

Transaction-Based Online Debug for NoC-Based Multiprocessor SoCs

Mehdi Dehbashi*

*Institute of Computer Science, University of Bremen
28359 Bremen, Germany
Email: dehbashi@informatik.uni-bremen.de

Görschwin Fey*†

†Institute of Space Systems, German Aerospace Center
28359 Bremen, Germany
Email: goerschwin.fey@dlr.de

Abstract—As complexity and size of Systems-on-Chip (SoC) grow, debugging becomes a bottleneck for designing IC products. In this paper, we present an approach for online debug of NoC-based multiprocessor SoCs. Our approach utilizes monitors and filters implemented in hardware. Monitors and filters observe and filter transactions at run-time. They are connected to a Debug Unit (DU). Transaction-based programmable Finite State Machines (FSMs) in the DU check assertions online to validate the correct relation of transactions at run-time. The experimental results show efficiency and performance of our approach.

Keywords—transaction-based online debug, system-on-chip (SoC), network-on-chip (NoC)

I. INTRODUCTION

Modern high-performance Systems-on-Chip (SoC) include many IP cores such as processors and memories. Network-on-Chips (NoC) have been proposed as a scalable interconnect solution to integrate large multiprocessor SoCs [1] [2]. Having a large SoC with complex communication among its cores, achieving complete verification coverage at pre-silicon stage is almost impossible. Therefore in addition to electrical bugs, some design bugs may also appear in the final prototype of an SoC.

The idea of transaction-based communication-centric debug is introduced in [3] to debug complex SoCs which interact through concurrent interconnects such as NoC. The transactions are observed using monitors [4] and the debug control unit can control the execution of the SoC (stopping, single stepping, etc). In [5], transactions are stored at run-time in a trace buffer using on-chip circuits. After an SoC run, the content of the trace buffer is read and analyzed offline with software. The analysis software tries to find certain patterns [6] in the extracted transactions that are defined by their *Transaction Debug Pattern Specification Language* (TDPSL). Because of limited size of a trace buffer, getting an execution trace of the transactions related to the time of bug activation is a challenging problem. To overcome this problem, the content of the trace buffer is utilized to backtrace the transactions along their execution paths [7]. The backtracing is performed in transaction-level states using *Bounded Model Checking* (BMC). However, backtracing needs formal pre-image computations which can blow up for large and complex designs [8]. To address this problem, we need to have online detection to stop the SoC close to the time of bug activation at the transaction level.

In this paper, we present a transaction-based debug infrastructure which can be used not only for online debug and online system recovery but also for interactive debug in which an external debug platform programs the FSMs and the filters according to the considered assertions at each round of debugging. Our hardware infrastructure contains monitors, filters, and a debug network including *Debug Units* (DU). Filters and DUs are programmed according to the transaction-based assertions defined by TDPSL. Transactions are monitored only at master interconnects. Slaves send information

to masters. This redundant information is used to observe the elements of transactions online. No modification of the internal components of the NoC is required. At run-time the programmable FSMs in the DUs investigate the assertions online and detect an error. Upon detection of an error, the DU recovers the SoC by informing the masters which have participated in the observed error. Then, the corresponding masters start the recovery process at run-time. Also we identify the requirements which a debug infrastructure has to fulfill in order to perform transaction-based online debug.

The main contributions of this paper are as follows:

- Introducing a debugging infrastructure to transaction-based online debug of NoC-based SoCs without modifying the internal components of the corresponding NoC (non-intrusive to the NoC).
- Analyzing and finding transaction-based debug patterns at-speed using debug units including programmable filters and FSMs.
- Presenting an ordering mechanism in the routers of the debug network to order the transactions online.
- Online system recovery without stopping and interrupting the NoC.

The experimental results show the efficiency of our approach using different assertion patterns defined by TDPSL such as race, deadlock, and livelock. An NoC-based SoC using a mesh network is setup in the Nirgam NoC simulator [9] to evaluate our approach. Also we show the effectiveness of the proposed online recovery in the experimental results.

The remainder of this paper is organized as follows. Section II introduces preliminary information on transactions and TDPSL. Our debug method including hardware and software parts is explained in Section III. The debug patterns and their corresponding FSMs are explained in Section IV. This section also presents experimental results on an NoC-based SoC. The last section concludes the work.

II. PRELIMINARIES

A. Transaction

In this section we shortly explain the transaction elements from [10] and [5]. Each transaction includes a request and a response. Masters request and slaves respond. Each transaction has four basic elements: *Start of Request* (SoRq), *End of Request* (EoRq), *Start of Response* (SoRp), and *End of Response* (EoRp). In TLM, SoRq corresponds to putting the request in the channel by the master. EoRq is getting the request by the slave. SoRp corresponds to putting the response in the channel by the slave. EoRp is getting the response from the channel by the master. Also there are two additional elements which are called: *Request Error* (ErrRq) and *Response Error* (ErrRp). These elements handle error conditions and correspond to any kind of error that causes a request or a response to fail.

B. Transaction Debug Pattern Specification Language (TDPSL)

TDPSL has three layers: Boolean layer, temporal layer, and verification layer [5]. The Boolean layer includes *trans_exp*

This work has been supported in part by the University of Bremen's Graduate School SyDe, funded by the German Excellence Initiative and in part by the German Research Foundation (DFG, grant no. FE 797/6-1).

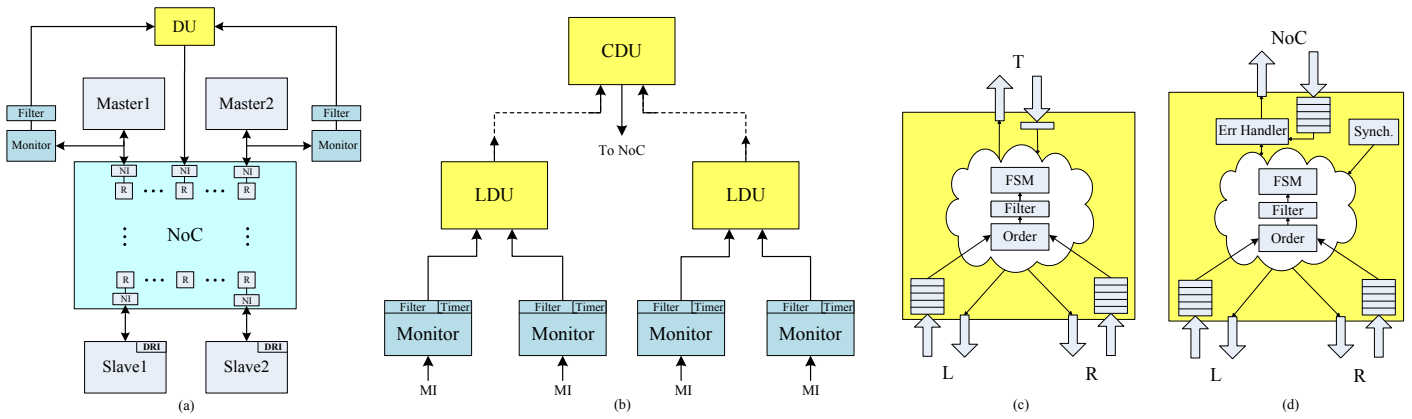


Fig. 1. (a) Debug Infrastructure, (b) Tree-Based Debug Infrastructure, (c) LDU Structure, (d) CDU Structure

which represents the basic elements of transactions. The *trans_exp* format is as follows:

$trans_type(master, slave, type, address, tag)$

Field *trans_type* can be any transaction element mentioned in Section II-A as well as the *Start of Transaction* (SoTr) and the *End of Transaction* (EoTr) which are similar to SoRq and EoRp respectively. Fields *master* and *slave* specify the ID of master and slave. Field *type* can be *Rd* or *Wr*. Field *address* indicates the slave address symbolically as SAME, SEQ, and OTHER. Field *tag* indicates the transaction number and is only used for buses that allow non-blocking requests and out-of-order responses [5]. In our paper, we show a transaction without considering the field *tag*.

The motivation to use symbols for the address field is to abstract and to compress the address bits. In this case, only the compact address information is stored or sent via network for debugging. The symbols can be defined with respect to the application and the granularity of debugging. SAME specifies that in the current transaction, slave address is same as the address in the previous transaction for this slave. SEQ specifies that in the current transaction, slave address has one word difference with the previous address for this slave. OTHER specifies that in the current transaction, the slave address in neither SAME nor SEQ.

For example transaction $EoTr(m1, s2, Rd, -)$ represents the end of a read transaction from master $m1$ to slave $s2$ with any address. The symbol “-” indicates that we leave the corresponding field as don’t care.

The properties in terms of transaction sequences are defined at the temporal layer. Different operators are utilized at this layer such as concatenation operator (;), fusion operator (:), or operator (!), and operator (&), and repetition operators [5]. In the verification layer, the *assert* statement is defined. Also a *filter* can be defined which specifies a filter over the execution path for the evaluation of the assertion statement.

Following is an example of a simple assertion in TDPDSL:

assert never

$EoTr(m2, s1, Wr, -); SoTr(m1, s1, Rd, -)$

This assertion specifies that start of a read transaction from master $m1$ to slave $s1$ must never be directly after the end of a write transaction from master $m2$ to slave $s1$.

III. DEBUG METHOD

Transaction level online debug aims at improving the observability and the controllability of the system. Whenever transactions conform to certain debug patterns, an error is detected. In this case, the *Debug Unit* (DU) sends the debug packets to the SoC nodes in order to control the network and to recover from the error state.

We have some requirements to enable transaction-based online debugging:

- 1) Our debug infrastructure has to be able to collect the elements of each transaction at run-time.
- 2) We have to be able to order the transactions online.
- 3) We need to assert debug patterns, i.e., the relation of transactions, at run-time.

If a debug infrastructure fulfills the three mentioned requirements, it can be used for transaction-based online debug.

In the following, we explain our debug infrastructure fulfilling the above mentioned requirements. To collect all elements of each transaction in a system based on NoC (first requirement), we need distributed monitors and *Debug Redundant Information* (DRI). Monitors and DRI are explained and discussed in Section III-A and Section III-B, respectively.

The transaction ordering mechanism in the *Debug Units* (DU) is responsible to order transactions (Section III-D) fulfilling the second requirement. A DU is the main part of the debug infrastructure which searches for certain debug patterns in the received transactions. We use a tree-based debug network structure in which all monitors have a short distance to DUs. In the debug network, the transactions are ordered using DUs. The ordered transactions are transferred on each link of the debug network from bottom to top such that the ordered transactions can be utilized in each level of the debug network for hierarchical and assertion-based debug.

FSMs in DUs are utilized to investigate transaction-based assertions at run-time to fulfill the third requirement (Section III-E). The filters in DUs and in monitors help FSMs by dropping unrelated transactions (Section III-E).

Figure 1(a) shows the hardware infrastructure of our approach for an SoC including four IPs, 2 masters and 2 slaves. The debug infrastructure has the following parts: monitors, filters, DU, and DRI. The internal structure of a DU has also three main parts: transaction ordering, filter, and FSM. Each IP in Figure 1(a) is connected via a *Network Interface* (NI) to the NoC.

Figure 1(b) shows the tree-based debug infrastructure. The lowest level of the infrastructure includes monitors and filters. Monitors are connected to *Master Interconnects* (MI) to observe the transactions. A *Central Debug Unit* (CDU) is only at the top level. The other levels have *Local Debug Units* (LDU). LDU and CDU structures are explained in Section III-C.

A. Monitor

Monitors extract the basic elements of a transaction as mentioned in Section II-B. They observe master interconnects to enable transaction-based debug [3]. In a packet-based protocol in an NoC, we can immediately extract the elements *master*, *slave*, and *type* by observing the master interconnects. But to extract the element *address* as SAME, SEQ, and OTHER, which is a comparison of the slave address in the current




Fig. 2. CDU procedure to recover the SoC

transaction and the previous transaction for the corresponding slave, we need some DRI. The next section explains the DRI.

A monitor in our infrastructure observes a master interconnect and signals a matching transaction expression explained in Section II-B as an output.

Each monitor includes also a timer. The timer is used to attach a timestamp to each observed transaction. The timestamp attached to a packet is utilized at DUs to order the transactions arriving from the left and right input links of DUs. As the transactions are consumed online using FSMs, large timestamps are not required. Timestamps only need to distinguish the order of transactions arriving at DUs.

In an SoC with asynchronous IPs, the CDU sends synchronization packets to monitors. The timers in monitors are synchronized according to the synchronization packets. As only the CDU sends the packets from the top level to the bottom level in the tree-based debug network, the delay of synchronization packets arriving at monitors are predetermined. In this case, the time in monitors is synchronized with the time in the CDU by incrementing the CDU time included in the synchronization packet with the delay of synchronization packet.

Another approach to synchronize the transactions is using the *relative timestamps*. In this approach, the time of each transaction is calculated in comparison to the time of the front transaction in the debug network. Then, this relative timestamp is attached to the corresponding transaction. To use this approach, some timers are required in LDUs and CDU.

B. Debug Redundant Information (DRI)

DRI is used to extract and to transfer the element *address* of a transaction. We can form the element *address* using slaves (slave-based approach) or using debug units (DU-based approach). In the following we discuss these two approaches.

In a slave-based approach, the element *address* is formed in the slaves and is sent as redundant information to masters through the NoC. Because the element *address* is a comparison of the address of the current transaction with the address of the previous transaction for the corresponding slave, this comparison can be simply done in each slave.

The slave should send two bits redundant information to masters. These two bits specify the symbols SAME, SEQ, and OTHER. We can also use more symbols for the slave address to have more accurate data depending on the applications running on the SoC.

The DRI section in each slave in Figure 1(a) compares the slave address of the current transaction with the slave address of the previous transaction in the corresponding slave. Then the DRI section selects a symbol (SAME, SEQ, or OTHER) and adds this symbol in the response packet as two redundant bits. On the master side these two bits are read by monitors to constitute a complete transaction expression. These two redundant bits are used only in monitors. The master applications should ignore these two redundant bits.

In this case, we can have the address information for the corresponding slave only in the *EoTr*. The element *address*




Fig. 3. An example for master recovery thread in the case of a software deadlock

is not available in *SoTr*. To have this information, we should wait to receive an *EoTr* by the corresponding slave.

The second approach to form the element *address* uses debug units, i.e., DU-based approach. In this approach, the slave addresses are observed by monitors and sent to the DUs. In the DU, there is one address register for each slave. The address registers keep the address of the previous transaction for each slave independently. When a new transaction is performed, the content of the address register related to the corresponding slave is compared to the new transaction address. Then the symbols SAME, SEQ, and OTHER are derived and the address register is updated to keep the slave address in the latest transaction for the corresponding slave.

In the DU-based approach, the element address is available for both *SoTr* and *EoTr*. The DU-based approach needs more memory in the debug units storing slave addresses. Also it needs more bandwidth for the debug network to transfer slave addresses to DUs. The advantage of this approach is being non-intrusive to the SoC.

C. Debug Unit (DU)

A debug unit can be an LDU or a CDU. A CDU is used at the top level in the tree-based debug network (Figure 1(b)). An LDU is used in other levels of the debug network. The structure of an LDU is suitable to build a tree-based network. An LDU has three ports (Figure 1(c)): top port T, left port L, and right port R. The right and left ports transfer the data observed by monitors to the top level. Also the synchronization packets sent by the CDU are transferred from the top port to the left and right ports reaching timers. The CDU controls the traffic of the packets sent from the top level to the leaves in the debug network. In this case, we use only one buffer in each LDU to transfer synchronization packets.

The packets arriving at the inputs of the right and left ports are stored in the right and left FIFOs. Then the transaction ordering selects a transaction packet such that the transactions are ordered based on their timestamps. The filter does not allow a transaction to be forwarded if the transaction is not related to the considered assertions. The related transactions regarding the considered assertions are used in the FSM to investigate the assertions. If an assertion fails, an error message is sent to the CDU.

The CDU is used at the top level of the debug network. The CDU has two additional tasks: synchronizing the timers and handling error cases (Figure 1(d)). Synchronization is performed by part *Synch* in Figure 1(d). When there is an error, the error handler in the CDU manages the network by sending some debug packets to other nodes in the SoC. The CDU is connected to the NoC communicating with other nodes in the network. In this case, the CDU can send the error state to all nodes or some special nodes in the network in order to collect more accurate debug information or to recover the SoC from the error state.

Figure 2 shows the CDU procedure to recover the SoC from an error. At the step of pattern detection, the CDU checks the debug patterns at run-time. If the CDU detects an error, the

second step is started. In this step, the CDU sends a recovery packet to the masters which have contributed to the observed error. A recovery packet contains an error type and additional information helping the masters to start a recovery process. In the third step, the CDU restarts its own FSM. Then, the CDU sends a restart packet to the LDUs restarting the LDU-FSMs. Afterwards, the procedure continues with the step of pattern detection.

Figure 3 shows an example for a master recovery process in the case of a software deadlock. In this case, when a master receives a recovery packet from the CDU with the error type deadlock, the master releases the locked resources. Then the master waits for a random time and proceeds its main function again. With this procedure, the system is recovered online from the error state without stopping and interrupting the NoC.

DUs include FSMs to investigate transaction-based assertions at run-time. To check an assertion in an efficient way, we need to program both LDU-FSMs and CDU-FSM. Distributed online assertion checking can be performed through programming the FSMs in different levels.

D. Transaction Ordering

When there is more transaction traffic on one link than another link in the debug network, some early-generated transactions are accumulated and buffered in the FIFO of the corresponding DU. This case may also occur when the bandwidth of the debug network is less than the bandwidth of the NoC. When there are some transactions in the FIFO of the left link which have been generated earlier than the transaction available in the FIFO of the right link, the transaction of the right link has to wait until the transactions with smaller timestamps on the other link have been transferred. By comparing the timestamp of a packet in the left FIFO and the right FIFO, the packets are ordered based on their generation time.

The length of the timestamp depends on the worst case delay of the debug network. A timestamp should only be able to distinguish the packets based on the time in which they have been generated or sampled. In Figure 1(c) and Figure 1(d), the part *Order* in DU compares the timestamps of a packet in the right and the left FIFOs and selects a packet which has a smaller timestamp.

The size of the left and right FIFOs may influence the accuracy of the debug pattern detection because if the FIFO becomes full, some transactions are lost. In this work, we assume that the size of the FIFOs is sufficient to process the transactions.

E. Filter and Debug FSM

A filter is located in monitors and DUs. A filter is used to filter unrelated transactions in a trace. In this way, the debug unit receives only the related transactions for the assertion statements. Filtering can be done over all parameters of a transaction expression, i.e., *trans_type*, *master*, *slave*, *type*, and *address*. The filter is programmable according to the main assertion statements.

Debug FSMs are programmable FSMs which are utilized in debug units to investigate the assertions online. Debug FSMs include local FSMs and global FSMs verifying local assertions and global assertions. Debug units can be programmed to implement distributed FSMs validating online assertions in different levels. To do this, first the transaction-based assertions should be analyzed based on their locality in the corresponding SoC. We need to know which task is running on which IP. Accordingly, the filters should be programmed and the assertions should be distributed among debug units (LDUs and CDU).

IV. IMPLEMENTATION

For the experiments we setup a 3x3 mesh network in the Nirgam NoC simulator [9]. Nirgam is a cycle-accurate simulator which is implemented in SystemC language. We have simulated the system for one million cycles. During the run-time of the SoC, our debug infrastructure asserts the debug FSMs which are mentioned in the next sections. We have implemented dining philosophers [5] and a random application [7] as example applications. In the random application, each master waits for a random time. Then the master selects a random list of slaves as resources. If the master can lock all the required resources, the processing is started. Afterwards, the resources are released and the procedure is repeated. If the master cannot lock all the required resources, the master waits for a random time and tries again [7].

In our experimental setup, the SoC has four masters (philosophers) and four slaves (chopsticks) which are divided into two groups communicating in parallel. Each group has two masters and two slaves (first group: $m1, s1, m2, s2$. second group: $m3, s3, m4, s4$). Four monitors are used to observe the master interconnects. Also two LDUs and one CDU constitute the debug network. In the following sections, we discuss debug patterns for race, deadlock, and livelock as an example.

A. Debug Pattern for Race

A race may occur when one write transaction to the same place occurs during the previous write. In TDPSL this case is written as follows:

```
assert never{
  SoTr(m1, s1, Wr, -); SoTr(m2, s1, Wr, SAME);
  EoTr(m1, s1, Wr, SAME)
}filter(*, *, *)
```

Filtering is done on the three first parameters of transaction expressions. Sign * in the filter means only the related transaction types, masters, and slaves should be considered. Therefore the transactions related to slave $s2$ are omitted. Also all transactions related to the second group, i.e. group of master $m3$ and master $m4$, have to be omitted. In our infrastructure, the filters are programmed such that the transactions related to slave $s2$, master $m3$ and master $m4$ are filtered online.

As explained in Section III-B, the DRI can be transferred using the slave-based approach or the DU-based approach. As in the race assertion we need the element *address* in *SoTr*, therefore we can only use the DU-based approach to investigate this assertion. In the slave-based approach, we can have the element *address* only for *EoTr*. If we use the slave-based approach, we need to change the race assertion such that only *EoTr*s include the element *address*. But this case causes some latency in the detection of the assertion fail.

To increase the verification coverage of the FSM, we need to have a more comprehensive pattern. In the following, we write an improved race pattern to cover more race conditions happening on slave $s1$:

```
assert never{
  SoTr(m1, s1, Wr, -); SoTr(m2, s1, Wr, SAME);
  EoTr(m1, s1, Wr, SAME)
|SoTr(m2, s1, Wr, -); SoTr(m1, s1, Wr, SAME);
  EoTr(m2, s1, Wr, SAME)
}filter(*, *, *)
```

In the first part of the assertion, the pattern checks a race condition in which master $m1$ starts a race. The second part of the assertion specifies a race condition in which master $m2$ starts a race.

TABLE I
NUMBER OF DEBUG PATTERNS DETECTED FOR EACH APPLICATION

	Rand. Application	Din. Application
Race Pattern	62	0
Deadlock Pattern	0	1
Livelock Pattern	0	0

B. Debug Pattern for Deadlock and Livelock

When some masters are waiting for other masters to release shared resources, deadlock happens. Here we show the case of two masters and two slaves as an example. Each slave has a semaphore which specifies its access permission. When semaphore is 0, the slave is free. When semaphore is 1, the slave is locked. Each master should first lock required slaves, then it can start its process using the corresponding slaves as resources. To lock a slave, a master has to first read the semaphore of the corresponding slave. If the semaphore is 0, then the master can write 1 to the semaphore to lock the corresponding slave. Therefore to lock a slave, a master needs two transactions, i.e., one read transaction and one write transaction. If the semaphore is 1, i.e., the slave is already locked, then the master should wait until the corresponding slave becomes released (in the application of dining philosophers). Both masters have access to the semaphore of each slave. Accessing a semaphore is equivalent to accessing the same address by different masters.

A simple deadlock scenario for two masters and two slaves is as follows [5]: 1) Master1 locks the first semaphore. 2) Master2 locks the second semaphore. 3) Master1 waits for the second semaphore. 4) Master2 waits for the first semaphore. 5) Steps 3 and 4 are repeated. This deadlock condition is written in TDPSL as follows [5]:

```
assert never{
  EoTr(m1, s1, Rd, -); EoTr(m1, s1, Wr, SAME);
  EoTr(m2, s2, Rd, -); EoTr(m2, s2, Wr, SAME);
  {EoTr(m1, s2, Rd, SAME); EoTr(m2, s1, Rd, SAME)
  |EoTr(m2, s1, Rd, SAME); EoTr(m1, s2, Rd, SAME)
  }[+]
  }filter(*, *, *)
```

This assertion is written for applications in which each master first locks the slave with the same ID. For example master $m1$ first locks slave $s1$. If it is successful, then it locks slave $s2$. To implement this assertion by our debug infrastructure, the filters are programmed such that transactions $SoTr$ are filtered online as unrelated transactions. Also all transactions related to the second group, i.e. group of master $m3$ and master $m4$, are filtered.

In the mentioned deadlock assertion, first the lock process from master $m1$ is checked, then the lock process from master $m2$. To illustrate this case better, we denote a read transaction (write transaction) from master m_x to slave s_x as Rxy (Wxy). In the previous deadlock assertion only the sequence $(R11, W11, R22, W22)$ is checked for the lock process. To increase the verification coverage of the deadlock assertion we check the following sequences for the lock process: $(R11, W11, R22, W22)$, $(R11, R22, W11, W22)$, $(R11, R22, W22, W11)$, $(R22, W22, R11, W11)$, $(R22, R11, W22, W11)$, $(R22, R11, W11, W22)$

A livelock is similar to a deadlock where two or more processes proceed accessing shared resources which are already locked. But in the case of a livelock, they release the locked resources permitting the other processes to continue. A simple livelock scenario for two masters and two slaves is as follows [5]: 1) Master1 locks the first semaphore. 2) Master2 locks the second semaphore. 3) Master1 waits for the second semaphore. 4) Master2 waits for the first semaphore. 5) Master1 unlocks the first semaphore. 6) Master2 unlocks the second semaphore.

TABLE II
EFFECT OF ONLINE RECOVERY

	Without Recovery	With Recovery
#Eating	6	3276
#Resolved Deadlock	0	77

7) Steps 1 to 6 are repeated. To implement the livelock FSM, the deadlock pattern is enhanced to check the steps 5 and 6.

Our debug infrastructure is programmed according to race, deadlock, and livelock debug patterns and detects the occurrence of each debug pattern at run-time. Table I shows the number of times a debug pattern is detected during the simulation time of one million cycles. In the random application, the race pattern is detected 62 times. Also at run-time, this information is sent to the corresponding masters (Figure 2). In the application of dining philosophers, the deadlock pattern is detected one time. In this case, after the first deadlock detection, the group of deadlocked masters cannot proceed with their process anymore.

As explained in Section III-C (Figure 3), the masters can start a recovery process after the CDU has sent them the error state. Table II indicates the effect of using the recovery process in the masters. Without recovery process in the application of dining philosophers, the masters in one group can eat only 6 times. After that a deadlock happens and the masters cannot pick up their required chopsticks. However, the recovery process of Figure 3 causes that the masters can continue their main process even if they get into a deadlock. In this case, a deadlock happens 77 times. But in each time, the CDU detects the deadlock at run-time and triggers the recovery process in the masters to recover the deadlocked network. Consequently, the masters can eat more often (3276 times) as shown in Table II.

V. CONCLUSION

We introduced an approach to online debug for NoC-based multiprocessor SoCs. Our approach contains a hardware infrastructure, debug redundant information, and FSMs. Monitors, filters, and debug units are considered in our debug hardware infrastructure. This infrastructure allows us to investigate and to debug the behavior of an NoC-based SoC at run-time. Filters and FSMs are programmed according to the transaction-based assertions defined by TDPSL. In the experimental results, we investigated the efficiency of our approach for the debug patterns race, deadlock, and livelock.

REFERENCES

- [1] L. Benini and G. D. Micheli, "Networks on chips: A new SoC paradigm," *IEEE Computer*, vol. 35, no. 1, pp. 70–78, 2002.
- [2] P. P. Pande, C. Grecu, A. Ivanov, R. A. Saleh, and G. D. Micheli, "Design, synthesis, and test of networks on chips," *IEEE Design & Test of Computers*, vol. 22, no. 5, pp. 404–413, 2005.
- [3] K. Goossens, B. Vermeulen, R. van Steeden, and M. T. Bennebroek, "Transaction-based communication-centric debug," in *International Symposium on Networks-on-Chips (NOCS)*, 2007, pp. 95–106.
- [4] B. Vermeulen and K. Goossens, "A network-on-chip monitoring infrastructure for communication-centric debug of embedded multi-processor SoCs," in *International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*, 2009, pp. 183–186.
- [5] A. M. Gharehbaghi and M. Fujita, "Transaction-based debugging of system-on-chips with patterns," in *Int'l Conf. on Comp. Design*, 2009, pp. 186–192.
- [6] W. Ecker, V. Esen, M. Hull, T. Steininger, and M. Velten, "Requirements and concepts for transaction level assertions," in *Int'l Conf. on Comp. Design*, 2006, pp. 286–293.
- [7] A. M. Gharehbaghi and M. Fujita, "Transaction-based post-silicon debug of many-core system-on-chips," in *Int'l Symp. on Quality Electronic Design*, 2012, pp. 702–708.
- [8] F. M. de Paula, A. Nahir, Z. Nevo, A. Orni, and A. J. Hu, "TAB-BackSpace: Unlimited-length trace buffers with zero additional on-chip overhead," in *Design Automation Conf.*, 2011, pp. 411–416.
- [9] NIRGAM NoC simulator, 2013, available online: <http://nirgam.ecs.soton.ac.uk/>.
- [10] *OSCI TLM-2.0 Language Reference Manual*, 2013, available online: <http://www.systemc.org>.