

A High-Level Modeling and Simulation Approach Using Test-Driven Cellular Automata for Fast Performance Analysis of RTL NoC Designs

Moon Gi Seok, Hessam S. Sarjoughian
 ACIMS lab., Arizona State University, Tempe, USA
 mseok@asu.edu, Hessam.Sarjoughian@asu.edu

Daejin Park
 Sch. of EE., Kyungpook National University, Rep. of Korea
 boltanut@knu.ac.kr

Abstract—The simulation speedup of designed RTL NoC regarding the packet transmission is essential to analyze the performance or to optimize NoC parameters for various combinations of intellectual-property (IP) blocks, which requires repeated computations for parameter-space exploration. In this paper, we propose a high-level modeling and simulation (M&S) approach using a revised cellular automata (CA) concept to speed up simulation of dynamic flit movements and queue occupancy within target RTL NoC. The CA abstracts the detailed RTL operations with the view of deciding a cell’s state of actions (related to moving packet flits and changing the connection between CA cells) using its own high-level states and those of neighbors, and executing relevant operations to the decided action states. During the performing the operations including connection requests and acceptances, architecture-independent and user-developed routing and arbitration functions are utilized. The decision regarding the action states follows a rule set, which is generated by the proposed test environment. The proposed method was applied to an open-source Verilog NoC, which achieves simulation speedup by approximately 8 to 31 times for a given parameter set.

I. INTRODUCTION

The network-on-chip (NoC) paradigm has been an important design methodology using on-chip communications among large numbers of intellectual properties (IPs) to mitigate the design complexity and scalability issue of a system on a chip (SoC) due to its advantages including communication throughput, flexibility, and scalability. Designers have commonly implemented various NoCs using hardware description languages (HDLs) at register-transfer level (RTL) for automatic synthesis to circuits.

Generally, RTL NoC design requires iterative simulations to find acceptable communication performance using a set of parameters for the participation of potential IPs or to optimize the design parameters (e.g., input buffer size) for designated IPs. However, slow RTL simulation speed causes reducing the parameter set space or limits having a sufficient parameter exploration for optimization. To increase the simulation speed, the high-level modeling can be a solution. Typically, the high-level NoC modeling approaches can be divided into two main categories: queuing modeling and architectural- and algorithm-level hardware modeling using high-level modeling languages, such as SystemC.

The queueing modeling studies have tried to develop various parametric equations to observe the average service time and queue utilization based on the queueing theory [1], [2]. The previous studies show relative sound accuracies ($\approx 90\text{-}92\%$)

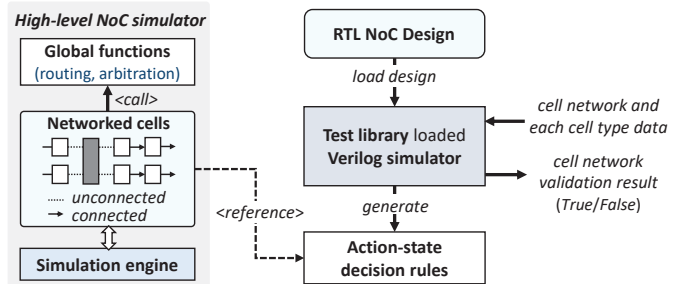


Fig. 1: Proposed CA modeling process using RTL NoC.

and negligible computation overhead due to their analytic characteristics comparing with the RTL simulation, which iteratively processes events for signal value transitions at every simulation step. However, the previous studies have limitations in evaluating dynamic changes in packet transmission, such as buffer status, so it is not able to monitor hotspot situations at a specific time. Moreover, the statistical approach of queuing models cannot guarantee the exact communication performance results compared to the RTL simulation.

The architectural- and algorithm-level NoC models are usually developed using SystemC at the early stage of NoC designs. The high-level models have loosely-/approximated-timed or cycle-level accuracy, and utilized to explore design spaces and verify functionalities for assisting implementation of RTL models as intermediate models [3], [4]. Due to the algorithm-level abstraction of RTL operations, the models increase simulation speed with reduced modeling efforts. However, some operations of final RTL models and their high-level intermediate models might differ, so it cannot guarantee matched packet transmission results between two models. Moreover, the high-level modeling requires in-depth knowledge about detailed and specific operations of target RTL design.

In this paper, we propose a high-level and cycle-accurate modeling and simulation (M&S) method for the fast evaluation of packet transmissions of target RTL NoC design using an extended concept of cellular-automata (CA) and the proposed process, as shown in Fig. 1. The proposed methods abstracts detailed RTL operations into 1) cell’s action-state decision based on high-level state (that are flit and capacity information) and 2) executes operations (that are flit fetching and connection change) corresponded to decided action state

using a built-in flit-fetching function and user-developed global functions. Compared to the previous algorithm-level modeling, the CA model is reversely derived from final RTL design, not an intermediates model toward the RTL design. Without in-depth knowledge of all RTL operations, functions of action-state decision (related to the action-execution timing) reference the derived rules by the proposed test environment using the user-specified information of CA network and RTL probing signals of cells' high-level states, and the implementation of global functions requires the knowledge about the routing and arbitration algorithm of target RTL design.

The CA network structure is identified by the designer using the information of RTL arbitration units, FIFO queues, or other units that cause propagation delay due to the flit-storing flipflops (FFs). Similar to the conventional CA concept, each cell determines the action state based on neighbors' and own state at each time step according to a rule [5]. The central CA extensions consist of two parts: 1) neighbors are dynamically determined in runtime to reflect the NoC arbitration and 2) each cell has an influence function, which enables the simulation engine to find and schedule only state-changeable neighbors after a cell's action state is changed in order to prevent unnecessary executions of CA to achieve scalability.

A decision rule consist of multiple coupled states of the target cell and its neighbors and associated action states. For the decision-rule generation, the proposed test environment observes RTL signals that are related to cells' high-level state and action state and stores the pair of high-level and action states during the RTL NoC simulation of randomly injected packets. Moreover, the CA network is validated by checking next position of flits to be moved during the test. After the test, the environment generates compact decision rules through checking a state dependency.

The rest of the paper is organized as follows. Section II describes the proposed CA modeling and simulation method. Section III describes the proposed test environment and rule generation. Section IV applies the proposed approach to an open source RTL NoC model. Section V concludes the paper.

II. CELLULAR AUTOMATA MODELING AND SIMULATION OF NOC RTL DESIGN

A. CA Modeling Perspectives and Formal Specifications

We propose high-level CA modeling for RTL NoC designs, which have the following characteristics.

- NoC designs consist of synchronous routers.
- NoC designs have a flit-based transmission mechanism, which breaks a large packet into small flits for low latency.

All CA can be classified into two types: a buffer cell (BC) and a coupling cell (CC). Each BC has a FIFO buffer to store one or multiple flits. The BC represents queues or intermediate units where each FFs holds a flit and causes cycle-level delay, as shown in Fig. 2(a). The CC is responsible for establishing the connection or disconnection between two specific BC s among two or multiple adjacent BC s. The cell

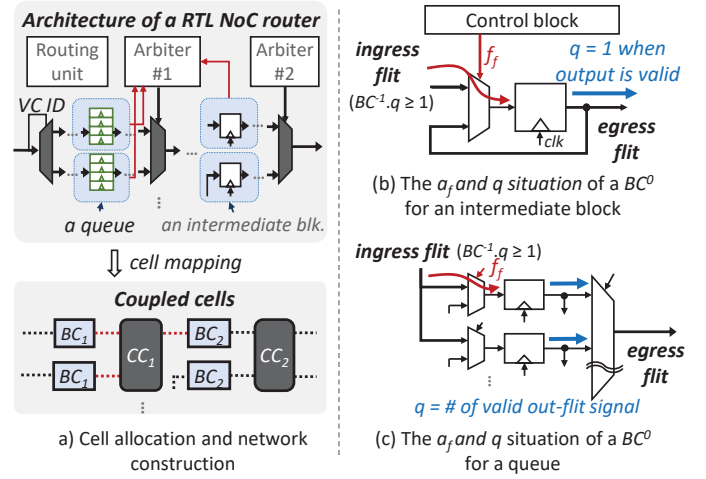


Fig. 2: Cell allocation, and f_f and q value of a BC (BC^0).

coupling between CC and BC can be established using input and output information of the representing arbiter unit of the target RTL NoC. The specific high-level states of BC and CC are defined as below.

Definition 1: A high-level state (S_b) of BC is represented using a 4-tuple $\langle id, F, Q, FT \rangle$ and a high-level state (S_c) of CC is represented as a 3-tuple $\langle id, C, P \rangle$, where

- id : the cell identification index;
- F : the set for the flit-fetching action states (f_s) of BC , $F = \{f_i, f_f\}$, f_i : idle, f_f : state of flit fetching at the next clock trigger;
- Q : the set of possible number of flits (q) in the FIFO buffer within BC ;
- FT : the set of the front flit types, $FT = \emptyset \cup \{header, payload, last, single\}$;
- C : the set of connectivity action states (c_s) of CC , $C = \{c_c, c_d, c_{c(0)}, c_{d(1)}\}$, c_c : connected state, c_d : disconnected state, $c_{c(0)}$: newly connected state at the current clock cycle, $c_{d(1)}$: to be disconnected state at the next clock trigger;
- P : the id pair of connected BC s, $\emptyset \cup \{BC.id, BC.id\}$.

The $f_f \in F$ means to bring a flit from previous BC (BC^{-1}) to current BC (BC^0) at the next clock trigger, and the $q \in Q$ is related to the number of valid output flits, as shown in Fig. 2(b)(c). The RTL signal of a cell's f_f is related to write-enable conditions of an input flit at the clock trigger. The action state of a cell, f_s and c_s , is decided by checking the current values of its own and neighbors' state based on the described mapping condition in its decision rule.

Depending on the existence the connection states of adjacent CC s, neighbors of a cell can be variously coupled, as Fig. 3(b). For a BC , if an adjacent CC is disconnected, the BC cannot reach another BC across the CC as a neighbor, and an adjacent CC can be a neighbor after requesting a connection. The connection request to next CC is executed by the simulation engine when a flit is fetched and a next CC is not designated. Depending on the number of next reachable

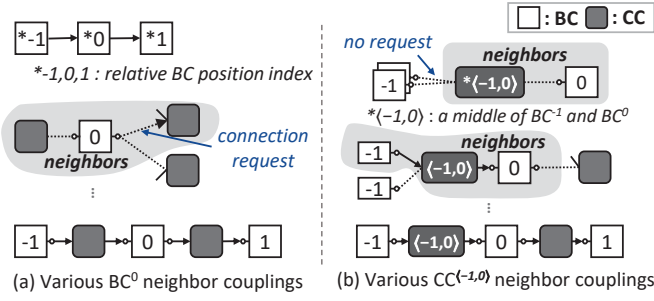


Fig. 3: Various neighbor couplings of BC and CC .

CC s, multiple connection requests can be performed based on the destination information of the front flit, which is obtained using the global routing function. The current BC (BC^0) and current CC ($CC^{(-1,0)}$, connecting BC^{-1} to BC^0 and BC^0) can reach BC^{-1} and BC^1 at most as neighbors.

The background philosophy referring three BC s at most is to find the relationship between status of sender (BC^{-1}) and availability of receiver (BC^0), and the receiver's availability sometimes depends on the action of BC^1 . For example, a BC^0 represents an intermediate unit that has a FF to store a flit. A new flit arrival is possible when the current FF's output flit is ready to exit at the next clock trigger to prevent the flit loss, the output-flit exit is related to the f_s of BC^1 . In general, two cells should be referred to decide the availability.

Based on the defined high-level states and neighbor couplings, we define the decision function of the action state and an influence function for BC and CC , as shown below.

Definition 2: The simulation engine invokes the BC functions δ_a and η as well as the CC functions δ_c and η , where

- $\delta_a : \overbrace{Q^{-1} \times Q^0}^{\text{prev. cur.}} \times \overbrace{F^1}^{\text{next}} \rightarrow F^0$, the f_s decision function;
- $\delta_c : \overbrace{(F \times Q \times FT)^{-1}}^{\text{prev. states}} \times \overbrace{C_s^{(-1,0)}}^{\text{cur.}} \times \overbrace{(Q^0 \times F^1)^n}^{\text{next states}}$, the c_s decision function, p/n : the number of previous/next neighbor BC s;
- $\eta : \overbrace{S_{b/c}^0}^{\text{cur.}} \times \overbrace{S_{b/c}^m}^{\text{prev./next}} \rightarrow \{id\}^r$, the influence function, m : the number of neighbors, r : the number of neighbors whose action state can be changeable due to the current state ($S_{b/c}^0$)

Each of the δ_a , δ_c , and η functions has an *if-then* structure based on corresponding rule-describing conditions of previous, current and next states, and the abstracted structures of those functions are shown in Fig. 4(a)(b)(c). In the decision rules of CC and BC , default conditions are defined for trivial action-change situations.

After simulation time has advanced, the $f_{s,i}$ of all BC s are set to f_i before making decisions. Depending on whether previous, current and next states meet any condition of f_f , the f_s can be newly updated to f_s . The c_s of CC s can be also updated to $c_{d(1)/c(0)}$ or remain unchanged depending on the

$A \otimes B$: returns true if cur. and neighbor's states of A cell match any cond. of B action
 $A \odot B.C$: returns true if cur. state of A cell matches any prev./next state cond. of C action of next/prev. B cell

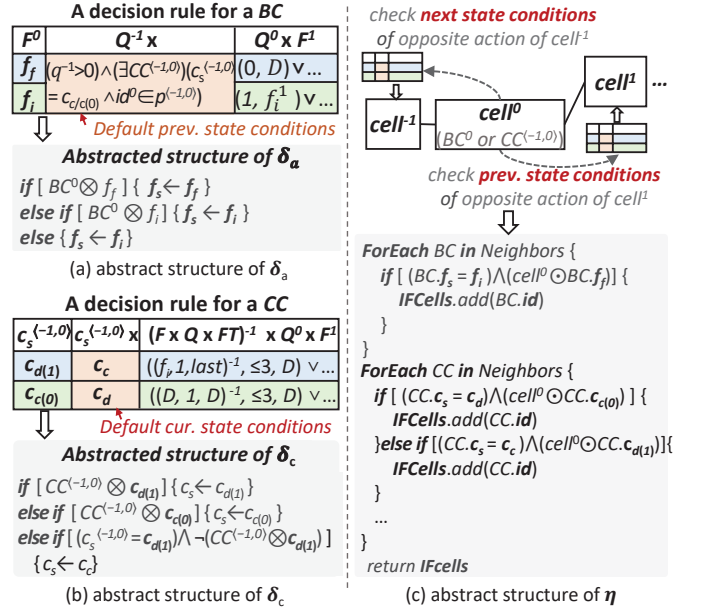


Fig. 4: Abstract structures of $\delta_{a/c}$ and η .

current and neighbor's states. The $c_{c(0)}$ affects the f_s state of a BC^1 , but the $c_{d(1)}$ does not affect any neighbor cells.

Due to the setting that the f_s of all BC s are initialized into f_i before action decisions, the $c_{c(0)}$ decision is performed conservatively. In detail, the $c_{c(0)}$ decision is related to the receiver's availability of a BC^0 , and the availability can depend on the f_s of BC^1 . In this situation, if a BC^0 of a CC is actually available at the time, and the δ_c of the CC is invoked before δ_a of BC^1 during simulation, c_s would be unchanged as c_d . This conservative decision leads that the c_s as $c_{c(0)}$ would be retained at the clock cycle. On the other hand, $c_{d(1)}$ is determined aggressively, so c_s can be set as $c_{d(1)}$ to c_c interchangeably at the current time.

After the decision of f_s or c_s and the execution of state-related operations by the engine, the changed current state can affect some neighbors' state decisions, and η is defined to provide the information of influenced neighbor cells from the current change. If the updated current cell state meets any pre-/next conditions of the opposite action of a next/previous cell, the cell is considered as a changeable cell, as shown in Fig. 4(c), and the group of changeable cells is informed to the simulation engine.

B. Scheduling-based Simulation Algorithm

The CA high-level states, except for c_s or f_s , are updated by the simulation engine after performing operations of the flit fetching and the connection change. The overall algorithm for the CA state transition at each iteration, is described in Algorithm 1. After the action state decisions of CA end at a current time, the time advances to the next clock-trigger time, and the action transitions of CA are repeated.

Algorithm 1: Simulation algorithm at each time step.

```

1  $L_p$  : a list of previously scheduled CA
2  $L_n$  : a list of newly scheduled CA for decisions
3  $L_i$  : a list of influenced CA
  /* Move each flit or reflect connection changes */
4 foreach cell in  $L_p$  do
5   if cell type is BC then
6      $BC^0$ .buffer.insert( $BC^{-1}$ .buffer.pop())
7      $BC^0$ . $f_s \leftarrow f_i$ ,  $BC^{-1}$ . $f_s \leftarrow f_i$ 
8     if  $BC^0$ . $q > 0$  and next CC is not designated then
9       request connection(s) to next CC(s) using the
10      global routing function
11      schedule next CC(s) to  $L_n$ 
12   else
13     if  $CC^{(-1,0)}$ . $c_s = c_{c(0)}$  then  $CC^{(-1,0)}$ . $c_s \leftarrow c_c$ 
14     else  $CC^{(-1,0)}$ . $c_s \leftarrow c_d$ 

  /* Decide new  $f_s$  and connection state */
15  $L_p \leftarrow \emptyset$ 
16 foreach cell in  $L_n$  do
17   if cell is not executed as an influenced cell then
18     if cell type is BC then
19       execute  $BC^0$ . $\delta_a$ 
20       if  $BC^0$ . $f_s = f_f$  then sched. cell to  $L_p$ 
21     else
22       execute  $CC^{(-1,0)}$ . $\delta_c$ 
23       if  $CC^{(-1,0)}$ . $c_s = c_{d(1)}$  then sched. cell to  $L_p$ 
24       else if  $CC^{(-1,0)}$ . $c_s = c_{c(0)}$  then
25         update  $CC^{(-1,0)}$ . $p$  using an arbitration
26         function, then schedule cell to  $L_p$ 

27   schedule cell. $\eta$  to  $L_i$ 
  /* Process the influenced cells */
28   while  $L_i$  is not empty do
29      $cell_i \leftarrow L_i$ .pop()
    /* Repeat upper sequential sentences
    from line No. 17 to 25 */
    exec.  $cell_i$ . $\delta_{a/c}$ , ... and sched.  $cell_i$ . $\eta$  to  $L_i$ 

```

The overall algorithm can be divided into two phases: the first phase to perform flit-fetching operations and reflect the connect change and the second phase to decide f_s and c_s state. Regarding the flit flow, the first phase is to move flits at the clock trigger and the second phase is to identify the flit-fetching BC s at the next clock trigger. After fetching a flit, f_s s of related BC s are initialized and connection requests to reachable CC s are performed. In the perspective of connection, the first phase removes the connection for new connectivity and stabilizes $c_{c(0)}$ by setting c_c to prevent the change of connected BC pair until c_s is set to $c_{d(1)}$.

Since the execution order of BC can be different at each iteration, the buffer of each BC has one redundant slot because a flit can arrive before an existing flit being fetched. The f_s or c_s of each cell in L_n , which is a list of state-changed cells

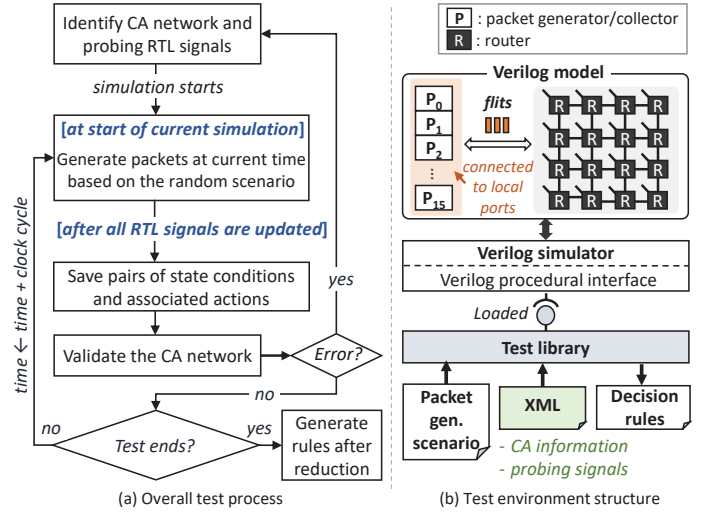


Fig. 5: Overall test process and environment structure.

while executing operations in the first phase, is decided if the cell is not already executed as an influenced cell. The execution priority of the influenced cells is higher because those cells are determined as changeable cells using η , so early detection of active cells whose action state are f_f or $c_{c(0)}$ can eliminate useless iterations. Therefore, the influenced cells are executed in the *while-loop*. Depending on the decided action state, cells are scheduled to L_p for the execution of action operations at the next time.

III. TEST-DRIVEN RULE SET GENERATION USING RTL NOC DESIGNS

The purpose of the test environment is to validate the CA network and to provide rules for the CA action-state decision. The overall test process is shown in Fig. 5(a). First of all, the designer identifies the RTL modules that hold the flits using FFs for BC s and arbiter modules for CC s by analyzing a simulation result of a few routers for a small sample of packet transmissions. Second, the designer constructs networks using the simulation result and the input and output information of arbiter modules. Using codes of the RTL modules for BC s, the designer identifies probing RTL signals that correspond to the CA states.

Based on the obtained CA network and probing RTL signal information, a network validation and saving pairs of conditions and their associated actions are performed. These are achieved using the proposed test environment with runtime NoC simulation for randomly injected packets. After the simulation, the environment generates decision rules for each cells after reducing the saved condition-action pairs. The overall structure of the test environment is illustrated in Fig. 5(b). To intervene during RTL simulation, the test library should be developed using a programming language interface (PLI) of the simulator, such as the Verilog procedure interface (VPI). The PLI enables an external library to access RTL signals and register functions of the library as callbacks to be invoked during the runtime simulation.

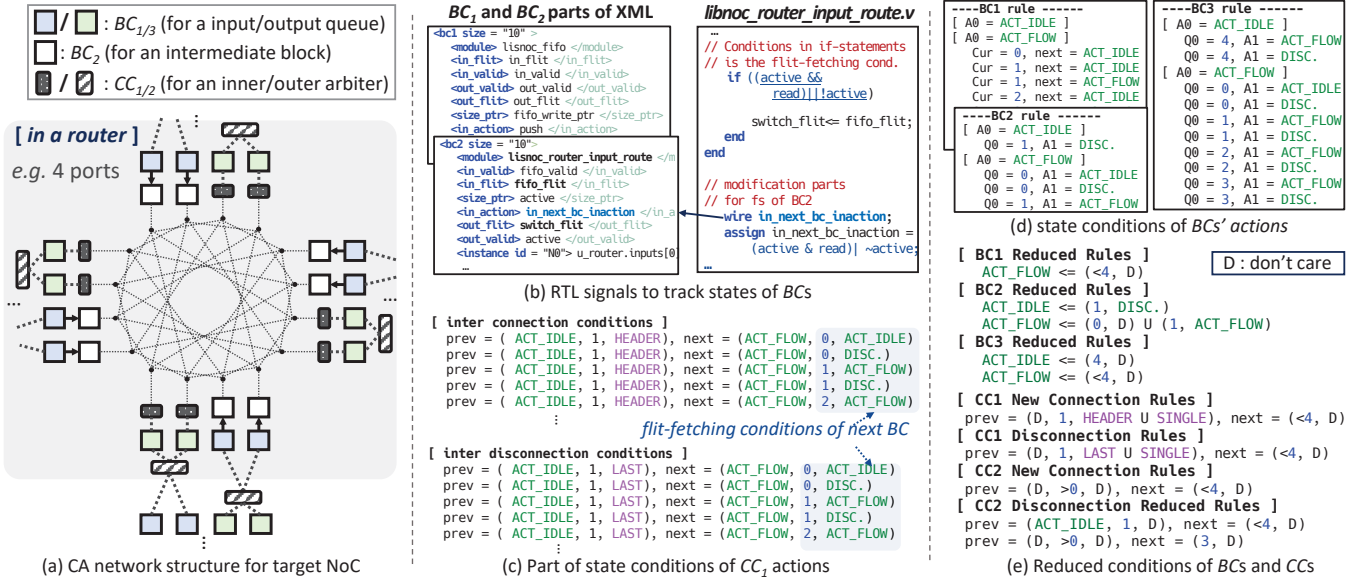


Fig. 6: Overall CA network of target RTL NoC and test-driven state conditions for action decision.

To achieve the test purpose, the proposed library references three types of time-related callbacks: a callback called at a specific time before any RTL signal updates, a callback called after all RTL signals are updated at the time, and a callback called at the end of the simulation. For Verilog, the callback types are *cbAtStartOfSimTime*, *cbReadOnlySynch* and *cbEndOfSimulation*, respectively.

At every clock cycle before RTL simulation, a callback in the test library is invoked to check the existence of scheduled packets at the invoked time based on the input scenario. If any scheduled packets exist, the callback requests packet generations to relevant RTL generators by manipulating their input signals.

After the current time RTL simulation, a callback is invoked to save states of BCs and connected BC pairs of CCs to identify state conditions of actions. The state of BCs can be determined by checking the values of probing signals, and the connected BC pairs can be obtained by comparing the valid output and input flit values between candidate BC pairs. The connection pair enables us to validate the CA network. The test library manages the current and previous states and connected BC pairs.

To detect conditions of $c_{c(0)}$ and $c_{d(1)}$, the test library compares current BC pairs with the pairs at the previous clock cycle. If the previous connection pair of the $CC^{(i-1,0)}$ does not exist or does not match the current connection, the connected state of the CC is newly changed and the neighbors' states at the clock cycle are one of the $c_{c(0)}$ conditions. If the previous pair does not match the current pair, the previous neighbor states are one of the $c_{d(1)}$ conditions if the previous f_s of the BC^0 was f_f , which means that a last transmission has happened at the current clock trigger. If the previous f_s of the BC^0 was f_i , we do not consider the current neighbor states of the CC as a recent disconnection. After updating the

connection pairs of CC, if the default pre-condition of a BC is met, related current and next states are saved as one of the conditions of the BC action. The conditions are managed by multiple rule sets that are shared by multiple cells based on the symmetry of representing RTL units.

After testing until the number of conditions of each rule is converged, we can reduce conditions using the following principle: if one or multiple state values of conditions are identical and all possible values of rest states are observed, then the rest states do not need to be considered during the action-state decision. For example, in the situation that two conditions $(0, f_i)$ and $(0, f_f)$ are observed, and f_i and f_f are all possible cases when the first state is 0, the action decision only reference the first state.

IV. EXPERIMENTATION

We applied the proposed M&S approach to an open source Verilog NoC, CONNECT [6]. In this experiment, we targeted an NoC with a 4-by-4 routers topology. We identified a CA network validated by the test environment, as shown in the Fig. 6(a). There are three kinds of BCs, which represents an input queue, an output queue, and an intermediate unit and two types of CCs, which serve an arbiter to resolve contention between the input ports and another arbiter to resolve contention between the virtual channels (VCs) at output ports. We developed three global functions: a routing function based on the XY routing scheme and two round-robin arbitration functions for the two arbiters.

To track the states of BCs, we identified RTL signals; the input XML for the signals is described in Fig. 6(b). The input and output queue instances of the routers share the same module. The test library was developed using the VPI provided by the *ModelSim 10.5c* in Quartus®Prime Pro 17.1. Using the test environment with input XML and various packet

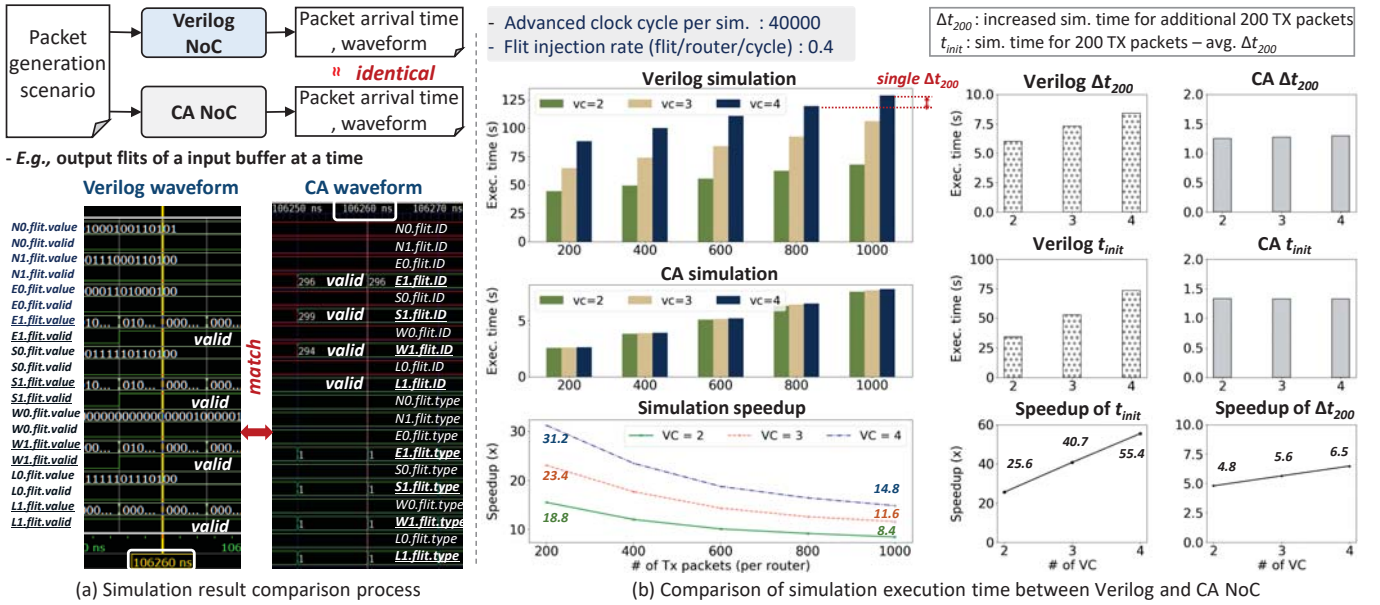


Fig. 7: Comparison of simulation results between Verilog and CA NoCs.

generation scenarios, we obtained various conditions for CA actions, and parts of the conditions are shown in Fig. 6(c)(d). After sufficient random testing and the convergence for the conditions, overall rules of each cell were obtained after the reduction, as shown in Fig. 6(e).

Based on the rule set, we developed the high-level CA and simulation engine following the proposed method using C++. To check and debug the equivalence of flit transmission between the Verilog and the CA NoC model, we enabled the NoC CA simulator to share packet generation scenarios with the Verilog test environment and to visualize the FT states of all BCs . For various packet generation scenarios, we confirmed identical arrival times of the packets between the Verilog and CA model, and an example of the matched output flit values of a queue is illustrated in Fig. 7(a).

For the simulation speed comparison, we measured the execution time of Verilog and CA NoC models by increasing the number of injecting packets per node, keeping the 0.4 average flit-injection rate per cycle to prevent packet loss caused by traffic congestion. We used an experimental machine whose CPU is Intel@Xeon@E3-1505M 3GHz and memory size is 64 GB. Using the proposed M&S approach, we obtained the speedup from 8.4 to 31.2 \times , as shown in Fig. 7(b).

We observed that the number of increased packets and increased execution time have a linear relationship. Based on the linearity, we defined two parameters, which are Δt_{200} to measure the runtime speedup and t_{init} to measure initialization speedup. Based on the obtained values of Δt_{200} and t_{init} , we can expect that if the number of injected packets goes zero, the CA speedup reaches about 24.6/40.7/55.4 \times (for the number of VCs is 2/3/4). If the number of injected packets grows arbitrarily large, the speedup will be converged to about 4.8/5.6/6.5 \times .

The proposed M&S method requires the designers to de-

velop the routing and arbitration functions. Depending on the target NoC, various routing and arbitration functions, such as adaptive routing and quality-of-service supporting arbitration algorithms can be implemented. We will validate the flexibility by applying various RTL NoCs. Moreover, we will improve the test environment to support the automatic identification of the CA network.

V. CONCLUSION

We proposed a high-level M&S method using a newly extended CA to reduce the RTL simulation execution time focusing on packet transmissions. The proposed CA defines detailed RTL operations in two steps: First, the action states are determined (that are related to the flit-fetching and connection changes) using the defined high-level states of current and neighbor cells based on test-driven rule sets. Second, the detailed operations corresponding to the action states are simulated. For the scalable simulation, we proposed an influence function to schedule changeable cells to be mainly activated by the simulation engine. Our evaluation showed various CA speedup (up to 31 times) relative to the fixed NoC size and the number of injected packets with valid transmission results.

REFERENCES

- [1] A. E. Kiasari, Z. Lu, and A. Jantsch, “An Analytical Latency Model for Networks-on-Chip,” *IEEE Trans. on Very Large Scale (VLSI) Syst.*, vol. 21, pp. 113–123, 2013.
- [2] E. Fischer and G. P. Fettweis, “An Accurate and Scalable Analytic Model for Round-Robin Arbitration in Network-on-Chip,” in *NoCS*, 2013.
- [3] A. Portero, R. Pla and J. Carrabina, “SystemC Implementation of a NoC,” in *ICIT*, 2005.
- [4] M. Briere et al., “System Level Assessment of an Optical NoC in an MPSoC Platform,” in *DATE*, 2007.
- [5] T. Toffoli, “Cellular Automata Machines,” *Cambridge, MA: The MIT Press*, 1987.
- [6] M. K. Papamichael and J. C. Hoe, “CONNECT: Re-examining conventional wisdom for designing NoCs in the context of FPGAs,” in *FPGA*, 2012.