

# Approximating Age-Based Arbitration in On-Chip Networks

Michael M. Lee,  
John Kim  
KAIST  
Daejeon, Korea  
{minjong7,jjk12}@kaist.edu

Dennis Abts,  
Michael Marty  
Google Inc.  
Madison, WI  
{dabts,mikemarty}@google.com

Jae W. Lee  
Parakinetics Inc.  
Princeton, NJ  
leejw@parakinetics.com

## ABSTRACT

The on-chip network of emerging many-core CMPs enables the sharing of numerous on-chip components. This on-chip network needs to ensure fairness when accessing the shared resources. In this work, we propose providing equality of service (EoS) in future many-core CMPs on-chip networks by leveraging distance, or hop count, to approximate the age of packets in the network. We propose *probabilistic* arbitration combined with distance-based weights to achieve EoS and overcome the limitation of conventional round-robin arbiter. We describe how nonlinear weights need to be used with probabilistic arbiters and propose three different arbitration weight metrics – fixed weight, constantly increasing weight, and variably increasing weight. By only modifying the arbitration of an on-chip router, we do not require any additional buffers or virtual channels and create a complexity-effective mechanism for achieving EoS.

## Categories and Subject Descriptors

C.1.2 [Computer Systems Organization]: Multiprocessors—*Interconnection architectures*

## General Terms

Algorithms, Design, Performance

## Keywords

on-chip network, age-based arbitration, fairness

## 1. INTRODUCTION

The on-chip network in many-core chip multiprocessors is crucial to providing equality of service among the shared resources; providing consistent latency and bandwidth characteristics for each processor. Applications should be insensitive to *where*, within the on-chip network, the thread is scheduled to execute. Achieving this goal will reduce the variance in the execution time among threads and will provide more efficient synchronization when transitioning between a parallel and sequential code regions. Stated simply as to provide “*equality of service* (EoS)” among the threads competing for shared network resources. Recently, cost-efficient quality of service (QoS) for on-chip networks has been proposed [2, 3]. Unlike QoS, which strives to provide differentiated service and guarantees for end-to-end latency or bandwidth profile, EoS does not provide guarantees yet provides equal access to shared on-chip resources.

Copyright is held by the author/owner(s).  
PACT’10, September 11–15, 2010, Vienna, Austria.  
ACM 978-1-4503-0178-7/10/09.




Figure 1: 8-ary 1-mesh example where all nodes are sending to P7 and merging traffic at each hop.

A motivation example of EoS in on-chip network is shown in Figure 1 with an 8-ary 1-mesh with processors P0 thru P6 sending traffic to P7. Newly injected packets merge with traffic from other directions and this merging of traffic causes each node to have unfair throughput. With a round-robin switch arbitration policy, the processor closest to the destination (P6 is only one hop away) will get the most bandwidth —  $1/2$  of the available bandwidth. The processor two hops away, P5, will get half of the bandwidth into router R6, for a total of  $1/2 \times 1/2 = 1/4$  of the available bandwidth. Thus, P0 and P1 each receive only  $1/64$  of the available bandwidth into P7. That is, bandwidth of each processor decreases exponentially as its distance increases. Reducing the variation in bandwidth is critical for application performance, particularly as applications are scaled to increasingly higher processor counts.

Age-based arbitration [1] is known to provide global fairness, but there are some disadvantages that it needs microarchitecture modification, additional packet overhead, increased arbitration complexity to consider the age metric as an input to arbitration, and updating/maintaining the age field which will require a new packet checksum to be computed on each hop.

In this work, we avoid the complexity with age-based arbitration by proposing to approximate the age of a packet with *distance* or hop count. By using information already present in the packet, such as source node, current node, or destination node and using distance as a proxy for the packet’s age, age-based arbitration is greatly simplified. We introduce *distance-based* arbitration to allow nodes located many hops from the edge to get equal service compared to a node close to the edge. We propose using *probabilistic* arbitration with a distance-based selection algorithm to achieve EoS while providing a low complexity, livelock-free arbitration allowing for consistent latency and bandwidth characteristics for all cores. Since nodes that are farther away are serviced at a ratio that is geometrically proportional to the hop count, we propose using *nonlinear* weights in probabilistic arbitration to provide fairness to nodes that are farther

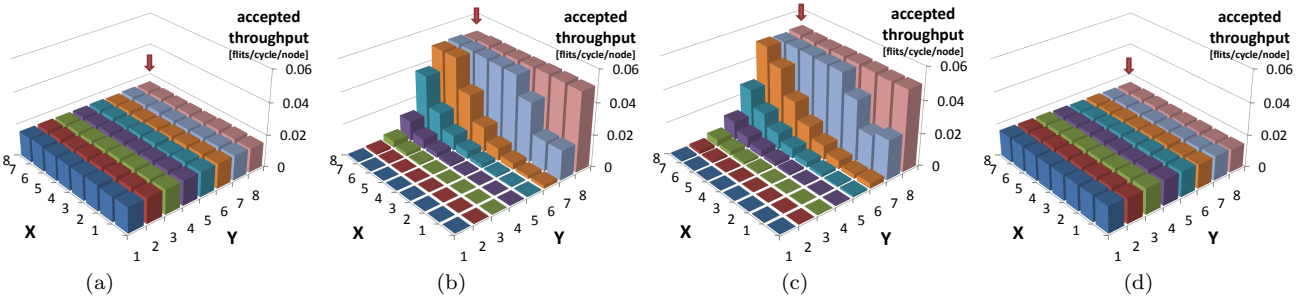


Figure 2: Results of hotspot traffic using (a) age-based arbitration, (b) round-robin arbitration, (c) probabilistic arbitration with linear weights, and (d) probabilistic arbitration with non-linear weights.

away. Three different arbitration weight metrics are proposed which all provide EoS but have varying trade-off in terms of complexity and performance degradation on different traffic patterns.

## 2. ARBITRATION DESIGN

Previous arbitration architectures are *deterministic*—that is, for priority-based arbitration such as age-based arbitration, the arbitration is made deterministically based on the relative age of the requests. Starvation is inherently not a problem with age-based arbitration. However, by using priority based on hop count, livelock and starvation issues are problematic because a packet with a lower priority (i.e., a lower hop count) can continually lose arbitration because of a constant injected traffic with a higher priority. To overcome this issue, while still using hop count as the weight, we propose *probabilistic* arbitration where the output of the arbitration is probabilistically determined based on the weight of the input requests.

Assume an arbiter with two request  $r_1$  and  $r_2$ , each with a corresponding weight  $w_1$  and  $w_2$ . The probability of each grant  $g_1$  and  $g_2$  being asserted with probabilistic arbitration is equal to the following.

$$P(g_1) = \frac{w_1}{w_1 + w_2}$$

$$P(g_2) = \frac{w_2}{w_1 + w_2}$$

In general, for a request  $r_i$  into an arbiter with  $m$  requests, the probability of  $r_i$  being granted is

$$P(g_i) = \frac{w_i}{\sum_{j=1}^m w_j}$$

The hop count can be used as weights for probabilistic arbiter. For example, *linear* weights can be implemented – i.e.,  $w = h_x$  or  $w = h_x + h_y$  where  $h_x$  and  $h_y$  represent the hop count from source to destination in each dimension. However, as shown in Figure 2(c), probabilistic arbitration using linear weight hop count cannot provide fairness since farther nodes will be serviced linearly instead of exponentially.

As shown earlier in Figure 1, packets that are  $h$  hops away are serviced at the rate of  $(1/2)^h$ . To provide fairness in this case, we introduce *nonlinear* weights based on the distance – i.e.,  $w = C^h$  where  $C$  is the *contention degree* or the number of flows contending for the same output port. For example, in Figure 1, each node of  $h$  hops away needs to be prioritized with a weight of  $w = 2^h$ . That is,  $C = 2$  value is used since for each output, there are two contending flows. With XY routing in hotspot traffic, packets traveling in the

Table 1: Arbitration metrics to determine weight of probabilistic arbitration where  $h$  is the hop count and  $C$  is the contention degree.

$h$	$C$	description
static	static	fixed weight (FW)
dynamic	static	constantly increasing weight (CW)
static	dynamic	N/A
dynamic	dynamic	variably increasing weight (VW)

$x$ -dimension will merge similar to the traffic in Figure 1. When traversing the  $y$ -dimension, there are 3 or 4 traffic flows merging at each router, resulting in each flow being serviced at a rate of  $1/3$  or  $1/4$ . Thus, weight used for  $y$ -dimension is  $w = C^h$  where  $C = 3, 4$ .

The hop count weight used can be categorized as either *static* or *dynamic*. With a static arbitration metric, the weight of the packet is known beforehand at the time the packet is injected into the network. On the other hand, dynamic metrics will cause the weight of a packet to change en route. Leveraging the nonlinear weight ( $C^h$ ), the different metric can be categorized based on whether  $C$  and  $h$  are either static or dynamic as summarized in Table 1.

For the fixed weight (FW), the weight of each packet remains constant or *fixed* throughout the network. However, for both constantly increasing weight (CW) and variably increasing weight (VW), all packets are injected with a weight of 1 and the weights are increased appropriately. With CW, the increase in the weight is based on the distance a packet has traversed while with VW, the increase in the weight is proportional to the contention encountered at each router.

Thus, by modifying only the arbitration aspect of an on-chip network through probabilistic arbitration with nonlinear weights, age-based arbitration can be approximated to provide equality-of-service (EoS) in many-core CMP on-chip network.

## 3. REFERENCES

- [1] D. Abts et al. Age-based packet arbitration in large-radix k-ary n-cubes. In *SC '07*, pages 1–11, New York, NY, USA, 2007. ACM.
- [2] B. Grot et al. Preemptive virtual clock: A flexible, efficient, and cost-effective qos scheme for networks-on-a-chip. In *MICRO 42*, New York, NY, 2009.
- [3] J. W. Lee et al. Globally-synchronized frames for guaranteed quality-of-service in on-chip networks. *SIGARCH Comput. Archit. News*, 36(3):89–100, 2008.