Theoretical Analysis and Evaluation of NoCs with Weighted Round-Robin Arbitration

Sumit K. Mandal[1], Jie Tong[1], Raid Ayoub[2], Michael Kishinevsky[2], Ahmed Abousamra[2], Umit Y. Ogras[1]

[1]Dept. of ECE, University of Wisconsin-Madison; [2]Intel Corporation, Hillsboro, OR

Abstract— **Fast and accurate performance analysis techniques are essential in early design space exploration and pre-silicon evaluations, including software eco-system development. In particular, on-chip communication continues to play an increasingly important role as the many-core processors scale up. This paper presents the first performance analysis technique that targets networks-on-chip (NoCs) that employ weighted round-robin (WRR) arbitration. Besides fairness, WRR arbitration provides flexibility in allocating bandwidth proportionally to the importance of the traffic classes, unlike basic round-robin and priority-based arbitration. The proposed approach first estimates the *effective service* time of the packets in the queue due to WRR arbitration. Then, it uses the effective service time to compute the average waiting time of the packets. Next, we incorporate a decomposition technique to extend the analytical model to handle NoC of any size. The proposed approach achieves less than 5% error while executing real applications and 10% error under challenging synthetic traffic with different burstiness levels.**

I. INTRODUCTION

Networks-on-chip continue playing a central role as many-core processors start dominating the server [1, 2] and deep learning [3–5] market. As the commercial solutions scale up, the latency, area, and power consumption overheads of NoCs become increasingly crucial. Designers need analytical power-performance models to guide complex design decisions during the architecture development and implementation phases. After that, the same models are required by virtual platforms, commonly used to develop and evaluate the software ecosystem and applications [6]. Hence, there is a strong demand for high-fidelity analytical techniques that accurately model fundamental aspects of industrial designs across all segments ranging from systems-on-chip to client and server systems.

NoCs can be broadly classified in terms of buffer usage as buffered and bufferless architectures [7–10]. Most early solutions adapted buffered techniques, such as wormhole and virtual-channel switching, where the packets (or their flits) are stored in intermediate routers. Area, latency, and energy consumption of buffers have later led to bufferless architectures, where the intermediate routers forward the incoming flits if they can and deflect otherwise. Bufferless NoCs save significant buffer area and enable ultra-fast, as low as single-cycle routing decisions [7, 8]. Therefore, many industrial NoCs used in server and client architectures employ bufferless solutions to minimize the communication latency between the cores, last-level caches (LLC), and main memory [2, 11]. These solutions give priority to the packets already in the network to enable predictable and fast communication while stalling the newly generated packets from the processing and storage nodes. However, buffer area savings and low communication latency come at the cost of the early onset of congestion. Indeed, the packets wait longer at the end nodes, and the throughput saturates faster when the NoC load increases. Moreover, all routers in the NoCs remain powered on, increasing the NoC power consumption. Therefore, there is a strong need to address these shortcomings.

Buffered NoCs with virtual channel routers have been used more commonly in academic work and most recent industry standards [12, 13]. Shared buffering resources, such as input and output channels, require arbitrating among different requesters. For example, suppose that packets in different input channels request the same output channel. An arbiter needs to resolve the conflicts and grant access to one of the requesters to meet performance target. The architectures proposed to date predominantly employ basic round-robin (RR) arbiter to provide fairness to all requesters [14–16]. Although the decisions are locally fair, the number of arbitrations a packet goes through grows with its path length. Hence, RR arbitration is globally unfair. More importantly, basic RR cannot provide preference to a particular input, which is typically desired since not all requests are equal. For example, data and acknowledgment packets can have higher priority than new requests to complete outstanding transactions, especially when the network is congested.

WRR arbitration provides flexibility in allocating bandwidth proportionally to the importance of the traffic classes, unlike basic round-robin and priority-based arbitration. Each requester has an assigned weight, which is a measure of its importance. A larger weight indicates that the requester is given more preference in arbitration. Due to its generality, WRR arbitration has been employed in several NoC proposals in the literature [17–19]. Indeed, WRR arbitration enables higher throughput than RR arbitration [18]. *Despite its potential, WRR arbitration has not been analyzed theoretically, especially for large-scale NoCs.* A large body of literature has proposed performance analysis techniques for buffered and bufferless NoCs since analytical models play a crucial role in fast design space exploration and pre-silicon evaluation [20–24]. In contrast, no analytical modeling technique has been proposed to date for NoCs with WRR arbitration. A formal analysis is required to understand the behavior of NoCs with WRR arbitration. At the same time, executable performance models are needed to guide many-core processor design and enable virtual platforms for pre-silicon evaluation.

This paper presents a fast, accurate, and scalable perfor-

mance analysis technique for NoCs with WRR arbitration. To the best of our knowledge, it is the first performance analysis technique for NoCs with weighted round-robin arbitration. Furthermore, the proposed technique supports bursty core traffic observed in real applications, which is typically ignored due to its complexity. It first estimates the effective service time of the packets in the queue due to WRR arbitration. Then, it utilizes the effective service time to obtain the average waiting time of the packets. We also propose a decomposition technique to extend the analytical model for any size of NoC. Extensive experimental evaluations show that the proposed analytical model has less than 5% with real applications and 10% error with synthetic traffic having different burstiness levels congesting the NoC.

The major contributions of the work are listed below:

- A novel performance analysis technique for NoCs that employ WRR arbitration,
- A decomposition technique to obtain a scalable analytical model for NoC of any size.
- Experimental evaluations with multiple NoC configurations with different traffic scenarios showing less than 5% error for real applications.

II. RELATED WORK AND NOVEL CONTRIBUTIONS

Analytical models are required to estimate the NoC performance for fast design space exploration and pre-silicon evaluation. Multiple prior studies have proposed NoC performance analysis techniques with basic round-robin arbitration. [21, 25, 26]. Authors in [25] first construct a contention matrix between multiple flows in the NoC. Then, the average waiting time of the packets corresponding to each flow is computed. Support vector regression-based analytical model for NoCs is proposed in [21]. The analytical model proposed in [26] estimates the mean service time of the flows with RR arbitration. The estimated mean service time is used to find the average waiting time of the flows. However, none of these techniques are applicable in the presence of both bursty traffic and WRR arbitration.

Analytical modeling of round-robin arbitration has also been studied outside NoC domain [20, 27, 28]. The techniques presented in [27, 28] incorporate a polling model to approximate the effective service time of a queue in the presence of RR arbitration. However, none of these approaches are applicable when the input distribution to the queue is not geometric. A Markov chain-based analytical model is proposed in [20] to account for bursty input traffic. However, the technique is not scalable for a network of queues. Moreover, none of these techniques are applicable for discrete-time queuing systems. Since each transaction in NoC happens at discrete clock cycles, the analytical models need to incorporate discrete-time queuing systems. The major drawbacks of the prior approaches are summarized in Table I.

The basic round-robin arbitration cannot provide fairness when requesters have widely varying data rate requirements and priorities. Therefore, weighted round-robin arbitration,

TABLE I
COMPARISON OF PRIOR RESEARCH AND OUR NOVEL CONTRIBUTION.

Research	Approach	WRR	Bursty Traffic	Scalable	Discrete Time
Boxma et al. [27]	Polling model	No	No	Yes	No
Wim et al. [28]	Extended polling model	No	No	Yes	No
Wang et al. [20]	Markov chain	No	Yes	No	No
Fischer et al. [26]	Heuristic	No	No	Yes	No
Vanlerbergee et al. [29]	Moment generating function	No	No	No	Yes
This work	**Queue decomposition**	**Yes**	**Yes**	**Yes**	**Yes**

i.e., WRR, has been used in on-chip communication architectures [17, 18]. Qian et al. compute delay bounds for different channels with different weights to assign appropriate weight to each input channel of the NoC [17]. They show that WRR delivers better quality of service than NoCs with strict priority-based arbitration. Authors in [18] propose a WRR-based scheduling policy. The proposed technique assigns larger bandwidth to input channels with higher weights. It achieves higher throughput compared to round-robin arbitration. *Although WRR has shown promise, no analytical modeling approach exists for NoCs with WRR to date.*

This paper presents the first performance analysis technique for NoCs with WRR arbitration. It fills an essential gap since WRR can address the shortcomings of priority-based bufferless NoC architectures and the basic round-robin arbitration. Furthermore, the proposed technique supports bursty traffic observed in real applications, which is typically ignored due to its complexity. Hence, it is a vital step towards comprehending the theoretical underpinnings of NoCs with WRR arbitration and enabling their deployment in industrial designs.

III. BACKGROUND AND OVERVIEW

A. Weighted Round-Robin Arbitration

This work uses weighted round-robin arbitration in the NoC routers. The basic operating principle of WRR arbitration is illustrated in Figure 1 for three traffic classes. Packets from each class are first written to a dedicated input queue (also known as a channel). Suppose there are N input queues $Q_1, Q_2, \ldots, Q_N$. WRR technique assigns Q_i a positive integer weight denoted as $\omega_i \in \mathbb{Z}^+$ for $1 \leq i \leq N$. The WRR arbiter serves up to ω_i consecutive packets from Q_i before moving to the next queue. If Q_i has less than ω_i packets, then WRR serves Q_i until it becomes empty. Then, the WRR arbiter serves the subsequent queues following the same principle. After a cycle is completed, the weights are reset to their initial values, as illustrated in Figure 1, and the same arbitration cycle is repeated. In NoCs with WRR arbitration, whenever two


Fig. 1. Illustration of weighted round-robin arbitration. 'A' is served first, 'I' last. In this example it is assumed that no new packets arrived until all prior packets (A-I) have been served.

or more requesters compete for the same resource, they are arbitrated following WRR. WRR arbitration can be used both for arbitrating different virtual channels and different ports in the network.

B. Usage of the Proposed Performance Analysis Technique

WRR arbitration is promising for NoCs since it can tailor the communication bandwidth to different traffic classes. Furthermore, it provides end-to-end latency-fairness to different source-destination pairs, unlike basic round-robin and priority arbitration techniques. However, these capabilities come at the expense of a vast design parameter space. An $n \times m$ mesh with P-port routers has $n \times m \times P$ tunable weights, e.g., an 8×8 2D mesh with 5-port routers would have 320 WRR weights. Due to this ample design space, the current practice is limited to assigning two weights to each router (e.g., one weight to local ports and another weight to packets already in the NoC).

The benefits of the proposed theoretical analysis are two-fold. *First*, it can enable accurate pre-silicon evaluations and design space exploration without time-consuming cycle-accurate simulations. *Second*, it can be used to find the combination of weights that optimizes the performance, i.e., to solve the optimization problem in the vast design space described in the previous paragraph. *This paper focuses on constructing the proposed analysis technique and its evaluation against cycle-accurate simulation for two reasons.* First, the theoretical analysis is complex, and its evaluation deserves a dedicated treatment on its own. Moreover, its application for solving optimization problems requires demonstration of its fidelity first. Multi-objective optimization of the WRR weights using the proposed model is one of our future research directions.

IV. PROPOSED METHODOLOGY AND APPROACH

Figure 2 shows a weighted round-robin arbiter with N queues. The packets that belong to traffic class-i are stored in queue Q_i. The corresponding average arrival rate of class-i packets is denoted by λ_i, as summarized in Table II. The burstiness of traffic class-i can be captured by their squared coefficient of variation of inter-arrival time, denoted by C_{Ai} [17, 30]. Finally, the weight assigned to class-i is denoted as $\omega_i \in \mathbb{Z}^+$ for $1 \leq i \leq N$, as illustrated in Figure 2.

To provide a step-by-step derivation, Section IV-A starts with a particular case of the proposed technique tailored to the basic round-robin arbitration, i.e., all weights are set to one. Then, Section IV-B extends the formulation to weighted round-robin arbitration.

Fig. 2. The original WRR arbiter with its traffic parameters (on the left) and a transformed WRR (on the right) that comprises fully decomposed queue nodes using our effective service time transformations.

TABLE II
LIST OF THE IMPORTANT PARAMETERS USED IN THIS WORK.

N	Number of traffic classes to an arbiter
λ_i	Injection rate of class-i
ω_i	Weight assigned to class-i
T	Original mean service time of all traffic classes
$\widehat{T}_i$	Mean value of effective service time of class-i
$\rho_i, \widehat{\rho}_i$	($\rho_i = \lambda_i T_i$) and ($\widehat{\rho}_i = \lambda_i \widehat{T}_i$) server utilizations of class-i
C_{Ai}	Squared coeff. of variation of inter-arrival time of class-i
C_S	Squared coeff. of variation of orig. service time (all classes)
$\widehat{C}_{Si}$	Squared coeff. of variation of eff. service time of class-i
C_A	Sq. coeff. of variation of inter-arrival time (merged traffic)
C_D	Squared coeff. of variation of inter-departure time
n_i	Mean queue occupancy of class-i
$p(n_i > k|\widehat{T}_j)$	Probability that queue occupancy of class-$i > k$ given a mean effective service time of class-$j = \widehat{T}_j$
W_i	Average waiting time of class-i
$\widehat{R}_i$	Mean effective residual service time of class-i

A. Analytical Model for **Basic** Round-Robin Arbitration

In the basic round-robin arbitration, all weights are equal to one, i.e., $\omega_i = 1$, $1 \leq i \leq N$. Each traffic class experiences an additional delay until the arbiter serves the head packets in the other queues. The proposed technique has two steps:

1) The first step is to compute the first two moments of the *effective service time* for each class-i ($\widehat{T}_i$, $\widehat{C}_{Si}$) of the fully decomposed queue nodes illustrated in Figure 2.
2) In the second step, we use the transformed effective service times ($\widehat{T}_i$, $\widehat{C}_{Si}$, $1 \leq i \leq N$) to find the total waiting time (including the queuing delay) of each traffic class.

For the clarity of notation and illustration, the derivation below assumes the original two moments of service time (T, C_S) are the same across all traffic classes.

*1) **Mean effective service time of RR** ($\widehat{T}_i$):* The effective service time accounts for the delay experienced by the packets at the head of each queue. This delay includes its own service time and the service time of the packet at the head of other queues since the round-robin arbiter serves them one by one. When all queues have packets, class-i packets will be served every $N \times T$ cycles, i.e., they will wait for $(N-1) \times T$ cycles after being served before winning the arbitration again. In general, a traffic class will only contribute to the extra service time if it has a packet waiting for service. Thus, assuming packet arrival of different classes are independent, we can express the mean effective service time as:

$$\widehat{T}_i = T + \sum_{j=1, j \neq i}^{N} p(n_i > 0|\widehat{T}_i)p(n_j > 0|\widehat{T}_i)T$$

$$= T + Tp(n_i > 0|\widehat{T}_i)\sum_{j=1, j \neq i}^{N} p(n_j > 0|\widehat{T}_i), 1 \leq i \leq N \quad (1)$$

where $p(n_j > 0|\widehat{T}_i)$ denotes the probability that the occupancy of class-j is greater than zero given a mean effective service period of class-i is $\widehat{T}_i$, for $1 \leq i, j \leq N$. We approximate $p(n_j > 0|\widehat{T}_i)$ using Little's law as $\lambda_j \widehat{T}_i$ following the busy-cycle approach [27]. Little's law gives the expected number of class-j packets that arrive during the period of $\widehat{T}_i$, which

Algorithm 1: Basic round-robin arbitration: Effective service time ($\widehat{T}_i$) computation.

---
1 Input: Injection rate of each class (λ), service time (T), number of classes (N), Convergence $Tolerance$
2 Output: Effective service time of each class ($\widehat{T}$)
3 for $i = 1{:}N$ **do**
4   Find smaller root of the quadratic equation derived from Equation 2:
5 $T + T\lambda_i(\widehat{T}_i)^2 \sum_{j=1, j\neq i}^{N} \lambda_j - \widehat{T}_i = 0$
6 $\delta = Tolerance, k = 1$
7   **while** $\delta \geq Tolerance$ **do**
8 $\widehat{T}_i^{k+1} \leftarrow T + Tmin(1, \lambda_i\widehat{T}_i^k) \sum_{\substack{j=1 \\ j\neq i}}^{N} min(1, \lambda_j\widehat{T}_i^k)$
9 $\delta = \widehat{T}_i^{k+1} - \widehat{T}_i^k$
10 $k \leftarrow k + 1$
11   **end**
12 end

---

approximates the probability that the occupancy of class-j packets is greater than zero during this period. To ensure that the probability term does not exceed one, we approximate $p(n_j > 0|\widehat{T}_i)$ as $min(1, \lambda_j\widehat{T}_i)$. Therefore, Equation 1 can be rewritten as:

$$\widehat{T}_i = T + Tmin(1, \lambda_i\widehat{T}_i) \sum_{\substack{j=1 \\ j\neq i}}^{N} min(1, \lambda_j\widehat{T}_i), \ 1 \leq i \leq N \quad (2)$$

We solve Equation 2 through a simple iterative approach described in Algorithm 1. First, Equation 2 is solved by removing the nonlinearity introduced by the min operation (lines 4 and 5). We consider the smaller solution out of the two solutions, since we observed through experiments that the smaller solution is more accurate. The solution is used as the initial estimate of the iterative approach. Then, this estimate is plugged into Equation 2 within an iteration loop to obtain a better estimate (line 7). The iterations continue until the change in the effective service time is within a user-provided $Tolerance$ value, set to 0.01 in this work. Algorithm 1 typically takes less than ten iterations due to the quadratic nature of Equation 2, which takes only a few microseconds to compute. After the iterations are completed, Algorithm 1 returns the effective service time $\widehat{T}_i$.

2) Coefficient of variation of effective service time for RR ($\widehat{C}_{Si}$): This section derives $\widehat{C}_{Si}$ by leveraging the conservation of work principle and property of equal residual service time of individual classes [27]. We leverage these properties to compute the mean effective residual service time of class-i ($\widehat{R}_i$). Then, $\widehat{R}_i$ is used to calculate $\widehat{C}_{Si}$ using the relation found in [30]:

$$\widehat{C}_{Si} = \frac{1}{\widehat{\rho}_i}\left(\frac{2\widehat{R}_i}{\widehat{T}_i} + 1 - C_{Ai} - \widehat{\rho}_i\right) \quad (3)$$

where $\widehat{\rho}_i = \lambda_i\widehat{T}_i$ is the effective server utilization. In addition, we show that leveraging the conservation of work principle gives our model the capability of handling bursty traffic.

We exploit the *conservation of work principle*, stating that the same amount of packets are served regardless of how the bandwidth is divided between the traffic flows assuming

all packets have same moments of service time. Due to this fundamental principle, the total queue occupancy n_{sum} is invariant under the arbitration policy. Furthermore, this sum can be computed with near-perfect accuracy using maximum entropy-based analytical queuing model of multi-class given in [30]. One strong aspect of this model is that it can take bursty traffic that conforms to General Geometric distribution (GGeo). By leveraging this model, we can embed the capability of handling bursty traffic into our analytical model. The total queue occupancy (n_{sum}) can be calculated as:

$$n_{sum} = \frac{1}{2}\sum_{i=1}^{N}\left(\rho_i(C_{Ai} - 1) + \frac{\sum_{k=1}^{N}\frac{\lambda_i}{\lambda_k}\rho_k^2(C_{Ak} + C_S)}{1 - \sum_{k=1}^{N}\rho_k}\right) \quad (4)$$

using the parameters summarized in Table II. We do not delve into the derivation of this equation for clarity since it is not our major focus. The detailed derivation can be found in the seminal work by Kouvatsos et al. [30].

Next, we derive n_{sum} using the first two moments of transformed effective service time. The occupancy of each queue Q_i can be expressed using Little's law as: $n_i = \lambda_i W_i$, where W_i is the average waiting time of class-i. Therefore, we obtain:

$$n_{sum} = \sum_{i=1}^{N}\lambda_i W_i \quad (5)$$

The average waiting time W_i consists of two components. The first one is the waiting time in the queue till the packet reaches the head of the queue ($\widehat{W}_i$), i.e. until all other class-i packets leave the server completely. The second component accounts for the waiting time at the head of the queue when the server is busy serving other classes (class-j, $j \neq i$). This component is the additional time captured by the effective service time $\widehat{T}_i$, i.e., the difference between the effective and original service times: $\Delta T_i = \widehat{T}_i - T$. Furthermore, the waiting $\widehat{W}_i$ is expressed as a function of the residual time [31]. Hence,

$$W_i = \frac{\widehat{R}_i}{1 - \lambda_i\widehat{T}_i} + \Delta T_i, \ for \ 1 \leq i \leq N \quad (6)$$

We leverage the property that all individual mean effective residual service time of classes are equal [27]; hence we write ($R = \widehat{R}_i, 1 \leq i \leq N$). Thus, we can rewrite the total occupancy by plugging Equation 6 into Equation 5 and then solve for the residual time R:

$$n_{sum} = \sum_{i=1}^{N}\lambda_i\left(\frac{R}{1 - \lambda_i\widehat{T}_i} + \Delta T_i\right)$$

$$\implies R = \left(n_{sum} - \sum_{i=1}^{N}\lambda_i\Delta T_i\right)\left(\sum_{i=1}^{N}\frac{\lambda_i}{1 - \lambda_i\widehat{T}_i}\right)^{-1} \quad (7)$$

Then, we can compute the coefficient of variation $\widehat{C}_{Si}$ by substituting the residual time R from Equation 7 in Equation 3.

3) Average waiting time of RR (W_i): Finally, we compute the mean waiting time of individual classes by plugging $\widehat{T}_i$ and $\widehat{C}_{Si}$ found for RR arbitration into Equation 3 and Equation 6 as:

$$W_i = \frac{0.5\widehat{T_i}(\widehat{\rho_i} - 1 + C_{Ai} + \widehat{\rho_i}\widehat{C}_{Si})}{1 - \lambda_i\widehat{T_i}} + \Delta T_i, \ 1 \le i \le N \quad (8)$$

B. Analytical Model for **Weighted** Round-Robin Arbitration

This section extends the analytical model constructed in Section IV-A to WRR using the same steps. It first describes the derivation of the first two moments of effective service time. Then, it will use these moments to find the average waiting time.

*1) **Mean effective service time of WRR** ($\widehat{T_i}$):* In WRR, the weight of each class can be larger than or equal to one, i.e., $\omega_i \ge 1, \ 1 \le i \le N$. To illustrate the generalization, we start with a system of two queues, where $\omega_1 > 1$ and $\omega_2 = 1$. Then, we generalize it to N queues with arbitrary weights.

*Mean effective service time of **class**-1 with $\omega_1 > 1$ and $\omega_2 = 1$:* Since $\omega_1 > 1$, a batch of up to ω_1 number of class-1 packets are served without interruption. Let $\widehat{T_i^b}$ be the total service time of ω_1 class-1 packets. When there is a class-2 packet in Q_2, the next batch of ω_1 class-1 packet will wait. Therefore, we can express $\widehat{T_1^b}$ as:

$$\widehat{T_1^b} = \omega_1 T + p(n_1 > 0|\widehat{T_1^b})p(n_2 > 0|\widehat{T_1^b})T \quad (9)$$

Note that this equation generalizes Equation 1 to arbitrary ω_1 for two queues. Using Little's law and the methodology described in Section IV-A, we approximate $\widehat{T_1^b}$ as:

$$\widehat{T_1^b} = \omega_1 T + \frac{T}{\omega_1}min(1, \lambda_1\widehat{T_1^b})min(1, \lambda_2\widehat{T_1^b}) \quad (10)$$

*Mean effective service time of **class**-2 with $\omega_1 > 1$ and $\omega_2 = 1$:* The batch size of class-2 packet is one in this case since $\omega_2 = 1$. They need to wait in the queue for a maximum of ω_1 class-1 packets. Therefore, $\widehat{T_2}$ is expressed as:

$$\widehat{T_2} = T + \sum_{i=0}^{\omega_1 - 1} p(n_1 > i|\widehat{T_2})p(n_2 > 0|\widehat{T_2})T \quad (11)$$

We can approximate $p(n_1 > i|\widehat{T_2})$ as $\frac{1}{i+1}p(n_1 > 0|\widehat{T_2})$ to obtain:

$$\widehat{T_2} = T + \sum_{i=0}^{\omega_1 - 1} \frac{1}{i+1}p(n_1 > 0|\widehat{T_2})p(n_2 > 0|\widehat{T_2})T$$

$$= T + Tmin\left(1, \sum_{i=0}^{\omega_1 - 1} \frac{1}{i+1}\lambda_1\widehat{T_2}\right)min(1, \lambda_2\widehat{T_2}) \quad (12)$$

To get to the final form of Equation 12 above, we approximate the conditional probabilities based on [27].

Generalization to N queues and arbitrary weights: In general, there are N classes with weights $\omega_i, 1 \le i \le N$. Consider the class-i packets with the total service time $\widehat{T_i^b}$ for a batch of ω_i packets. The first part of the effective service time will be due to the *packets of same class*. This effect is captured in our two-queue illustration for class-1 packets in Equation 9. The second part of the effective service time will be due to the *packets of other classes*. This effect is captured in our two-queue illustration for class-2 packets in Equation 11. By combining them, we can express the generalized effective service time as:

$$\widehat{T_i^b} = \omega_i T + p(n_i > 0|\widehat{T_i^b}) \sum_{j=1,j\neq i}^{N} \left(\sum_{k=0}^{\omega_j - 1} p(n_j > k|\widehat{T_i^b}) \right)T \quad (13)$$

Algorithm 2: Obtaining effective service time ($\widehat{T_i}$) of weighted round-robin

1 **Input:** Injection rate of each class (λ), WRR weights (ω_i), service time (T), number of classes (N)
2 **Output:** Effective service time of each class ($\widehat{T}$)
3 **for** $i = 1:N$ **do**
4   Find smaller root of the quadratic equation derived from Equation 14 for $\widehat{T_i^b}$:
5     $(\widehat{T_i^b})^2 \frac{T}{\omega_i}\lambda_i \sum_{\substack{j=1 \\ j\neq i}}^{N} \sum_{k=0}^{\omega_j - 1} \frac{1}{k+1}\lambda_j - \widehat{T_i^b} + \omega_i T = 0$
6     $\delta = Tolerance, k = 1$
7   **while** $\delta \ge Tolerance$ **do**
8       $Z = \sum_{\substack{j=1 \\ j\neq i}}^{N} min\left(1, \sum_{k=0}^{\omega_j - 1} \frac{1}{k+1}\lambda_j\widehat{T_i^{b,k}}\right)$
9       $\widehat{T_i^{b,k+1}} \leftarrow \omega_i T + \frac{T}{\omega_i}min(1, \lambda_i\widehat{T_i^{b,k}})Z$
10      $\delta = \widehat{T_i^{b,k+1}} - \widehat{T_i^{b,k}}$
11      $k \leftarrow k + 1$
12   **end**
13     $\widehat{T_i} = \frac{\widehat{T_i^b}}{\omega_i}$
14 **end**

Note that this equation reduces to Equation 9 for $i = 1, N = 2, \omega_2 = 1$, and to Equation 11 for $i = 2, N = 2, \omega_2 = 1$.

Finally, we can obtain the generalized effective service times by replacing $p(n_i > 0)$ with $min(1, \lambda_i\widehat{T_i^b})$ as follows:

$$\widehat{T_i^b} = \omega_i T + \frac{T}{\omega_i}min(1, \lambda_i\widehat{T_i^b}) \sum_{\substack{j=1 \\ j\neq i}}^{N} min\left(1, \sum_{k=0}^{\omega_j - 1} \frac{1}{k+1}\lambda_j\widehat{T_i^b}\right)$$

$$\widehat{T_i} = \frac{\widehat{T_i^b}}{\omega_i} \quad (14)$$

Due to the non-linearity introduced by the min operation, we compute $\widehat{T_i}$ using the iterative Algorithm 2 just like the basic round-robin case. The quadratic nature of Equation 14 enables a fast convergence within 10 iterations under $5\mu s$.

*2) **Coefficient of variation of effective service time of WRR** ($\widehat{C}_{Si}$) [1]:* In the basic round-robin case, the residual service time R_i accounts for the second-order moment of the service time, i.e., C_{Si}. However, effective mean residual service time of different classes are not necessarily equal to each other for arbitrary weights, which is the case in WRR arbitration. The mean effective residual service time of class-i packets is expressed as:

$$\widehat{R_i} = 0.5\widehat{T_i}(\widehat{\rho_i} - 1 + C_{Ai} + \widehat{\rho_i}\widehat{C}_{Si}) \quad (15)$$

where $\widehat{T_i}$ is found from Algorithm 2 based on Equation 14. Note that the residual service time depends on the first two moments of service time, which are the average service time $\widehat{T_i}$ and the coefficient of variation $\widehat{C}_{Si}$, respectively. To approximate $\widehat{C}_{Si}$ for WRR, we leverage its monotonically decreasing behavior as a function of the corresponding weight. This behavior is expected since a larger weight corresponds to fewer arbitration stages, decreasing the effective service time variation. This observation indicates that the coefficient

---

[1] We use the same notation for RR and WRR arbitration since WRR results are generalizations of basic RR expressions.

of variation for round-robin arbitration is an upper bound of coefficient of variation for weighted round-robin arbitration. This upper bound, $\widehat{C}_{Si}(RR)$ for class-i, is computed using Equation 3. Given this, we first apply linear approximation to calculate the $\widehat{C}_{Si}$ under WRR as $\alpha\frac{\widehat{C}_{Si}(\text{RR})}{\omega_i^2}$, note we use the squared value of ω_i since $\widehat{C}_{Si}$ is the square of coefficient of variation as defined in Table II. We add the parameter α to tighten the approximation. Next, we recall Equation 4 and replace C_{Sk} by $\alpha\frac{\widehat{C}_{Sk}(RR)}{\omega_k^2}$ to obtain:

$$n_{sum}=\frac{1}{2}\sum_{i=1}^{N}\left(\widehat{\rho}_i(C_{Ai}-1)+\frac{\sum_{k=1}^{N}(\frac{\lambda_i}{\lambda_k})\widehat{\rho}_k^2(C_{Ak}+\alpha\frac{\widehat{C}_{Sk}(RR)}{\omega_k^2})}{1-\sum_{k=1}^{N}\widehat{\rho}_k}\right) \tag{16}$$

The left hand side of the Equation 16 (n_{sum}) is obtained by applying Equation 4. Therefore, α is the only unknown in Equation 16. Once we obtain α (by solving Equation 16), we compute $\widehat{C}_{Si}$ under WRR as $\alpha\frac{\widehat{C}_{Si}(\text{RR})}{\omega_i^2}$. This expression captures traffic burstiness addressing one of the significant drawbacks of the prior work.

*3) **Average waiting time of WRR** (W_i):* So far, we obtained first two moments of effective service time, i.e., $\widehat{T}_i$ (Equation 14) and $\widehat{C}_{Si}$. We obtain the average waiting time under WRR arbitration by plugging them to Equation 17:

$$W_i = \frac{0.5\widehat{T}_i(\widehat{\rho}_i-1+C_{Ai}+\widehat{\rho}_i\widehat{C}_{Si})}{1-\lambda_i\widehat{T}_i}+\Delta T_i \tag{17}$$

C. Estimation of end-to-end latency

The previous section described the canonical model for WRR arbitration, where all traffic classes go through a single arbiter. The packets in NoCs go through a sequence of WRR arbiters while traveling from their sources towards destinations. For example, an 8×8 2D mesh has 64 routers with at least one arbiter per output queue in each of them. Figure 3(a) illustrates an example with two stages, where the flow that wins the first stage is routed to the second one. The major challenge in modeling multiple stages is obtaining the inter-departure distribution, which becomes the arrival flow at the subsequent stages. For instance, the coefficient of variation of arrival flow 2 in the second stage is equal to the coefficient of variation of the departing flow from stage-1 (C_D) in Figure 3(a). We first find the squared co-efficient of variation of inter-departure time for each traffic class (C_{Di}):

$$C_{Di}=\rho_i^2(\widehat{C}_{Si}+1)+(1-\rho_i)C_{Ai}+\rho_i(1-2\rho_i),\forall 1\leq i\leq N \tag{18}$$

Next, we apply the decomposition technique [32] to obtain the departure distribution at each stage (C_D).

Algorithm 3: Estimation of end-to-end latency

1 **Input:** NoC size ($n \times m$), injection rates and co-efficient of variation of inter-arrival times for all classes (λ, C_A), Service time (T), WRR weights
2 **Output:** Average NoC ($\bar{L}$) and source–dest. latencies (L_{sd})
3 Number of routers $N_p \leftarrow n \times m$, $\lambda_{sum} \leftarrow 0, L_{sum} \leftarrow 0$
4 **for** $s = 1{:}N_p$ **do**
5   **for** $d = 1{:}N_p$ **do**
6         $W_{sd} = compute\_waiting\_time(s,d)$
7         $L_{sd} = W_{sd} + free\_packet\_delay(s,d)$
8         $L_{sum} \leftarrow L_{sum} + L_{sd}, \lambda_{sum} \leftarrow \lambda_{sum} + \lambda_{sd}$
9   **end**
10 **end**
11 $\bar{L} = \frac{L_{sum}}{\lambda_{sum}}$
12 **function** $compute\_waiting\_time(s,d)$
13     $S \leftarrow$ The number of stages for the flow with source s and destination d, $W_{sd} \leftarrow 0$
14   **for** $j = 1{:}S$ **do**
15         $W_j =$ Waiting time at stage j using Equation 17
16         $W_{sd} \leftarrow W_{sd} + W_j$
17         $C_D =$ Coefficient of variation using Equation 19
18   Use C_D as the C_A in the next stage
19   **end**
20   **return** W_{sd}
21 **end**

$$C_D = \frac{\sum_{i=1}^{N}\lambda_i C_{Di}}{\sum_{i=1}^{N}\lambda_i} \tag{19}$$

The departure distribution (C_D) is the arrival distribution to the next stage (C_{A2}). For a given source-destination pair, these calculations are performed at each arbitration stage on the path of the packets.

Algorithm 3 presents a step-by-step procedure to obtain end-to-end latency of a given NoC with WRR arbitration and deterministic routing. The input to the algorithm is the NoC size, injection rate, and coefficient of variation of inter-arrival time of all classes, service time of the queues, and weights assigned to each arbiter. The output of the algorithm is the latency of each source-destination pair (L_{sd}) as well as the average latency $\bar{L}$.

Algorithm 3 first computes the waiting time of each source-destination pair (W_{sd}). The function $compute\_waiting\_time$ in lines 12–22 describes the procedure for this computation. The procedure starts with finding the number of stages (S) for the flow from source s to destination d. At each stage, it computes the average waiting time using Equation 17 (line 16). Next, it adds the waiting time at the current state to the cumulative waiting time found so far (line 17). Then, it computes the coefficient of variation of inter-departure time


Fig. 3. Illustration of extending the proposed analysis to multiple stages. Only two consecutive stages are shown for clarity. The departure statistics at a given stage become the arrival statistics of the subsequent stage. The blue line denotes that the class which wins the arbitration goes to the next stage.

Fig. 4. Verification of the analytical model for basic round-robin with (a) 8×1 and (b) 8×8 NoC.

(C_D) using Equation 18 and C_D is used as the coefficient of variation of inter-arrival time in the next stage (lines 18–19). Finally, the procedure returns the total waiting time from source s to destination d (W_{sd}) (line 21). After retrieving the total waiting times, Algorithm 3 adds the free packet delay (t_{sd}), i.e., the latency due to the links and router micro-architecture, which can be found using the number of hops and the router pipeline depth [10, 33]. Then, end-to-end latencies and flow rates are accumulated (line 8). Finally, the average latency of the network is obtained by computing the weighted average of the latency of all source-destination pairs (line 11). We note that different WRR arbiters in the network can be assigned different arbitration weights.

V. Experimental Results

A. Experimental Setup

This section presents a thorough evaluation of the proposed analytical modeling approach under various traffic scenarios and NoC configurations. We employ geometric and bursty traffic distributions in addition to real application traces feeding 1×8, 6×6, and 8×8 NoCs (note that the typical size of state-of-the-art industrial NoCs is 6×6 [2, 11]). The synthetic traffic simulations use 100% last-level cache (LLC) hit and 100% LLC miss scenarios, commonly used corner cases. Furthermore, applications from SPEC CPU 2017 [34] and PARSEC [35] benchmark suits demonstrate the effectiveness of the proposed technique in a broader range of scenarios. The latency results of our analytical models are compared against an in-house cycle-accurate simulator, which is calibrated with an industrial simulator [36]. All simulations run for 200K cycles to reach a steady state and have 20K cycles of warm-up. The source code of the simulator and the analytical model are publicly released at [37].

B. Results with Basic Round-Robin Arbitration

We first evaluate the proposed analytical model under basic round-robin arbitration. Figure 4 depicts representative average end-to-end latency comparisons between analytical model and the cycle-accurate simulations for 8×1 ring and 8×8 mesh NoC. The average injection rates for both NoC configurations


Fig. 5. Verification of the analytical model for weighted round-robin with 8 × 8 NoC. [x y] denotes that the WRR weights associated with channels connected to the internal routers is x and traffic channels of external input to the NoC is y.

are varied until the NoC becomes highly congested, following geometric distribution. On average, the proposed technique incurs only 5% error for 8×1 ring and 7% error for 8×8 mesh. We also compare the proposed analytical model with a polling-based model for round-robin arbitration [28]. The polling-based model grossly overestimates the latency both for 8×1 and 8×8 NoC. A comprehensive summary for other NoC sizes and traffic patterns is presented in Section V-E. These results demonstrate that the proposed approach is accurate for difference NoC, traffic, and WRR parameters.

C. Results with Weighted Round-Robin Arbitration

The most important aspect of the proposed technique is considering WRR arbitration. Figure 5 shows the average end-to-end latency comparison of our analytical model against cycle-accurate simulations with an 8×8 mesh for three different weight configurations. The proposed analytical model incurs on average 8% error when the WRR weights associated with the traffic channels of external input to the NoC and channels connected to internal routers are 1 and 2, respectively. The average error slightly increases to 9% when the arbitration weights of the packets in the NoC increase to 3. This configuration decreases the network congestion by providing higher priority to the packets already in the NoC. Hence, the average waiting time decreases. Since there are no other analysis approaches for WRR, this section does not provide comparisons to state-of-the-art.

D. Results with Bursty Traffic and Real Applications

This section evaluates the proposed model in the presence of bursty traffic. We define probability of burstiness as p_{burst} of the general geometric distribution [30]. Increasing p_{burst} denotes increasing burstiness in the traffic. We note that no other prior work considered weighted round-robin arbitration together with bursty traffic. Figure 6 compares the end-to-end average latency results against cycle-accurate simulations for 6×6 and 8×8 mesh when the arbitration weights associated with the packets in the NoC and external packets arrivals are 3 and 1, respectively. When the burst probability is $p_{burst} = 0.1$ (Figure 6(a)), the average modeling errors for 6×6 and 8×8 mesh are 7% and 4%, respectively. When the burst probability increases to 0.3 (Figure 6(b)), the average errors become 6% and 5%, respectively.

Finally, we evaluate the proposed analysis technique while running applications from the SPEC CPU2017 [34] and PARSEC benchmark [35]. The applications from SPEC CPU2017 benchmarks show burstiness in the range of $p_{burst} = 0.1-0.6$.

Fig. 6. Verification of analytical model for bursty traffic with (a) $p_{burst} = 0.1$ and $p_{burst} = 0.3$.

Fig. 7. Verification against real applications.

The applications from PARSEC benchmark are 16-threaded and do not show any burstiness. Figure 7 compares the analysis results against cycle-accurate simulations on an $8{\times}8$ NoC. The WRR weights of packets already in the NoC are set to 3, while new packets from cores have a lower priority set with a weight of 1. We observe that our proposed analysis technique consistently achieves a modeling error of less than 5% for these applications. These results demonstrate that the proposed analysis technique can accurately model the end-to-end NoC latency under WRR arbitration and bursty traffic.

E. Summary of the Evaluation Results

This section summarizes the accuracy of our performance analysis technique systematically for different NoC sizes, WRR weights, and traffic burstiness. To capture the most commonly observed corner cases, we provide results for two extreme cases – 100% LLC hit and 100% LLC miss traffic. In both cases, cores send packets to each LLC at an equal rate.

Table III summarizes the comparisons between analysis and simulations for 100% LLC hit. When $p_{burst} = 0$ and the traffic load is moderate ($\lambda = 0.1$), the modeling error ranges from 1.4% to 5.5% considering both RR and WRR arbitration. Even when the traffic load increases congesting the NoC ($\lambda = 0.3$), the modeling error remains below 11% for all configurations. A higher level of burstiness to $p_{burst} = 0.3$ increases the NoC load. Consequently, the error range increases 3.6%–9.0% for the moderate traffic load ($\lambda = 0.1$). Increasing the load further congests the NoC ($\lambda = 0.3$) pushing the worst-case error only to 13%, which is acceptable since practical systems hardly operate this load due to congestion control mechanisms.

Finally, Table IV summarizes the modeling error of the proposed technique for 100% LLC miss. In this case, the missed core requests are forwarded to memory controllers (one in $8{\times}1$ NoC and two in $6{\times}6$ and $8{\times}8$ NoCs). Hence, the memory controllers become hotspots. Despite the high

degree of skewness, the error for moderate NoC loads ranges from 4.3% to 8.1% for all burstiness levels, WRR weight configurations, and NoC sizes. Even when the NoC is pushed to congestion, the modeling error remains under 9.0% for all configurations, except for $8{\times}8$ NoC with WRR weights 3 and 1. The error in this specific case is 13%; the error for 8×1 and 6×6 NoCs are 4.6% and 5.4%, respectively.

F. Execution Time of the Proposed Analysis Technique

We implemented the proposed technique in C++ to obtain end-to-end latency. The computational complexity of the proposed model is $O(n^2 \times m^2 \times p)$, for $n \times m$ NoC ($n \geq m$) and WRR arbiters with maximum p ports. The execution time as a function of NoC sizes with 3 output ports per router are summarized in Table V. We observe that even for $32{\times}32$ NoC, the analytical model takes only 45.07 ms to execute, while the simulation time of the 8×8 NoC is in the order of minutes. These results show that the proposed analysis technique is lightweight and provides four orders of magnitude speed-up compared to cycle-accurate simulations.

TABLE V
ANALYSIS ON EXECUTION TIME OF THE PROPOSED MODEL.

NoC size	4×4	6×6	8×8	16×16	32×32
Exe. time (ms)	0.56	1.41	2.25	13.52	45.07

VI. CONCLUSION AND FUTURE WORK

This paper presented a scalable performance analysis technique for NoCs with WRR arbitration. WRR arbitration is promising since it enables latency fairness and custom bandwidth allocation to meet the requirements of individual traffic flows, unlike basic round-robin and priority-based arbitrations. The proposed technique handles bursty traffic and provides high accuracy for both ring topologies used in client systems and large 2D mesh topologies used in servers. Extensive evaluations show the proposed approach achieves less than 5% error while executing real applications and 10% error under challenging synthetic traffic with different burstiness levels. Hence, it can be used for fast and accurate design space exploration as well as pre-silicon evaluations. One of our future directions is finding the WRR weights to meet individual traffic flows' bandwidth and latency requirements.

TABLE III
SUMMARY OF RESULTS FOR SYNTHETIC APPLICATIONS WITH **100% HIT**.

Topo.	$\omega_{ring} = 1, \omega_{src} = 1$				$\omega_{ring} = 3, \omega_{src} = 1$			
	$p_{burst} = 0$		$p_{burst} = 0.3$		$p_{burst} = 0$		$p_{burst} = 0.3$	
	$\lambda = 0.1$	$\lambda = 0.3$	$\lambda = 0.1$	$\lambda = 0.3$	$\lambda = 0.1$	$\lambda = 0.3$	$\lambda = 0.1$	$\lambda = 0.3$
$8{\times}1$	1.4%	11%	3.6%	7.8%	1.5%	13%	9%	11%
$6{\times}6$	5.5%	7.2%	7.4%	11%	5.2%	11%	5.9%	12%
$8{\times}8$	2.6%	7.8%	5.2%	10%	3.5%	11%	4.8%	7.2%

TABLE IV
SUMMARY OF RESULTS FOR SYNTHETIC APPLICATIONS WITH **100% MISS**.

Topo.	$\omega_{ring} = 1, \omega_{src} = 1$				$\omega_{ring} = 3, \omega_{src} = 1$			
	$p_{burst} = 0$		$p_{burst} = 0.3$		$p_{burst} = 0$		$p_{burst} = 0.3$	
	$\lambda = 0.1$	$\lambda = 0.2$	$\lambda = 0.1$	$\lambda = 0.2$	$\lambda = 0.1$	$\lambda = 0.2$	$\lambda = 0.1$	$\lambda = 0.2$
$8{\times}1$	4.3%	4.6%	5.3%	5.7%	4.4%	4.6%	4.5%	4.8%
$6{\times}6$	5.0%	5.4%	5.5%	6.0%	5.0%	5.4%	5.0%	5.8%
$8{\times}8$	7.1%	8.0%	7.4%	7.6%	7.4%	13%	8.1%	9.0%

REFERENCES

[1] Mohamed Arafa et al. Cascade Lake: Next Generation Intel Xeon Scalable Processor. *IEEE Micro*, 39(2):29–36, 2019.

[2] Jack Doweck et al. Inside 6th-generation Intel Core: New Microarchitecture Code-named Skylake. *IEEE Micro*, 37(2):52–62, 2017.

[3] Hanbo Sun et al. Reliability-Aware Training and Performance Modeling for Processing-In-Memory Systems. In *26th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 847–852. IEEE, 2021.

[4] Sumit K Mandal et al. A Latency-Optimized Reconfigurable NoC for In-Memory Acceleration of DNNs. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 10(3):362–375, 2020.

[5] Seyed Morteza Nabavinejad et al. An overview of Efficient Interconnection Networks for Deep Neural Network Accelerators. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 10(3):268–282, 2020.

[6] Ming-Chao Chiang, Tse-Chen Yeh, and Guo-Fu Tseng. A QEMU and SystemC-based Cycle-accurate ISS for Performance Estimation on SoC Development. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(4):593–606, 2011.

[7] Thomas Moscibroda and Onur Mutlu. A Case for Bufferless Routing in on-chip Networks. In *Proc. of Intl. Symp. on Computer Architecture*, pages 196–207, 2009.

[8] Bhavya K Daya, Li-shiuan Peh, and Anantha P Chandrakasan. Quest For High-Performance Bufferless Nocs With Single-Cycle Express Paths And Self-Learning Throttling. In *Proc. of Design Automation Conf.*, pages 1–6, 2016.

[9] Radu Marculescu et al. Outstanding Research Problems In Noc Design: System, Microarchitecture, And Circuit Perspectives. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Syst.*, 28(1):3–21, 2008.

[10] William James Dally and Brian Patrick Towles. *Principles and Practices of Interconnection Networks*. Elsevier, 2004.

[11] Avinash Sodani et al. Knights Landing: Second-generation Intel Xeon Phi Product. *Ieee micro*, 36(2):34–46, 2016.

[12] Umit Y Ogras and Radu Marculescu. *Modeling, Analysis And Optimization Of Network-On-Chip Communication Architectures*, volume 184. Springer Science & Business Media, 2013.

[13] Andrea Pellegrini et al. The Arm Neoverse N1 Platform: Building Blocks for the Next-gen Cloud-to-Edge Infrastructure SoC. *IEEE Micro*, 40(2):53–62, 2020.

[14] Eung S Shin, Vincent J Mooney III, and George F Riley. Round-robin Arbiter Design and Generation. In *Proc. of Intl. Symp. on System Synthesis*, pages 243–248, 2002.

[15] Yun-Lung Lee, Jer Min Jou, and Yen-Yu Chen. A High-speed and Decentralized Arbiter Design for NoC. In *Proc. of Intl. Conf. on Comp. Systs. and Applications*, pages 350–353, 2009.

[16] Gao Xiaopeng, Zhang Zhe, and Long Xiang. Round Robin Arbiters for Virtual Channel Router. In *Proc. of Multiconference on Computational Engineering in Systems Applications*, volume 2, pages 1610–1614, 2006.

[17] Yue Qian, Zhonghai Lu, and Qiang Dou. Qos Scheduling for NoCs: Strict Priority Queueing versus Weighted Round Robin. In *Proc. of Intl. Conf. on Computer Design*, pages 52–59, 2010.

[18] Jan Heißwolf, Ralf König, and Jürgen Becker. A Scalable NoC Router Design Providing QoS Support using Weighted Round Robin Scheduling. In *Proc. of Intl. Symp. on Parallel and Distributed Processing with Applications*, pages 625–632, 2012.

[19] Cesar Albenes Zeferino and Altamiro Amadeu Susin. SoCIN: A Parametric and Scalable Network-on-Chip. In *Symp. on Integrated Circuits and Systems Design, 2003. SBCCI 2003. Proceedings.*, pages 169–174, 2003.

[20] Lan Wang, Geyong Min, Demetres D Kouvatsos, and Xiaolong Jin. Analytical Modeling of an Integrated Priority and WFQ Scheduling Scheme in Multi-service Networks. *Computer Communications*, 33:S93–S101, 2010.

[21] Zhi-Liang Qian et al. A Support Vector Regression (SVR)-based Latency Model for Network-on-Chip (NoC) Architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(3):471–484, 2015.

[22] Abbas Eslami Kiasari, Zhonghai Lu, and Axel Jantsch. An Analytical Latency Model for Networks-on-Chip. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 21(1):113–123, 2012.

[23] Sumit K Mandal et al. Analytical Performance Modeling of NoCs under Priority Arbitration and Bursty Traffic. *IEEE Embedded Systems Letters*, 2020.

[24] Sumit K Mandal et al. Performance Analysis of Priority-aware NoCs with Deflection Routing under Traffic Congestion. In *Proc. of Intl. Conf. on Computer-Aided Design*, pages 1–9, 2020.

[25] Umit Y Ogras, Paul Bogdan, and Radu Marculescu. An Analytical Approach for Network-on-Chip Performance Analysis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 29(12):2001–2013, 2010.

[26] Erik Fischer and Gerhard P Fettweis. An Accurate and Scalable Analytic Model for Round-robin Arbitration in Network-on-Chip. In *Proc. of Intl. Symp. on Networks-on-Chip (NoCS)*, pages 1–8, 2013.

[27] Onno J. Boxma and Bernd W. Meister. Waiting-time Approximations in Multi-queue Systems with Cyclic Service. *Performance Evaluation*, 7(1):59–70, 1987.

[28] Wim P Groenendijk and Hanoch Levy. Performance Analysis of Transaction Driven Computer Systems via Queueing Analysis of Polling Models. *IEEE Computer Architecture Letters*, 41(04):455–466, 1992.

[29] Jasper Vanlerberghe. *Analysis and Optimization of Discrete-time Generalized Processor Sharing Queues*. PhD thesis, Ghent University, 2018.

[30] Demetres D Kouvatsos. Entropy Maximisation and Queueing Network Models. *Annals of Operations Research*, 48(1):63–126, 1994.

[31] Dimitri P Bertsekas, Robert G Gallager, and Pierre Humblet. *Data Networks*, volume 2. Prentice-Hall International New Jersey, 1992.

[32] Guy Pujolle and Wu Ai. A Solution for Multiserver and Multiclass Open Queueing Networks. *INFOR: Information Systems and Operational Research*, 24(3):221–230, 1986.

[33] Umit Y Ogras and Radu Marculescu. "It's a Small World After All": Noc Performance Optimization Via Long-Range Link Insertion. *IEEE Trans. on Very Large Scale Integration (VLSI) Systs.*, 14(7):693–706, 2006.

[34] James Bucek, Klaus-Dieter Lange, and Jóakim v. Kistowski. SPEC CPU2017: Next-Generation Compute Benchmark. In *Companion of the 2018 ACM/SPEC Intl. Conf. on Performance Engineering*, pages 41–42, 2018.

[35] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In *Proc. of Intl. Conf. on Parallel Architectures and Compilation Techniques*, pages 72–81, 2008.

[36] Umit Y Ogras, Yunus Emre, Jianping Xu, Timothy Kam, and Michael Kishinevsky. Energy-guided Exploration of On-chip Network Design for Exa-scale Computing. In *Proc. of Intl. Workshop on System Level Interconnect Prediction*, pages 24–31, 2012.

[37] Sumit K. Mandal et al. Theoretical Analysis and Evaluation of NoCs with Weighted Round-Robin Arbitration: An Open Source Release. https://github.com/sumitkmandal/WRR_NoC_analytical_model, 2021. Accessed 12 August 2021.