due to low computation to communication ratio. As a result, the RC model allows more outstanding data operations, which overlap and pipeline with each other. Similarly, the overhead and efficiency are also affected. The overhead (Figure 11(c)) in the 8x8 network is reduced by 84.6%, 4.1% and 0.9% under the RC, PSO and TSO models over the SC model, respectively. The efficiency (Figure 11(d)) on the 8x8 system for the RC model is high 0.88, while for the PSO, TSO and SC models are 0.58, 0.57 and 0.56, respectively.

3) Bit Count/Data Analysis

The application analyzes a data vector M[1024] and calculates the number of set bits in each integer data item. The M data elements are initialized, read, analyzed and the output values (number of 1s) are stored in the distributed shared memory. Five different traffic patterns are used. Each node operates on the data items in its (*Bit complement*, *local shared*, *random*, *tornado of degree one* and *transpose*) node and also writes the output results in to the same node.

Figure 12. Bit Count: a) AETs b) Speedup

The AETs reduction for the relaxed models over the SC model is more under the *bit complement* traffic pattern, while least is observed under the *local shared* traffic pattern (Figure 12(a)). The AETs are dependent on the *physical distance* between the initiator and target nodes in the network. The AETs are high under random traffic pattern, because, irregular traffic pattern increases the network congestion and delay. The AETs under the RC, PSO and TSO models on the average are reduced by 44.5%, 27.2% and 24.9% over the SC model in the 8x8 network. As expected, the AET reduction is more under the RC model. Likewise, the speedup is high under the local shared and least under random traffic pattern (Figure 12(b)).

VII. CONCLUSION

The hardware support (transaction counter, address stack) for the three memory models and their performance comparison are shown in the McNoC systems. The results show that under synthetic workloads, the average execution time for the RC, PSO and TSO models in the 8x8 network is reduced by 35.8%, 22.7% and 16.5% over the SC model, respectively. For the application workloads, the execution time in the 8x8 network under the RC, PSO and TSO models is decreased on average by 36.4%, 10.7% and 9% compared to the SC model. The area cost for the relaxed models is increased by less than 2% over the SC model at the processor interface.

REFERENCES

- [1] Axel Jantsch et al., "Memory architecture and management in an NoC platform," in: Axel Jantsch and D. Soudris, editors, Scalable Multicore Architectures: Design Methodologies and Tools. Springer, 2011.
- [2] S. V. Adve et al., "Shared Memory Consistency Models: A Tutorial," Digital Western Research Laboratory, report no. 95/7, USA, 1995.
- [3] L. Lamport, "How to Make a Multiprocessors Computer That Correctly Executes Multiprocessor Programs," IEEE Transaction on Computers, Vol. C-28. No. 9, pp. 690-691, September 1979.
- [4] A. Jantsch "The Nostrum NoC," in: <http://www.ict.kth.se/nostrum>.
- [5] David L. Weaver and Tom Germond, editors. "The SPARC ArchitectureManual," Prentice Hall, 1994, SPARC International, Ver. 9.
- [6] P. Sewell et al., "x86-TSO: A Rigorous and Usable Programmer's Model for x86 Multiprocessors," Communications of the ACM, 2010.
- [7] M. Dubois et al., "Memory access buffering in multiprocessors," in: Proc. of 13th Ann. Inter. Symp. on Comp. Arch. (ISCA'86), 1986.
- [8] K. Gharachorloo et al. "Memory consistency and event ordering in scalable shared-memory multiprocessors," Computer Architecture News, 18(2): 15-26, June 1990.
- [9] D. Lenoski et al., "The Stanford Dash Multiprocessor," Computer, 87(3), March 1992, pp. 418- 429.
- [10] M. Ferdman et al., "Cuckoo directory: A scalable directory for many-core systems," in: Proc. of 17th Int. Symp. on (HPCA), 2011.
- [11] Romanescu, B.; Lebeck, A.; Sorin, D.J.; , "Address Translation Aware Memory Consistency," IEEE Micro, , vol.31, no.1, pp.109-118, 2011.
- [12] F. Petrot et al. "On cache coherence and memory consistency issues in NoC based shared memory multiprocessor SoC architectures," in: Proc. of 9th Euromicro Conf. on Digital System Design, DSD, 2006.
- [13] J.W. van den Brand and M. Bekooij, Streaming consistency: a model for efficient MPSoC design, in: Proc. of Euromicro Conf. on DSD, 2007.
- [14] "AMBA AXI Protocol Specification," in: <http://infocenter.arm.com/>
- [15] OCP International Partnership. OCP Specification 2.2, 2007.
- [16] A. Naeem et al., "Realization and Performance Comparison of Sequential and Weak Memory Consistency Models in Network-on-Chip based Multi-core Systems," In Proc. of the ASP-DAC, Japan, 2011.
- [17] A. Naeem et al., "Realization and scalability of release and protected release consistency models in noc based systems," in: Proc. of Euromicro Conf. on Digital Systems Design, DSD, Oulu, Finland, 2011.
- [18] http://jorisvr.nl/leon3_insntiming.html