

Static Timing Analysis for Modeling QoS in Networks on Chip

Evgeni Krimer^{a,*}, Isaac Keslassy^b, Avinoam Kolodny^b, Isask'har Walter^b,
Mattan Erez^a

^a*Department of Electrical and Computer Engineering, University of Texas at Austin,
Austin, TX 78712, USA*

^b*Department of Electrical Engineering, Technion — Israel Institute of Technology,
Haifa 32000, Israel*

Abstract

Networks-on-chip (NoCs) are used in a growing number of SoCs and multi-core processors. Because messages compete for the NoC's shared resources, quality of service and resource allocation are major concerns for system designers. In particular, a model for the properties of packet delivery through the network is desirable. We present a methodology for packet-level static timing analysis in NoCs. Our methodology quickly and accurately gauges the performance parameters of a virtual-channel wormhole NoC without simulation. The network model can handle any topology, link capacities, and buffer sizes. It provides per-flow delay analysis that is orders-of-magnitude faster than simulation while being significantly more accurate than prior static modeling techniques. Using a carefully derived and reduced Markov chain, the model can statically represent dynamic network state. Usage of the model in a placement optimization problem is shown as an example application.

Keywords: Network-on-Chip, Wormhole Routing/Switching, Virtual Channels, Static-Timing-Analysis, Analytical Model, Quality of Service

*Corresponding author

Email addresses: krimer@mail.utexas.edu (Evgeni Krimer), isaac@ee.technion.ac.il (Isaac Keslassy), kolodny@ee.technion.ac.il (Avinoam Kolodny), zigi@tx.technion.ac.il (Isask'har Walter), mattan.erez@mail.utexas.edu (Mattan Erez)

1. Introduction

Networks-on-chip (NoCs) are increasingly used instead of buses and dedicated signal wires in large-scale processors and, even more so, in modern systems-on-chip (SoC) [5]. In NoC-based systems, data transmission takes the form of multi-packet flows routed through the NoC over multiple links and routers. Because network resources are shared, per-flow quality of service (QoS) becomes a major concern for the system architect. Packet delay and data throughput are primary metrics of QoS [9].

The purpose of this paper is to rigorously derive a delay and throughput model for packet-level *static timing analysis (STA)* for NoC-based SoCs. Static timing analysis in a shared network is a non-simulation-based technique to estimate the average delay of each flow in the network, given the network topology, link capacities, router architecture, and the bandwidth requirements and characteristics of all flows.

The motivation for a per-flow STA technique is to enable a range of design optimizations that can rely on accurate and fast network analysis. Methods such as module placement and resource allocation [39, 1] require a large number of iterations, and thus the evaluation of network performance within each iteration must be very efficient. Until now, an accurate and complete modeling of advanced NoCs has only been possible with detailed and time-consuming simulations. The main reason is that network resources, including links, routers, buffers, and ports, are shared between several information flows. Thus, contention can arise inducing *statistical uncertainty* in the delay of each packet. Detailed simulation, however, is too slow to be effective within an optimization inner loop because all internal buffers and states must be modeled on a cycle-by-cycle basis.

Contributions

We present a rigorous analytical model that relies on a carefully constructed and reduced Markov chain to represent network state, including the occupancy

of all buffers. Our model is inspired by industrial work-flow modeling techniques and, to the best of our knowledge, is the first that can accurately account for arbitrary network topology, link capacities, and buffering, when using wormhole routing with virtual channels. We rely on the well-developed theory of stochastic processes and show that our technique faithfully predicts network queuing delay for both synthetic and real-world SoC traffic scenarios. We present results and validate the model for the delay analysis of flows with fixed-length packets that are composed of a large number of flits. We discuss extensions to these assumptions as future work. Note that our model is valid for any interconnection network, and not just an NoC, but the type of optimizations it enables are particularly well suited for the tight resource constraints and design flow of NoCs and SoCs.

To summarize our contributions:

- We present the first rigorous NoC model that is based on stochastic theory and show how to represent and solve for the network state using a Markov chain.
- We show how to account for arbitrary and finite buffering, as well as support wormhole routing and virtual channels. We use network delay analysis as an illustrative example of the modeling technique
- We validate our model for statistical traffic behavior using synthetic and real-world scenarios, and discuss why it is more complete and more accurate than prior analytical models.
- We demonstrate that our model can serve at the core of a design optimization method by showing that it can faithfully choose between multiple placement options in a real-world SoC example, and do so while requiring *orders of magnitude* less time than simulation. We also show that the most advanced prior-art model fails to make the correct optimization decision.

The rest of the work is organized as follows. We start by discussing the

related work in Section 3. Then, in Section 4, we establish a general analytical model for the average delay of each flow in an arbitrary NoC topology. We evaluate the delay model in Section 5 by comparing it with accurate simulation results and with previous delay models. Finally, Section 6 uses the model in a placement optimization tool and provides more insights and simulation results for the model.

2. Background and Modeled Network

The methodology we present in this paper is general and can be used with a wide array of NoC configurations, including arbitrary topology and buffering. In particular, we target NoCs with efficient *Wormhole Routing* in conjunction with *Virtual Channels* that use deterministic routing algorithms. This section provides a brief introduction to these techniques. Later in the paper we provide more details on other assumptions we make (Section 4). We also validate the model using a specific 4×4 mesh NoC example. The full details of this example NoC, which are presented in Section 5, are not pertinent to explaining the model and methodology because it is valid for any topology.

Wormhole Switching, or *Wormhole Routing* [33], is a flow control technique that can improve latency while reducing required buffering. In wormhole routing, large packets are broken into small (constant length) units called *flits* (flow control digits) and two control flits are added, a *head* flit at the beginning and a *tail* flit at the end of each packet. The head flit contains all the routing relevant information and establishes a path through the network. The subsequent flits follow the head flit in a pipelined fashion through the intermediate routers, which need only buffer a few flits at a time rather than the entire packet. In this way, the overall routing latency is reduced with pipelining, while intermediate nodes require smaller buffers. The tail flit is used to signal the end of the message and for various bookkeeping purposes. Wormhole routing is used widely in applications ranging from local computer clusters [31], through SoC NoCs [15], to SpaceWire [37, 36] chips.

Wormhole routing has several disadvantages, with the most serious being the long duration in which a channel in the network is occupied. A particular channel in the network is occupied from the time a head flit arrives until the time the corresponding tail is processed, and the link cannot be used for any other packet. Because flit transmission is pipelined, multiple links in the network are occupied by a single packet at any given time, detrimentally impacting performance.

To mitigate this effect, *Virtual Channels* (VCs) [16] allows a router to serve several packets simultaneously, by time multiplexing the physical link resource. This is done by additional logic and buffering in the router that maintains the status of several flows. Each of these flows is referred to as a virtual channel, and is allocated when header flit is received by a router. The router can continue processing other packets as long as virtual channels and buffers are available. The different VCs arbitrate for the physical link and many arbitration policies have been studied (e.g., [40]).

3. Related Work

Much of the prior work on analytical delay modeling in wormhole-enabled networks approximates the mean delay of packets in the entire system rather than estimating the delay of each source-destination flow separately [35, 30, 26, 24, 3]. Such gross approximations are often inadequate, and in such cases cannot be used in the NoC design process to efficiently optimize the allocation of resources.

In addition, while state-of-the-art NoC architectures multiplex multiple packets on the network links using virtual channels [29, 7, 32, 43, 6], most existing analytical models do not support virtual channels [18, 12, 45, 14, 34]. Further, in [44, 22, 27, 41], the authors formulate worst-case latencies of flows in the NoC. While this approach is suitable for real-time flows with hard deadlines, the vast majority of communication in typical SoCs has a set of more relaxed timing requirements, which can be satisfied with statistical guarantees.

Network calculus was recently proposed for analytical modeling and evaluation of NoC traffic [42, 4, 41]. This technique, however, does not currently model virtual channels, which are common to practically all state-of-the-art NoC architectures.

A heuristic approach to estimate the average delay of each flow was taken in [21]. The authors developed a *heuristic delay model* (HDM) that takes into account the capacities of all the links traversed by the modeled flow, as well as the bandwidth consumed by all other network flows which share some links with the modeled flow. Their heuristics attempt to estimate the serialization and head-of-line blocking, which add to the delays of a link because of congestion further downstream. This approach is useful, providing a closed-formed formula to estimate the delay for each flow in the network based on traffic parameters of all the flows. However, the model uses heuristics and its accuracy has not been confirmed in a rigorous fashion. In addition, [21] does not capture the effects of finite buffer sizes. In Section 5, we compare this model against ours and show how its heuristic approach can lead to a wrong optimization decision.

4. Analytical Model

In this section, we derive an analytical approximation of the delay of a flow in a wormhole network. The model supports any network topology and arbitrarily-set link capacities, as well as VC buffer sizes. Likewise, we place no restriction on the routing algorithm except that it be deterministic. To simplify the analysis, we assume that all packets have a fixed length, and are long enough so that the time it takes the head flit to arrive at the destination is considerably shorter than the time it takes the source to inject the packet. We further assume that there is no blocking in the network due to a lack of virtual channels, and that the destination node can always receive packets from the network. We also assume that the system is stable, i.e. that the number of packets in the system does not grow unbounded.

In the following analysis, we consider the latency induced by routers to be

only caused by the contention for shared VC buffers and by the time-domain multiplexing of VCs. Other (fixed-duration) latencies (e.g., routing and switching latencies) can be easily accounted for with minor modifications to the presented model. We first derive the model assuming that the packet arrival times at the injection port of each node can be modeled by a Poisson random process. Then, in Section 4.3, we generalize the model to any random process for packet arrival times. In Section 7 we discuss extensions to our model that further relax the current assumptions on packet length.

Our technique follows three main steps:

1. We focus on the NoC service for a particular flow of interest, which we generically call *flow X*, and model it using a Markov chain (MC) [10]. The Markov chain represents the network state of the routers and buffers on the path of flow *X*, as well as the impact of *interfering flows*, i.e., those flows that share at least one link with flow *X*.
2. We derive the flit propagation characteristics by computing the stationary distribution of the Markov chain. We use discrete time to model the internals of the NoC (Section 4.1.1).
3. We use the derived properties and standard analysis of M/G/1 (extended to G/G/1 in Section 4.3) queues to model the expected packet delay and the throughput of flow *X*.

4.1. Constructing the Markov Chain

4.1.1. The Reduced Configuration

To fully represent the NoC as a Markov chain, the internal state of each router (and in particular the buffer occupancies) as well as the characteristics of all flows need to be expressed as states in the chain. Unfortunately, this naive approach would result in an enormous and intractable number of states.

As shown in the transition from Figure 1a to Figure 1b, to reduce this Markov chain to a manageable size, we generate a separate model for each "isolated" flow, generically represented as flow *X*. We call this model the *reduced*

Table 1: Definition of symbols used in model derivation.

λ_α	packet arrival rate of flow α [packets per unit time]
M_α	packet length of flow α [flits]
$\phi^{(l)}$	capacity of link l [flits per unit time]
s_i	Markov chain state i
π_i	stationary distribution probability of state s_i
γ_i	fraction of packets of a measured flow (X) served while in state s_i
ρ_i	NoC throughput associated with state s_i for flow X [packets per unit time]
η_i	head-flit propagation delay of packets in flow X incurred while in state s_i [time units]
τ_α	expected time to fully transmit a packet of flow α [time units]
$\Delta^{(l)}$	buffer capacity on link l for flow X [flits]
$\delta^{(l)}$	buffer <i>occupancy</i> on link l for flow X [flits]
ξ	error in delay
F	total number of flows in the network
R	total number of nodes (routers)
Ψ_α	number of interfering flows for flow α
P_α	number of hops between the source and the destination for flow α

configuration. In the reduced configuration, we limit the analysis to the routers on the path of X , and only consider those other flows that share network links with X , such as flows A and B , but not flow C (this particular reduction method was also used in [20, 21, 46], and is further analyzed in Section 5.1).

To further simplify the Markov chain, we restrict our analysis to epochs in which the flits of flow X are waiting to be served by the NoC. This last assumption eliminates the need to model the large buffers at the network injection points and permits to model them as infinite buffers. Without assuming that flow X is active, we would need to track the state of the network during periods of inactivity, which would complicate the MC representation. We discuss the

implications of this simplification in Section 5.2.

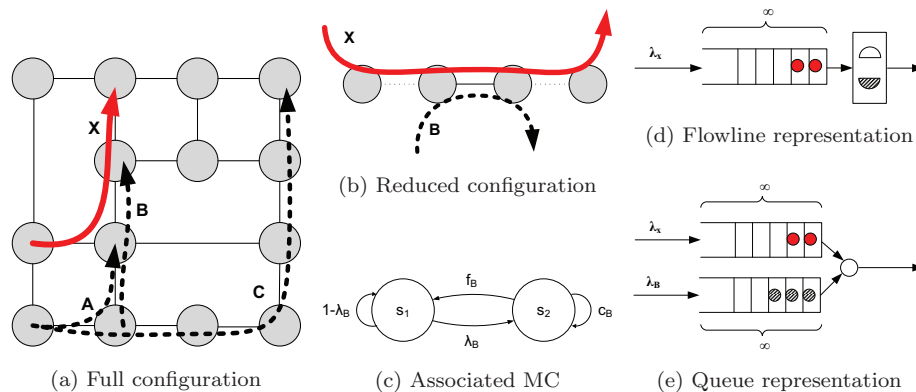


Figure 1: Flow X sharing a single link with a single interfering flow.

The reduced configuration for flow X can be viewed as a Markov chain in which each state is defined by the buffer occupancies and the existence of interfering flows along the links, as shown in Figure 1c and explained in Section 4.1.2. We construct a specific Markov chain for flow X to model its reduced configuration. This Markov chain representation accounts for both the extra queuing delay caused by interfering flows, which share links and ports with flow X , as well as for the delay of serialization and back-pressure within flow X and between the flits of a single packet (Figure 1e).

This system is equivalent to an open queuing network (Figure 1e) and to a manufacturing flowline (Figure 1d) with unreliable parallel machines [17]. By casting our problem in terms of unreliable machines, we can leverage a large body of works on stochastic theory and modeling methodology. We represent each hop taken by flow X as a production station consisting of a group of parallel machines. Likewise, in each cycle, the router is modeled as choosing a new machine in this group of parallel machines, in a round-robin fashion. A functioning machine processes flow X and contributes to its throughput, while a malfunctioning machine is equivalent to the link being used by an interfering flow.

Although the whole system operates in the continuous time domain, we use

discrete time to model the internals of the network. This simplifies the analysis of network properties, such as throughput, which we then apply back to the continuous domain. In our approximation, let t be the (NoC) cycle time and λ be the arrival rate of the packets in the continuous Poisson process, then the probability of packet arrival during a single cycle is $\lambda t e^{-\lambda t}$. Assuming that the cycle time is very short compared to the inter-arrival times, it can be simplified to a Bernoulli i.i.d processes for each cycle, making the probability of packet arrival during a single cycle simply λ .

In the following subsections we show how to construct the Markov chain for a number of representative interference patterns and conclude the analytical model section by deriving delay and throughput using the Markov chain (Section 4.2).

4.1.2. Sharing a Single Link with a Single Flow

In this scenario, measured flow X shares a single link with interfering flow B , as shown in Figure 1, which also depicts the flow-line and queue representations and the associated MC. The MC, which represents this single-interferer case, requires only two states (Figure 1c). State s_1 represents a “no interference” situation, where flow X can use the entire capacity of the shared link. If there are flits of both flows X and B waiting in the buffers, as illustrated in Figure 1e, the round robin arbitration mechanism allocates the shared link to each flow on every other cycle. As a result, each flow can utilize only half of the link capacity. State s_2 represents the situation where flow B interferes with the measured flow X . The MC does not need to represent a situation where only flow B is active, because it need only model the network as seen by flow X , and a separate MC is constructed to model the network properties of flow B .

The MC representation also expresses the probabilities of transitioning between the states, denoted in Figure 1c by the labels on the arrows connecting the two states. Starting from state s_1 , where only flow X is active, the probability of transitioning to state s_2 is simply the probability that a new packet of flow B arrives. Assuming Poisson arrival times, this probability is simply

(λ_B). Conversely, the probability of staying in s_1 is $(1 - \lambda_B)$. Starting from state s_2 , where both flows are active, the probability of transitioning to s_1 is the probability that the current packet of flow B is fully transmitted (f_B) and that no new packet of flow B has arrived during this transmission time. The time required to transmit a packet is the length of the packet (M_B) divided by the available link capacity ϕ , which is only half of the total capacity because flow X is also active: ($\tau_B = \frac{M_B}{\frac{1}{2}\phi}$). Thus, the probability of fully transmitting a packet at any time is $\frac{1}{\tau_B}$. The probability for another packet to appear during this time, which prevents transitioning back to s_1 is $(\lambda_B \tau_B)$. Therefore, the probability for transitioning from state s_2 to state s_1 is:

$$f_B = \frac{1}{\tau_B} \max(1 - \tau_B \lambda_B, 0) = \max\left(\frac{\phi}{2M_B} - \lambda_B, 0\right)$$

Finally, the probability for flow B to continue being active, and the MC to remain in state s_2 is:

$$c_B = 1 - f_B = \min\left(\left(1 - \frac{\phi}{2M_B}\right) + \lambda_B, 1\right)$$

Section 5.2 shows how to use this MC to derive delay and throughput properties for flow X , validates the results with detailed simulation, and discusses the implications and comparison to HDM.

4.1.3. Sharing a Single Link with Multiple Flows

A configuration where the measured flow shares a link with two other flows is shown in Figure 2 along with its flowline/queue equivalence and associated MC. Here, state s_1 represents no interference, states s_2 and s_3 represent inference by only flow A or only flow B respectively (i.e., flits of a single flow other than flow X are being multiplexed on the same link), and s_4 represents the state in which flits of all three flows (A , B and X) are multiplexed on the same link.

With the additional states, the MC is more complex and has a larger number of possible transitions. The evaluation of the transition probabilities is similar to the derivation discussed above. For state s_1 , the probability of staying in this state of no interference is the probability that no new packets arrive on

flow A and no new packets arrive on flow B ($(1 - \lambda_A)(1 - \lambda_B)$). Conversely, a transition from s_1 to s_4 occurs when both flows A and B become active at the same time ($\lambda_A \lambda_B$). The probability of transitioning from s_1 to s_2 is that of a packet arriving on flow A and not arriving on flow B , while the opposite is true for a transition from s_1 to s_3 .

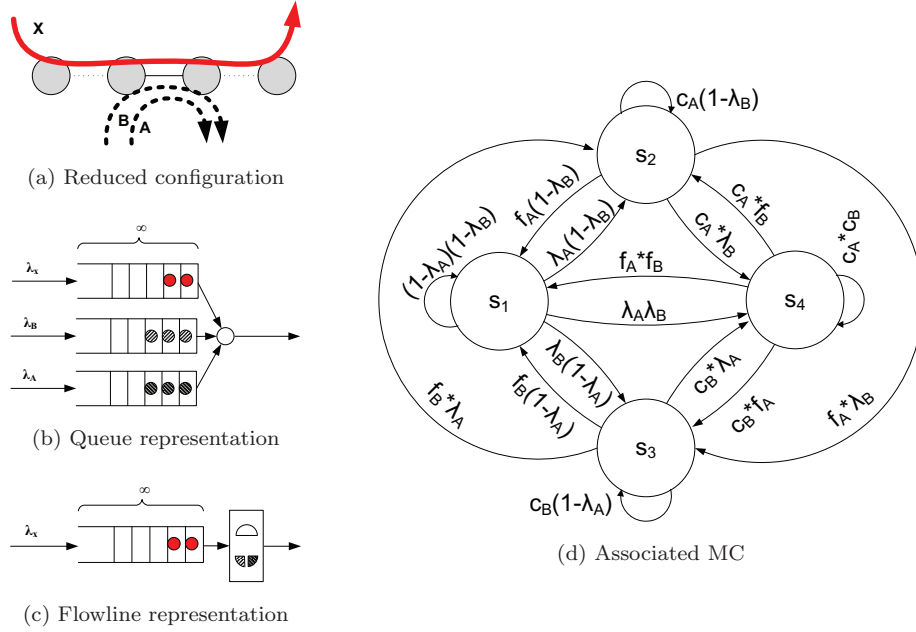


Figure 2: Flow X sharing a single link with multiple interfering flows.

We now discuss the transition probabilities from state s_2 . When a packet of flow A is fully transmitted and no other packets of both flows A and B arrived during the transmission time, then the MC transitions back from s_2 to s_1 . To calculate the probability of fully transmitting a packet, marked as f_A , we need to first derive the expected transmission time of a packet from flow A (τ_A). When a packet of flow A is being serviced, it can be done while at state s_2 at a rate of $\frac{\phi}{2}$, or in state s_4 with a rate of $\frac{\phi}{3}$ (because all three flows are active in s_4 but only two in s_2). Therefore, the expected transmission time is given by the expected fraction of time the packet is in states s_2 and s_4 , which are related to

the stationary distribution probabilities of these states (π_2 and π_4 respectively):

$$\begin{aligned}\tau_A &= \frac{M_A}{\frac{1}{2}\phi} \cdot \frac{\pi_2}{\pi_2 + \pi_4} + \frac{M_A}{\frac{1}{3}\phi} \cdot \frac{\pi_4}{\pi_2 + \pi_4} = \\ &= \frac{2M_A}{\phi} \left(1 + \frac{1}{2} \frac{\pi_4}{\pi_2 + \pi_4} \right)\end{aligned}$$

Given the expected transmission time, we can now write the probability of fully transmitting a packet as:

$$f_A = \frac{1}{\tau_A} \max(1 - \tau_A \lambda_A, 0) = \max\left(\frac{1}{\tau_A} - \lambda_A, 0\right)$$

Thus, the probability of transitioning from s_2 to s_1 is $(f_A(1 - \lambda_B))$. The probability for moving from s_2 to s_3 is that of fully transmitting the packet from flow A while a packet from flow B arrives at the same time, or $(f_A \lambda_B)$. The final two transitions are of remaining in s_2 or changing to s_4 , which occur when flow A continues and either a packet from flow B does not arrive (stay in s_2), or does arrive (move to s_4). As before, we mark the probability of flow A continuing activity as:

$$c_A = 1 - f_A = \min\left(\left(1 - \frac{1}{\tau_A}\right) + \lambda_A, 1\right)$$

The transitions out of state s_3 are derived in exactly the same way as for state s_2 , reversing the roles of flow A and flow B . With respect to state s_4 , the probability of transitioning to s_1 is the probability that both flows A and B are fully transmitted, or $f_A f_B$. Transitioning from s_4 to s_2 occurs when flow A continues and flow B finishes and the probability is $c_A f_B$. In symmetric fashion, the probability of s_4 to s_3 is $c_B f_A$. Finally, the probability of staying in state s_4 is $c_A c_B$ signifying that all flows remain active.

4.1.4. Sharing Multiple Links

A configuration of interfering flows over multiple links, including the equivalent queuing and flowline representations, is shown in Figure 3. Unlike the previous cases, this scenario includes a finite buffer that has to be taken into account. Flow X passes through two routers, each with an interfering flow, and therefore can experience back-pressure from the intermediate node that does not

have the infinite buffers assumed on the network injection and ejection packet queues. We assume that the intermediate flit queue has a depth of Δ flits, and will block transmission of an upstream node when it is full (Figure 3b).

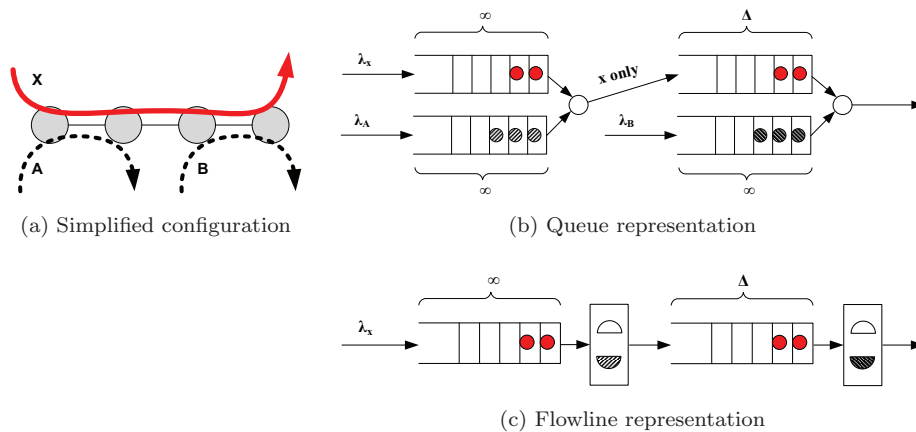


Figure 3: Flow X sharing multiple links with multiple interfering flows.

We explicitly model the occupancy of the intermediate buffer in the MC, by dedicating states to each possible buffer occupancy level. We show this in Figure 4 for the case where all link capacities are equal. The figure has three parts: Figure 4a shows only those states of the MC that correspond to an occupancy level of δ ; Figure 4b shows a schematic symbol that represents the partial MC of Figure 4a; and Figure 4c is the entire MC, using the schematic representation to simplify the figure.

Focusing on Figure 4a, state $s_{1,\delta}$ represents the case where only flow X is active and the buffer has occupancy level δ . State $s_{2,\delta}$ represents the case where flows A and X are active but flow B is inactive and $s_{3,\delta}$ is for when flows B and X are active but flow A is inactive. Finally, $s_{4,\delta}$ represents activity on all three flows and occupancy level δ . Starting from $s_{1,\delta}$ the buffer occupancy is not going to change in any scenario, because both routers are servicing flow X at the same rate and the buffer is emptied and filled at the same rate. The transition probabilities out of this state follow the same derivation described in Section 4.1.3. Similarly, state $s_{4,\delta}$ cannot change the occupancy because the

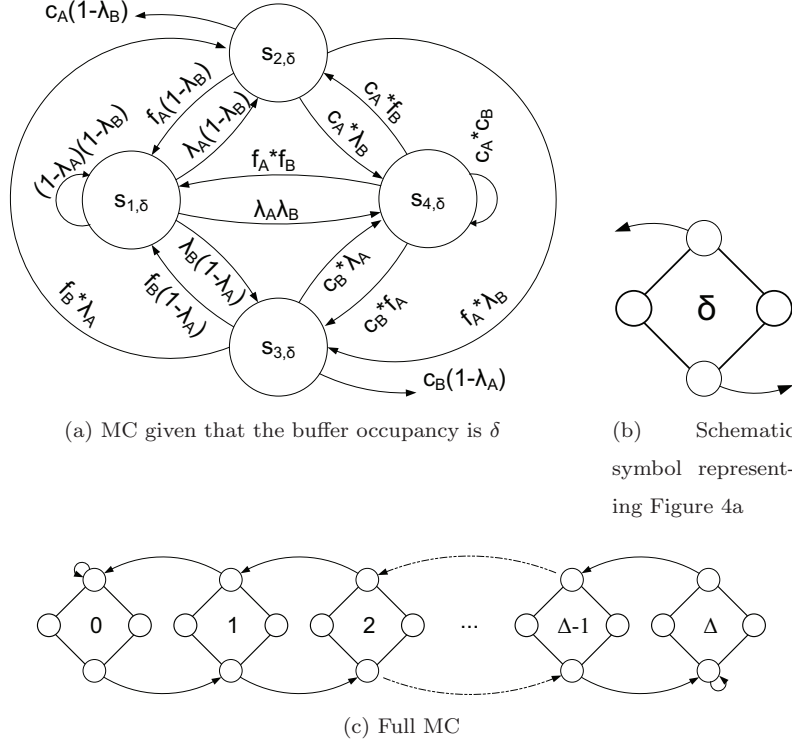


Figure 4: Markov chain representing flow X sharing multiple links with multiple interfering flows (Figure 3a).

service rate for flow X is equal to $\frac{\phi}{2}$ at both routers. Again, the transition probabilities out of $s_{4,\delta}$ follow the reasoning presented in Section 4.1.3. Because each link has only two multiplexed flows, however, the probability for fully transmitting a packet of the flow interfering with flow X is derived in the same manner as in Section 4.1.2:

$$f_\alpha = \max\left(\frac{\phi}{2M_\alpha} - \lambda_\alpha, 0\right) \quad \alpha \in \{A, B\}$$

$$c_\alpha = 1 - f_\alpha = \min\left(\left(1 - \frac{\phi}{2M_\alpha}\right) + \lambda_\alpha, 1\right)$$

Note that these probabilities of fully transmitting, or continuing with, an interfering flow are independent of the buffer occupancy. This is true because we assume that flow X is always active and that the ejection port of the network

is always available for each flow.

When flows A and X are active and flow B is inactive (state $s_{2,\delta}$), the service rate for flow X is higher in the downstream router, which is only servicing X , than in the upstream router that is servicing two flows. Therefore, the buffer occupancy decreases, which is represented by the transition from $s_{2,\delta}$ to $s_{2,\delta-1}$. This transition occurs with probability $(c_A(1 - \lambda_B))$, which is the probability that flow A remains active and that flow B remains inactive. The other transition probabilities of $s_{2,\delta}$ represent transitions that do not change buffer occupancy, either because flow B becomes active or flow A is fully transmitted and the service rate for flow X becomes equal at both routers. A similar analysis of state $s_{3,\delta}$ shows that if flow B continues and flow A does not become active, then the buffer occupancy increases and the MC transitions to state $s_{3,\delta+1}$. This is because the downstream node is now servicing flow X at half the rate of the upstream node.

The edge states $s_{2,0}$ and $s_{3,\Delta}$ are connected to themselves as shown on Figure 4c. Essentially, the buffer can never have fewer than zero flits and cannot exceed Δ flits. When the buffer is full, flow X is still being serviced by the downstream buffer and is not stalled (remember that we assume ejection is always possible). A full buffer causes back-pressure, which reduces the transmission rate in the upstream node. Our model inherently accounts for this back-pressure through the stationary distribution of the states, which directly determines the throughput as explained in the following subsection.

Observe that swapping the order of the interfering flows, i.e. swapping flows A and B in Figure 3a, would result in a symmetrical MC leading to exactly the same stationary distribution and estimated network properties. This is a significant improvement over the prior HDM technique, in which the order of the interfering flows affected the estimated delay. We further discuss this issue and show an example in Section 5.4.

4.2. Deriving Throughput and Delay

To derive the expected throughput of the NoC observed by flow X in the presence of interference we rely on the fact that the Markov chain we construct is positive recurrent and aperiodic. These properties imply that the random process corresponding to the state transitions, which represents the NoC, is ergodic. Using the ergodic theorem [10, 25] the expected throughput is given by:

$$T_X = \sum \pi_i \rho_i \quad (1)$$

$\bar{\pi}$ (a vector) is the stationary distribution of the MC, which can be computed by solving a system of linear equations.

Next, we express the expected delay of packets in flow X as:

$$L_X = W_X + H_X + \frac{1}{T_X} \quad (2)$$

W_X is the expected waiting time in the source node input buffer, which we derive in Equation 3. H_X , is the average propagation time of the head flit that we approximate in Equation 5. The final term, T_X , is the expected throughput observed by flow X as computed earlier and expressed in units of packets per unit time. Note that $\frac{1}{T_X}$ includes the serialization latency, which is simply the ratio of packet length (which is fixed in our model for simplicity) to link bandwidth (which can vary from link to link).

To express the input delay (W_X) we use classic results from the analysis of M/G/1 queues [25], which the NoC conforms to through the assumptions on Poisson packet arrival times and continuous service:

$$W_X = \frac{(1 + C_{S_X}^2)\lambda_X}{2T_X(T_X - \lambda_X)} \quad (3)$$

Note that we analyze the waiting time of flow X in the continuous domain using the M/G/1 model to match the continuous nature of the process used to model the flow. The throughput parameter was derived from a discrete-time model of the network (Section 4.1.1), which is the more appropriate representation of the internal NoC buffers and routing decisions.

The waiting time is dependent on the throughput (T_X), arrival rate (λ_X), and the coefficient of variation ($C_{S_X}^2$):

$$C_{S_X}^2 = \frac{\sigma_{S_X}^2}{S_X^2} \quad (4)$$

To calculate the average service time (S_X) and its variance ($\sigma_{S_X}^2$) we first compute the expected number of packets serviced at each state of the Markov chain, $\bar{\gamma}$. $\bar{\gamma}$ is the fraction of packets at each state, rather than the fraction of time spent at each state, which is $\bar{\pi}$.

$$\frac{\gamma_i}{\pi_i \rho_i} = \frac{\gamma_j}{\pi_j \rho_j} \quad \forall i, j$$

$$\sum \gamma_i = 1$$

We can now compute the average service time and its variance, and use the result to compute the coefficient of variation (Equation 4) and finally the waiting time (Equation 3):

$$S_X = \sum \frac{\gamma_i}{\rho_i} \equiv \frac{1}{T_X}$$

$$\sigma_{S_X}^2 = \sum \frac{\gamma_i}{\rho_i^2} - S_X^2$$

The last component of the expected packet delay (Equation 2) is the head-flit propagation delay:

$$H_X = \sum \gamma_i \eta_i \quad (5)$$

η_i is the head-flit propagation delay of packets in flow X incurred while in state s_i , measured in units of time.

Finally, for $S_X \gg H_X$ we can approximate the end-to-end packet delay as:

$$L_X \approx W_X + \frac{1}{T_X} \quad (6)$$

4.3. Model Generalization

In the previous section we derived the end-to-end delay assuming a Poisson arrival process. Therefore, we applied the M/G/1 queuing model. In this section, we generalize our method for any arrival process by applying the G/G/1 queuing model [28].

Let σ_a^2 be the variance of the inter-arrival time, then the input delay (W_X) (from Equation 3) can be generalized to:

$$W_X \cong \frac{1 + C_{S_X}^2}{\left(\frac{T_X}{\lambda_X}\right)^2 + C_{S_X}^2} \cdot \frac{\lambda_X (\sigma_a^2 + \sigma_{S_X}^2)}{2 \left(1 - \frac{\lambda_X}{T_X}\right)} \quad (7)$$

It is easy to show that for Poisson traffic, applying $\sigma_a^2 = \frac{1}{\lambda_X^2}$ to Equation 7 matches Equation 3 exactly.

5. Model Validation

This section illustrates how to apply the model to compute network properties such as delay and throughput. We use the scenarios described in Sections 4.1.2–4.1.4 and, for each case, compare the results of our model with detailed simulations and with HDM [21]. Our cycle-accurate, discrete-event, NoC simulator uses the OMNET++ framework [47]. The simulator simulates wormhole switching with virtual channels [15], deterministic XY/YX routing [8], and configurable network topology, buffers, and traffic parameters. The simulations were 8×10^7 cycles long, resulting in a confidence interval of 0.35% on average for a 95% confidence level. The simulated NoC implements the QNoC architecture, described in [9] using a single QoS level, with virtual channels added to each router port. The network properties are summarized in Table 2.

5.1. Isolation

In Section 4.1.1 we explained how to isolate the investigated flow X from the rest of the flows. By doing this we neglect to account for possible indirect interactions between other flows and flow X in different parts of the NoC, as illustrated in Figure 5. In this figure, flow C does not interfere with flow X directly, but flows A and B , which do interfere with flow X , interact through flow C .

We investigate the potential impact of indirect interference using the configuration shown in Figure 6, where flow X is interacting directly only with flow A , however, flow A might interact with other flows (B, C, D). We simulated

Table 2: Simulated NoC properties.

Dimensions	4×4
Message length	256 flits
Flit length	32 bit
Virtual Channels	4
Buffer size	5 flits
Routing	wormhole XY/YX
$Node \leftrightarrow Router$ capacity	400Gbps
$Router \leftrightarrow Router$ capacity	10Gbps
Router frequency	333MHz




Figure 5: Example of possible neglected interaction between interfering flows.




Figure 6: Configuration showing the isolation method

multiple configurations of indirect interference described in Table 3 and varied the arrival rate of flow X (λ_X), keeping other arrival rates fixed for simplicity ($\lambda_A = \lambda_B = \lambda_C = \lambda_D = 0.2$). The results (Figure 7) show that regardless of whether flows B, C, D are present, the impact on the end-to-end delay for flow X packets is only affected by flow A . This observation holds since the examined system is stable, and hence, indirect interference does not change the arrival properties of flow A as it interacts with flow X . Indirect interference is likely to impact the delay of flow X in unstable systems, but analysis of such systems

Table 3: Different interference configurations

configuration	A	B	C	D	configuration	A	B	C	D
I	x	x	x	x	IV	x	x		
II	x		x	x	V	x			x
III	x	x	x		VI	x			

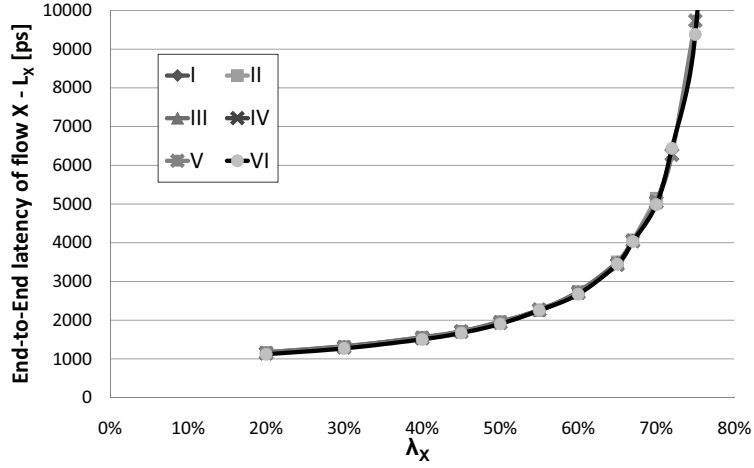


Figure 7: End-to-end latency for flow X packets in different interference configurations.

is beyond the scope of this paper.

5.2. Sharing a Single Link with a Single Flow

The simplest analyzed test case models a single link shared with a single flow, as shown in Figure 1b. The stationary distribution $\bar{\pi}$ of the Markov chain derived in Section 4.1.2 is the solution of the following set of linear equations:

$$\pi_1 \lambda_B = \pi_2 f_B \quad \pi_1 + \pi_2 = 1$$

Resulting in:

$$\pi_1 = \frac{f_B}{f_B + \lambda_B} \quad \pi_2 = \frac{\lambda_B}{f_B + \lambda_B}$$

Applying Equation 1 for the expected throughput provides:

$$\begin{aligned} T_X &= \pi_1 \rho_1 + \pi_2 \rho_2 = \frac{f_B}{f_B + \lambda_B} + \frac{\frac{1}{2} \lambda_B}{f_B + \lambda_B} = \\ &= \max \left(\frac{\phi}{M_X} - \lambda_B \frac{M_B}{M_X}, \frac{\phi}{2M_X} \right) \end{aligned}$$

T_B can be calculated in exactly the same way and following the steps in Section 4.2 L_X and L_B can be calculated (not shown). An important observation of this resulting throughput is that the worst-case throughput observed by flow X is $\frac{1}{2}$ of the link capacity, which corresponds to high contention with the interfering flow B .




Figure 8: Flow X sharing a single link with a single flow - variable λ_X .

Figure 8 shows the expected delay of packets in flow X as its throughput requirement, controlled by the arrival rate (λ_X) is increased, and where the arrival rate of the interfering flow B is constant ($\lambda_B = 0.4$). For these parameters, both our model and HDM match the simulation results well. As explained above, the minimum throughput observed by flow X is half of the link capacity, and this low throughput is reached when the interfering flow arrival rate exceeds 0.5. This phenomenon is not accounted for in HDM, because its heuristics were developed and tuned for low arrival rates. This error in delay estimation is apparent in Figure 9, which shows the case of fixed ($\lambda_X = 0.4$) and variable λ_B .




Figure 9: Flow X sharing a single link with a single flow - variable λ_B .

The HDM allows the interfering flow B to consume more than half of the link capacity, resulting in a sharp increase of the estimated delay of flow X , which tends to infinity for ($\lambda_X \rightarrow 0.6$). Our model, which inherently accounts for the minimum observed throughput, on the other hand, matches simulations well.

In addition, an interesting observation can be made regarding the degree of estimation error of our model (see Figure 9). The end-to-end latency of flow X when ($\lambda_X = 0.4 < \lambda_B$) has linear behavior with respect to λ_B . As can be seen, our presented model overestimates the end-to-end delay for this case, albeit with bounded error. The reason is that we assume that flow X is always active and has flits available for transmission. This assumption, however, partially fails when the interfering flow B has a higher rate than the investigated flow X . When flow X is not active, flow B observes higher service rate and is transmitted more quickly. This reduces the probability that it actually interferes with flow X and leads to the overestimation error.

While leaving more in-depth investigation of this phenomenon to future work, both in terms of curbing it and with respect to providing a tight bound, we now characterize the error behavior. First, we notice that because of this error, the end-to-end latency can only be over estimated and not under estimated. Second, the error can be roughly bounded by:

$$\xi < L_B(\lambda_X = 0.5) - L_B(\lambda_X = \lambda_B)$$

5.3. Sharing a Single Link with Multiple Flows

We now analyze the test case of a single link with multiple interfering flows, which corresponds to Section 4.1.3. In this scenario, HDM only considers the sum of the arrival rates of all the competing flows; hence, it cannot distinguish between a single interfering flow with ($\lambda = 0.4$) and two interfering flows with ($\lambda = 0.2$) each, for example.




Figure 10: Comparison of sharing a single link with three total flows (two interfering flows with ($\lambda_A = \lambda_B = 0.2$) (sim-triple and model-triple) and sharing the link with one interfering flow of equivalent rate ($\lambda_A = 0.4$) (sim-double and model-double). HDM models both cases as a single interferer with ($\lambda = 0.4$).

As shown in Figure 10, however, these two distinct cases result in very different throughput/delay characteristics. Our technique faithfully models the two cases, and closely follows simulation results.

5.4. Sharing Multiple Links

The last example is of sharing multiple links with different flows as shown in Figure 3a and discussed in Section 4.1.4. In this scenario, there is a finite buffer in an intermediate node that can back-pressure flow X .




Figure 11: Flow X sharing multiple links with multiple flows – variable buffer capacity.

Figure 11 demonstrates the impact of varying the buffer capacity for interfering flows of fixed rate ($\lambda_A = \lambda_B = 0.2$) and varying λ_X . Simulation results for buffer capacities of 5 and 300 flits are shown, along with estimates provided by the proposed model and by HDM. With lower buffer capacity, the peak throughput drops substantially, which our model accurately predicts. HDM [21], which

is oblivious to the buffer capacity, does not match the simulation results.

Finally, another aspect of this multiple interferer scenario relates to the order in which the interfering flows appear along the path of flow X . Following the conclusions of the analysis described in Section 4.1.4, our proposed model estimates the same network performance properties (throughput) regardless of the order of the interfering flows. HDM, on the other hand, is sensitive to the order in which the interfering flows are applied, as is evident in Figure 12, which shows results for two different configurations: Configuration A, where ($\lambda_A = 0.3$) and ($\lambda_B = 0.1$); and Configuration B where ($\lambda_A = 0.1$) and ($\lambda_B = 0.3$).




Figure 12: Flow X sharing multiple links with multiple flows – varying the order in which interfering flows interact along the path of flow X .

5.5. Non-Poisson Traffic

In this section we repeat some of the validation experiments for a different packet arrival time distribution. Results shown to this point use a Poisson

arrival process, and in this section we use the Uniform arrival time distribution — $Uniform[0.5 \cdot \frac{1}{\lambda}, 1.5 \cdot \frac{1}{\lambda}]$.

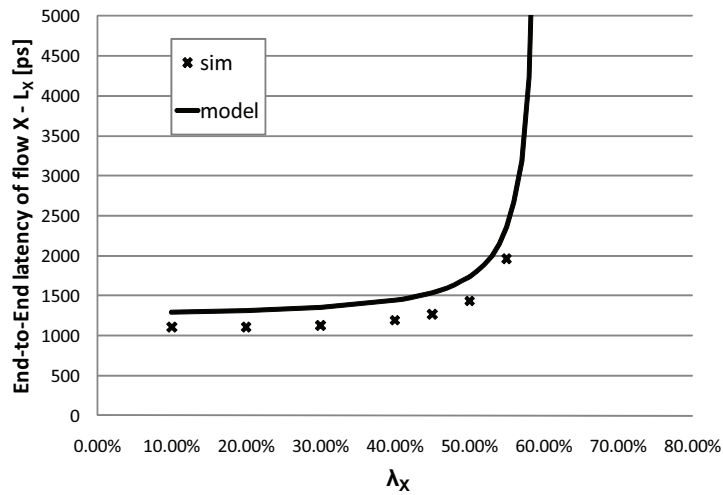


Figure 13: Uniformly distributed arrival times – Flow X sharing a single link with a single flow - variable λ_X .

First we repeat the experiments from Section 5.2, sharing a single link with a single flow as shown in Figures 13–14. We also show the results of sharing multiple links with different flows from Section 5.2 with uniformly distributed arrival times in Figure 15. Although these results are less precise compared to Poisson traffic, the model still presents a very good approximation for a general inter-arrival pattern.

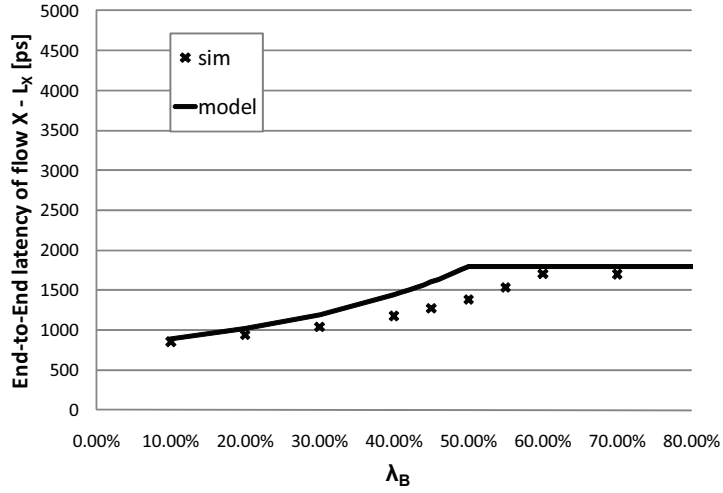


Figure 14: Uniformly distributed arrival times – Flow X sharing a single link with a single flow - variable λ_B .




Figure 15: Uniformly distributed arrival times – Flow X sharing multiple links with multiple flows.

Table 4: Audio-video benchmark traffic requirements from [23].

flow	src	dst	rate [kB/s]	flow	src	dst	rate [kB/s]
F_1	MEM1	ASIC4	1168730	F_{16}	ASIC4	DSP1	338480
F_2	ASIC2	ASIC1	800	F_{17}	DSP1	DSP2	338480
F_3	MEM1	CPU	752050	F_{18}	DSP8	DSP7	282650
F_4	MEM3	CPU	755840	F_{19}	DSP6	ASIC2	282480
F_5	ASIC4	CPU	1970	F_{20}	DSP1	CPU	203630
F_6	DSP3	DSP6	70610	F_{21}	DSP2	DSP1	203630
F_7	ASIC1	ASIC2	250	F_{22}	DSP3	DSP5	70610
F_8	DSP3	ASIC4	380160	F_{23}	DSP7	MEM2	70650
F_9	DSP8	ASIC1	800	F_{24}	MEM2	ASIC3	77050
F_{10}	DSP5	DSP6	269240	F_{25}	ASIC3	DSP8	6410
F_{11}	CPU	MEM1	380160	F_{26}	DSP4	DSP1	36720
F_{12}	CPU	MEM3	380160	F_{27}	ASIC2	MEM2	6400
F_{13}	CPU	ASIC3	380160	F_{28}	ASIC2	ASIC3	7650
F_{14}	DSP2	ASIC2	338480	F_{29}	ASIC3	DSP4	1440
F_{15}	DSP4	CPU	1970	F_{30}	ASIC1	DSP8	250

6. Benchmark Delay Model and Placement Optimization

In this section we demonstrate how our analytical model can be used in the inner-loop of many optimization algorithms, such as module placement, buffer allocation, link capacity allocation, and network topology selection. These inner loops cannot solely rely on simulations, as simulations take too long to complete, making analytical models crucial for an efficient design process. Further, the correctness of the analytical models directly affects the correctness of the optimization algorithm. Therefore, we show that our analytical model is both fast and accurate in real-world scenarios.

We first analyze the delay of all flows in a SoC using the audio-video benchmark presented in [23], with the traffic requirements summarized in Table 4. Using detailed simulations, Section 6.1 compares the accuracy and computation time of our proposed model against HDM. In this experiment the used simulations were 8×10^8 cycles long, resulting in a confidence interval of 0.44% on average for a 95% confidence level.

We then illustrate in Section 6.2 how our analytical model can be used in a *module placement* algorithm that attempts to minimize overall flow delay, whereas HDM fails to make a correct optimization decision.

6.1. Benchmark Delay Model

We evaluate our analytical delay model for the benchmark system and traffic requirements shown above assuming Poisson arrival times. We use a 4×4 mesh topology with the parameters detailed in Table 2, and the module placement shown in Figure 16, denoted *placement A*, where arrows indicate the different flows. For the simulations, we use the simulator described in Section 5, and run the simulations long enough for all performance characteristics of the different flows to stabilize.




Figure 16: *Placement A* of the components and flows of the audio-video SoC of [23].

Figure 17 shows the average end-to-end latency of each flow. It compares the simulation results with our proposed analytical model, as well as the analytical model presented in [21]. The total queuing delay is broken down into its components: the delay due to the source node queuing and the delay due to intermediate queuing, as well as the serialization delay (which is the ratio of packet length to link bandwidth). Due to space limitations, we only present the eight flows with the greatest relative slowdown, as ranked by simulation results and presented in Table 4 as (F_1, \dots, F_8) . In addition, we show the two flows with the longest distances traversed as presented in Table 4: F_{14} and F_{25} with 4 and 5 hops respectively.




Figure 17: End-to-end latency breakdown predicted by detailed simulation, our model, and HDM for the system requirements of Table 4 and Placement A (Figure 16).

Figure 17 shows that both the source node queuing delay and the intermediate queuing delay can be dominant (and significant) under different circumstances. One may observe that the dominance of the source node queuing delay component is correlated with the bandwidth requirement of the flow (for example in flow F_1). On the other hand, a low bandwidth flow competing over resources with other flows would suffer mostly from the second component, the intermediate queuing delay (for example in flow F_5).

As shown in Figure 17, our model approximates the simulation results significantly more accurately than HDM. In particular, Figure 18 depicts the absolute error of the queuing delay for each flow for both analytical models. We can see that the error for each flow is under 15% for our model, while the error of HDM is often above 50% and can be greater than even a *factor* of 10 times.

We also note that the time required to compute the results is orders of magnitude faster using our model than with the detailed simulation, requiring only $33ms$ as opposed to over 7 hours of simulation time.




Figure 18: Relative error in latency estimation between our model and HDM relative to detailed simulation for the system requirements of Table 4 and Placement A (Figure 16).

6.2. Placement Optimization

A possible use of the analytical delay model is to estimate network and flow properties within the inner-loop of a module placement optimization algorithm. As shown above, our model can quickly compute delay with high accuracy and

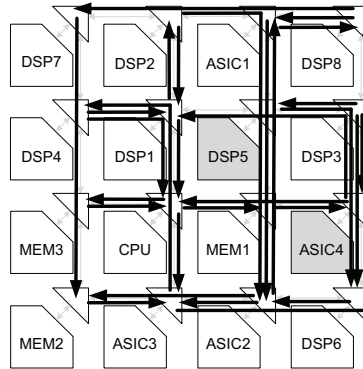


Figure 19: Placement B of the components and flows of the audio-video SoC of [23].

in this subsection we demonstrate that it also reflects the change in delay as a result of varying the module placement. Hence, our model can be used to predict, and correctly and efficiently choose between multiple placement options. Without loss of generality, we assume that there is a need to choose between two placements: *placement A*, illustrated in Figure 16, and *placement B*, shown in Figure 19, where modules *ASIC4* and *DSP5* have been swapped.

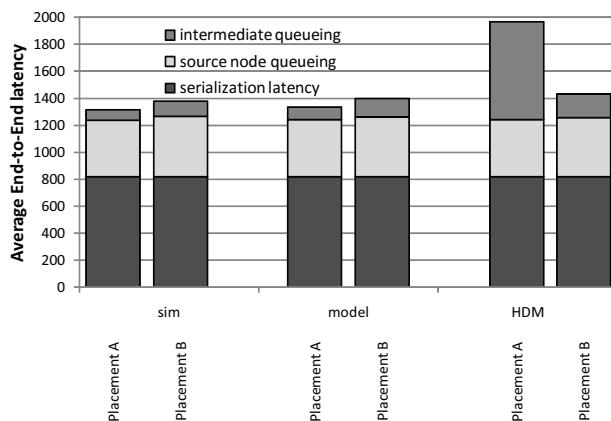


Figure 20: Comparison of estimated latency of placement *A* (Figure 16) and placement *B* (Figure 19) predicted by detailed simulation, the proposed model, and HDM.

Figure 20 shows the average end-to-end latencies for placements *A* and *B*. We break down the latencies to their components to demonstrate the relative impact of each latency source. The simulation columns in the figure show that the average flow delay is lower in placement *A*. This is accurately reflected in our analytical model, which closely approximates the simulation delays to within 3%, and would also have pointed to placement *A* as having a smaller overall delay. HDM, on the other hand, can be quite inaccurate, and as a result leads to an incorrect placement decision, predicting that placement *B* has lower overall latency than placement *A*. An interesting observation is that despite the fact that the intermediate queuing delay is the least dominant on average, it is

actually the delay component that changes the most between the placements A and B. In addition, in the analyzed example, the intermediate queuing delay is the component that determines the break even point, which makes its precise modeling important.

7. Discussion and Future Work

Our packet level static timing analysis (STA) model is constructed based on several important assumptions, which impact its accuracy as discussed below. We assume that flow X is always active, and is thus always competing for link capacity with the interfering flows. We discuss the implications in Section 5.2 and show that the resulting error is both bounded and limited to scenarios where the interfering flows require more throughput than flow X . We plan to extend our model to reduce this effect in future work, as well as provide a tighter bound on the possible error.

Unlike most prior modeling work, we show that our model is general and can apply to any stochastic process for packet inter-arrival times. We derived our model using the standard Poisson traffic assumption, but then showed how to apply a G/G/1 queuing model (Equation 7) that is accurate provided the first two moments of the distribution are known. We demonstrated the model's accuracy using both Poisson and Uniform inter-arrival time distributions.

We currently assume that packets are of constant length, and in future work we plan to investigate non-constant packet lengths. We have initial promising results for exponentially-distributed packet lengths, but presenting them here is beyond the scope of the paper. Finally, we approximated the end-to-end delay by neglecting the header flit propagation delay (Equation 6), which is only accurate for long packets. During our work, we verified that the queuing delay is accurately estimated for any arbitrary packet length and will fully validate the model for header-flit propagation delay (Equation 5) in the future. We will also test our model using a full-system simulator to validate the model's accuracy with real traffic.

An important aspect of our work is that we carefully reduced the number of Markov Chain states so that the model can be computed quickly and efficiently. First, we choose to ignore the network injection queues and assume that they are not the bottleneck in the system. Without this assumption we would need to model infinite queues (a common assumption in simulation), which would require an infinite number of states in the MC. Second, we assume that the flow under consideration can be isolated (Section 5.1) from all but the directly interfering flows. Without isolation, the MC would need to represent the states of all the available network flows (F) along with the occupancy of all intermediate buffers in all R routers. With a total number of F flows, there are exactly 2^{F-1} possible combinations of interfering flows. Each intermediate buffer occupancy can be in the range $[0..Δ]$, requiring $(Δ + 1)$ states. Taking into account that the number of intermediate buffers is $O(R)$, the upper bound on the number of states in the MC is an intractable $O(2^F * Δ^R)$. However, with isolation, we reduce the number of MC states to $O(2^{\Psi_\alpha} * Δ^{P_\alpha})$, where Ψ_α is the number of flows interfering with the specific flow of interest α) and P_α is the number of hops between the source and the destination for that flow. For instance, in the example discussed in Section 6, $F = 30$, $R = 16$, $\Delta = 5$, $\max(\Psi_\alpha) = 2$, and $\max(P_\alpha) = 6$. Thus, the number of states required (for each flow analysis) without isolation is $\approx 1.64 \cdot 10^{20}$ and only 20 (for the most complex case) after isolation is applied.

Our final assumption is that virtual channels are always available for any flow. Extending the model to account for this type of head-of-line blocking is left for future work. Additional further research should address more complex network topologies and routing schemes (e.g. the model may be generalized to allow adaptive routing).

Also as part of our future work, we will investigate the possibility of significantly simplifying the Markov chain by relying on MC decomposition methods developed for industrial flowline analysis [19, 11, 2, 38, 13]. This line of work promises to expedite the solution of complex NoCs by curbing the exponential growth in the number of states required to model a large number of interfering

flows.

Finally, while this paper demonstrated the application of our model for estimating throughput and end-to-end delay, the model inherently captures other network phenomena and parameters. For example, estimating buffer occupancy levels is a straight-forward extension, that uses the MC representation discussed in Section 4.1.4.

8. Conclusions

In this paper we introduced a packet-level *static timing analysis* (STA) for NoCs. We showed how it allows for a quick and precise evaluation of the performance parameters of a virtual-channel wormhole NoC without using any simulation techniques. It can handle any topology, link capacities, and buffer capacities — and unlike existing models, is able to evaluate the performance of a specific flow in a precise manner. Therefore, it enables per-flow QoS for NoCs.

Our new model allows for a per-flow STA that is orders-of-magnitude faster than simulation. Ultimately, the objective is for this packet-level STA model to be used in the inner-loop of NoC optimization tools — and become the packet-level equivalent of gate-level critical path analysis utilized in CAD tools.

Acknowledgments

We thank the anonymous reviewers whose comments helped to significantly improve the paper. We also thank Prof. Moshe Sidi for fruitful discussions at the early stages of this research. Finally, we thank Intel corp. and NVIDIA for providing equipment and funds in support of this research.

References

- [1] Ahonen, T., Sigüenza-Tortosa, D., Bin, H., Nurmi, J., 2004. Topology optimization for application-specific networks-on-chip. In: Proceedings of the 2004 international workshop on System level interconnect prediction. pp. 53–60.

- [2] Altıok, T., 1997. Performance Analysis of Manufacturing Systems. Springer, Berlin.
- [3] Alzeidi, N., Ould-Khaoua, M., Khonsari, A., 2008. A new modelling approach of wormhole-switched networks with finite buffers. *International Journal of Parallel, Emergent and Distributed Systems* 23 (1), 45–57.
- [4] Bakhouya, M., Suboh, S., Gaber, J., El-Ghazawi, T., 2009. Analytical modeling and evaluation of On-Chip Interconnects using Network Calculus. In: *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip-Volume 00*. IEEE Computer Society, pp. 74–79.
- [5] Benini, L., De Micheli, G., 2002. Networks on Chips: A New SoC Paradigm. *Computer* 35 (1), 70–78.
- [6] Bertozzi, D., Benini, L., 2004. Xpipes: a network-on-chip architecture for gigascale systems-on-chip. *Circuits and Systems Magazine, IEEE* 4 (2), 18–31.
- [7] Bjerregaard, T., Sparso, J., 2005. A Router Architecture for Connection-Oriented Service Guarantees in the MANGO Clockless Network-on-Chip. In: *Proceedings of the conference on Design, Automation and Test in Europe-Volume 2*. pp. 1226–1231.
- [8] Bolotin, E., Cidon, I., Ginosar, R., Kolodny, A., 2004. Cost considerations in network on chip. *INTEGRATION, the VLSI journal* 38 (1), 19–42.
- [9] Bolotin, E., Cidon, I., Ginosar, R., Kolodny, A., 2004. QNoC: QoS architecture and design process for network on chip. *Journal of Systems Architecture* 50 (2-3), 105–128.
- [10] Brémaud, P., 1999. *Markov Chains*. Springer, Berlin.
- [11] Buzacott, J., 1993. *Stochastic Models of Manufacturing Systems*. Prentice Hall, Englewood Cliffs.

- [12] Ciciani, B., Colajanni, M., Paolucci, C., 1997. An accurate model for the performance analysis of deterministic wormhole routing. In: Proceedings of the 11th International Symposium on Parallel Processing. p. 353.
- [13] Dallery, Y., Frein, Y., 1993. On decomposition methods for tandem queueing networks with blocking. *Oper. Res.* 41 (2), 386–399.
- [14] Dally, W., 1990. Performance analysis of k-ary n-cube interconnection networks. *IEEE Transactions on Computers* 39 (6), 775–785.
- [15] Dally, W., 2004. Principles and Practices of Interconnection Networks. Morgan Kaufmann Publishers, San Francisco.
- [16] Dally, W. J., 1992. Virtual-channel flow control. *IEEE Trans. Parallel Distrib. Syst.* 3 (2), 194–205.
- [17] Diamantidis, A., Papadopoulos, C., 2008. Exact analysis of a two-workstation one-buffer flow line with parallel unreliable machines. *European Journal of Operational Research*.
- [18] Draper, J., Ghosh, J., 1994. A comprehensive analytical model for wormhole routing in multicomputer systems. *Journal of Parallel and Distributed Computing* 23 (2), 202–214.
- [19] Gershwin, S. B., 1987. An efficient decomposition method for the approximate evaluation of tandem queues with finite storage space and blocking. *Oper. Res.* 35 (2), 291–305.
- [20] Guz, Z., Walter, I., Bolotin, E., Cidon, I., Ginosar, R., Kolodny, A., 2006. Efficient link capacity and qos design for network-on-chip. In: DATE '06: Proceedings of the conference on Design, automation and test in Europe. pp. 9–14.
- [21] Guz, Z., Walter, I., Bolotin, E., Cidon, I., Ginosar, R., Kolodny, A., 2007. Network delays and link capacities in application-specific wormhole nocs. Special Issue of the Journal of VLSI Design.

- [22] Hary, S., Ozguner, F., 1997. Feasibility test for real-time communication using wormhole routing. *IEE Proceedings-Computers and Digital Techniques* 144 (5), 273–278.
- [23] Hu, J., Ogras, U., Marculescu, R., 2006. System-Level Buffer Allocation for Application-Specific Networks-on-Chip Router Design. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems* 25 (12), 2919.
- [24] Kiasari, A., Rahmati, D., Sarbazi-Azad, H., Hessabi, S., 2008. A Markovian Performance Model for Networks-on-Chip. In: *Parallel, Distributed and Network-Based Processing, 2008. PDP 2008. 16th Euromicro Conference on*. pp. 157–164.
- [25] Kleinrock, L., 1975. *Theory, Volume 1, Queueing Systems*. Wiley-Interscience.
- [26] Loucif, S., Ould Khaoua, M., Min, G., 2007. A queueing model for predicting message latency in uni-directional k-ary n-cubes with deterministic routing and non-uniform traffic. *Cluster Computing* 10 (2), 229–239.
- [27] Lu, Z., Jantsch, A., Sander, I., 2005. Feasibility analysis of messages for on-chip networks using wormhole routing. In: *Proceedings of the 2005 conference on Asia South Pacific design automation*. pp. 960–964.
- [28] Marchal, W., 1974. Some Simple Bounds and Approximations in Queueing.
- [29] Mello, A., Tedesco, L., Calazans, N., Moraes, F., 2005. Virtual channels in networks on chip: implementation and evaluation on hermes NoC. In: *Proceedings of the 18th annual symposium on Integrated circuits and system design*. pp. 178–183.
- [30] Moadeli, M., Shahrabi, A., Vanderbauwhede, W., Ould-Khaoua, M., 2007. An analytical performance model for the Spidergon NoC. In: *21st IEEE International Conference on Advanced Information Networking and Applications*. pp. 1014–1021.

- [31] Mohapatra, P., 1998. Wormhole routing techniques for directly connected multicomputer systems. *ACM Computing Surveys (CSUR)* 30 (3), 374–410.
- [32] Mullins, R., West, A., Moore, S., 2006. The design and implementation of a low-latency on-chip network. In: *Proceedings of the 2006 conference on Asia South Pacific design automation*. pp. 164–169.
- [33] Ni, L., McKinley, P., 1993. A Survey of Wormhole Routing Techniques in Direct Networks. *IEEE Computer* 26 (2), 62–76.
- [34] Ogras, U. Y., Marculescu, R., 2007. Analytical router modeling for networks-on-chip performance analysis. In: *DATE '07: Proceedings of the conference on Design, automation and test in Europe*. pp. 1096–1101.
- [35] Ould-Khaoua, M., Sarbazi-Azad, H., 2001. An analytical model of adaptive wormhole routing in hypercubes in the presence of hot spot traffic. *IEEE Transactions on Parallel and Distributed Systems* 12 (3), 283–292.
- [36] Parkes, S., Armbruster, P., 2005. SpaceWire: a spacecraft onboard network for real-time communications. In: *IEEE-NPSS Real time conf.* pp. 6–10.
- [37] Parkes, S., Rosello, J., 2001. SpaceWire- Links, nodes, routers and networks. In: *DASIA 2001- DATA Systems in Aerospace; Proceedings of the Conference*.
- [38] Perros, H., 1994. *Queueing Networks with Blocking*. Oxford University Press, Oxford Oxfordshire.
- [39] Pestana, S., Rijpkema, E., Radulescu, A., Goossens, K., Gangwal, O., 2004. Cost-performance trade-offs in networks on chip: A simulation-based approach. In: *Proc. DATE*. pp. 764–769.
- [40] Pirvu, M., Bhuyan, L., Ni, N., 1999. The impact of link arbitration on switch performance. In: *Proceedings of the Fifth Symposium on High-Performance Computer Architecture*.

- [41] Qian, Y., Lu, Z., Dou, W., 2009. Analysis of worst-case delay bounds for best-effort communication in wormhole networks on chip. In: Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip-Volume 00. IEEE Computer Society, pp. 44–53.
- [42] Qian, Y., Lu, Z., Dou, W., 2009. Applying network calculus for performance analysis of self-similar traffic in on-chip networks. In: Proceedings of the 7th IEEE/ACM international conference on Hardware/software codesign and system synthesis. ACM, pp. 453–460.
- [43] Rostislav, D., Vishnyakov, V., Friedman, E., Ginosar, R., 2005. An asynchronous router for multiple service levels networks on chip. In: Proceedings of the 11th IEEE International Symposium on Asynchronous Circuits and Systems. pp. 44–53.
- [44] Shi, Z., Burns, A., 2008. Real-Time Communication Analysis for On-Chip Networks with Wormhole Switching. In: Networks-on-Chip, 2008. NoCS 2008. Second ACM/IEEE International Symposium on. pp. 161–170.
- [45] Thakur, R., Choudhary, A., 1994. All-to-all communication on meshes with wormhole routing. In: Parallel Processing Symposium, 1994. Proceedings., Eighth International. pp. 561–565.
- [46] Uitert, M. J. G., 2003. Generalized processor sharing queues. Ponsen and Looijen BV.
- [47] Varga, A., et al., 2001. The OMNeT++ discrete event simulation system. In: Proceedings of the European Simulation Multiconference (ESM 2001).