An Efficient Path Setup for a Hybrid Photonic Network-on-Chip

Cisse Ahmadou Dit ADI

The University of Electro-Communications,
Graduate School of Information Systems
IS-635, 1-5-1 Chofugaoka, Chofu-shi
Tokyo, 182-8585, Japan

Hiroki Matsutani

The University of Tokyo,
7-3-1, Hongo,Bunkyo-ku
Tokyo, 113-8656, Japan

Michihiro Koibuchi

National Institute of Informatics,
2-1-2, Hitotsubashi, Chiyoda-ku
Tokyo, 101-8430, Japan

Hidetsugu Irie

The University of Electro-Communications,
Graduate School of Information Systems
IS-635, 1-5-1 Chofugaoka, Chofu-shi
Tokyo, 182-8585, Japan

Takefumi Miyoshi

The University of Electro-Communications,
Graduate School of Information Systems
IS-635, 1-5-1 Chofugaoka, Chofu-shi
Tokyo, 182-8585, Japan

and Tsutomu Yoshinaga

The University of Electro-Communications,
Graduate School of Information Systems
IS-635, 1-5-1 Chofugaoka, Chofu-shi
Tokyo, 182-8585, Japan

Abstract

Electrical network-on-chip (NoC) faces critical challenges in meeting the high performance
and low power consumption requirements for future multicore processors interconnection. Re-

cent tremendous advances in CMOS compatible optical components give the potential for photonics to deliver an efficient NoC performance at an acceptable energy cost. However, the lack of in flight processing and buffering of optical data made the realization of a fully optical NoC complicated. A hybrid architecture which uses optical high bandwidth transfer and an electrical control network can take advantage of both interconnection methods to offer an efficient performance-per-watt infrastructure to connect multicore processors and system-on-chip (SoC). In this paper, we propose a predictive switching and a reservation based path setup techniques to reduce the path setup latency of such hybrid photonic network-on-chip (HPNoC). By using these techniques, it is possible to reduce the latency for end-to-end communication in a HPNoC improving its overall performance. In the simulation, we use a cycle accurate simulator under uniform, neighbor, and bitreversal traffic patterns for a 64-node torus topology. The results show that the proposed techniques considerably improve the overall latency of HPNoC.

Keywords: nanophotonics, photonic NoC, predictive switching, multicore processors.

1 Introduction

Transistor size is continuously shrinking down, leading to better chip integration capabilities. According to the international technology roadmap for semiconductors (ITRS), hundreds of cores can be integrated in a single chip in near future. Therefore, the communication infrastructure should be improved to deal with the enormous increases in complexity, energy consumption, and bandwidth demand. Today's electrical network-on-chips (NoCs), which consume a huge amount of power for electrical signaling, face critical challenges to provide the required communication performance within the available power budget. These limitations direct current research activities on finding alternative approaches with better energy efficiency for multicore processors and system-on-chip (SoC) interconnections.

Applying optical technology for data communication inside a chip has attracted considerable attention as an efficient alternative for the electrical interconnection due to its possibility to provide high-bandwidth at a lower power cost. With recent development of CMOS compatible nanophotonic technology and 3D stack, the integration of optical components on a chip becomes more realistic. However, optical interconnection on chip still faces a number of difficulties. Data buffering and in-flight[1] signal processing are not currently viable at the chip level [10]. Despite these limitations, several researches that take advantage of the unique low power and high bandwidth abilities provided by optical components have been proposed [1, 2, 4, 8, 9, 12].

Ventrease et al. proposed a clustered architecture that uses optical arbitration via a wavelength-routed token ring to reserve access to a fully optical crossbar made from waveguides, modulators, and detectors. Pan et al. proposed a clustered architecture using a dragonfly topology in which nodes in the same cluster are connected using a conventional electrical interconnect and node from different cluster are connected using an optical crossbar. Hendry et al. proposed a time division multiplexing (TDM) to remove the need of electrical path setup network by scheduling statically photonic transmission in TDM slot [4].

To take full advantage of a wavelength division multiplexing (WDM) technique, which allows the parallel transmission of many optical signals on different wavelengths on a single waveguide, Shacham et al. have proposed a circuit switching hybrid photonic network-on-chip (HPNoC) [10]. The architecture consists of a photonic layer, which uses a high-bandwidth circuit switching, controlled by an electrical packet switching layer. The HPNoC removes the need for buffering of optical data and the high power consumption of optical-electrical-optical (O-E-O) conversions at intermediate node for routing computation. With the combination of the optical circuit-switching network and electrical packet-switching network, the HPNoC provides a better interconnection bandwidth and transmission speed at a lower power consumption in comparison with the all-electrical NoC [9].

In this paper, we propose predictive switching [6] and a reservation based path setup techniques for the electrical control network to reduce the setup latency of a HPNoC. Since the circuit setup latency plays a key role in the overall performance of HPNoC [11], we use these techniques to

---

[1]Routing computation for packet switching normally requires buffering of the packet header.

reduce the path setup latency. In the simulation, we use a cycle accurate simulator under uniform, neighbor, and bitreversal traffic patterns for a 64-node torus topology. The obtained results show that the predictive switching and reservation based path setup techniques considerably reduce the $end-to-end-latency$ for communication, leading to a better performance HPNoC.

The paper is organized as follows. In Section 2, we describe our proposed HPNoC based on predictive switching and reservation based path setup. In Section 3, we first estimate power consumption cost of an all-electrical and a hybrid photonic architecture, and then present simulation results using a cycle accurate simulator. Finally, we conclude in Section 4.

2 Hybrid Photonic NoC
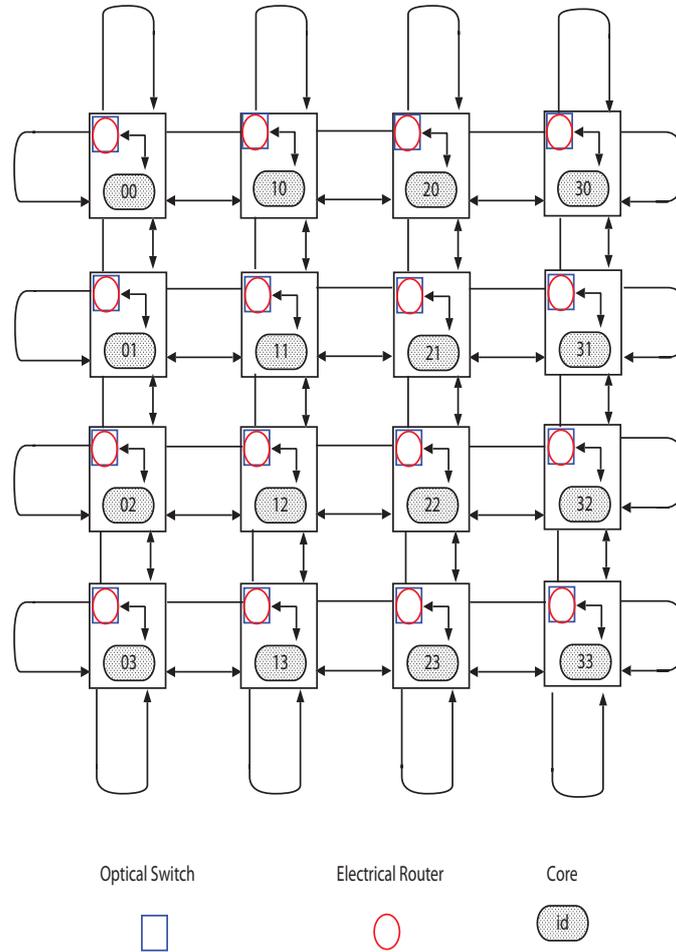
Figure 1: A 4×4 Hybrid Torus Photonic NoC.

While the hybrid photonic NoC offers unique advantages in terms of bandwidth and energy compared to fully electrical NoC, its implementation requires extra hardware to support the optical communication such as : light source (laser), modulators, waveguides, optical switches, and demodulators [11]. Fig. 1 shows a 4×4 torus HPNoC. The topology consists of 2 layers: an optical high-bandwidth data transfer circuit switching network, and an electrical packet switching control network. Nodes in the HPNoC communicate as follows:

- Firstly a path setup message is sent by the source node in the electrical network to establish a path for the optical network.

- After the path is set, an acknowledgment pulse is sent back to the source node by the destination node in the optical network, and optical data can be transfered without need for buffering at intermediate nodes.

- Finally when all data are sent, a teardown message is sent by the source node in the electrical control network to release the optical circuit.

Similarly to a circuit switching flow control, the HPNoC performs better with larger message sizes because of the high speed data transfer in the optical network once the communication path is established. When only a few small-sized data transmissions occur, the HPNoC is not needed, while a cheap simple electrical NoC fits with such a case.

2.1   Optical Network

The optical network comprises optical switches connected by optical waveguides. At each node, an optical modulator and detector are needed for electrical-optical-electrical (E-O-E) conversions. At the source node, an external laser light is modulated in the optical modulator from electrical to optical data signal. The modulated optical signal is transmitted on the optical waveguides. At the destination node, the optical signal is detected by the optical detector and ejected from the optical network. To build a $2D$ torus topology, a 5×5 optical switch is necessary for each node: one input/output port for each direction (WEST, NORTH, EAST, and SOUTH) and one for the processing element. To remove the need for extra injection and ejection gateways in the switch used in [10], we use the optical switch proposed in [4] shown in Fig. 2. The switch consists of micro-ring resonators, waveguides and a control unit. By turning ON/OFF the state of a resonator, light can be directed in the switch from one direction to another according to the control unit which is set by the electrical network. For instance in Fig. 2(a), optical data coming from the GATEWAY port is guided to the WEST output port by turning "ON" the resonator 4. The same data can be guided to the EST port by turning "ON" the resonator 2 shown in Fig. 2(b).
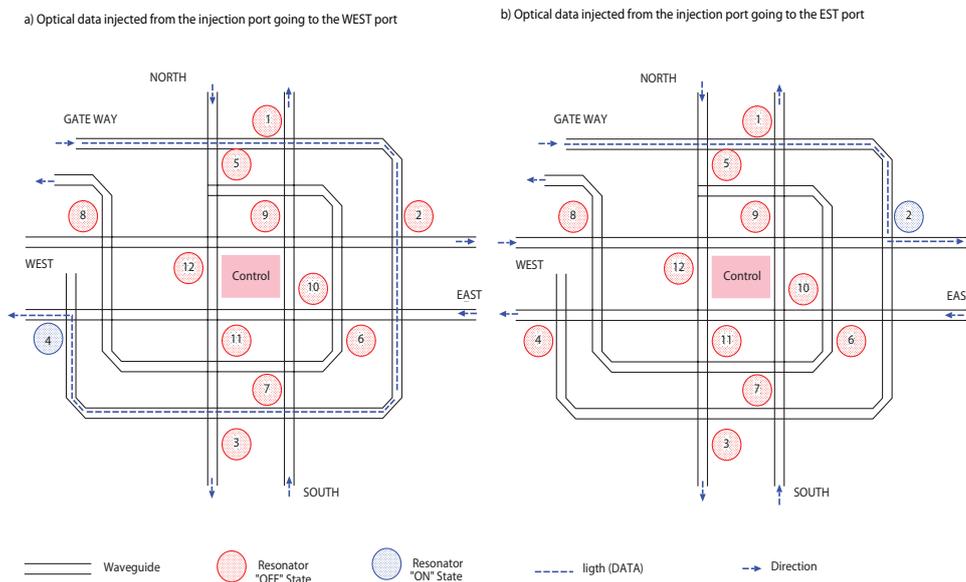
Figure 2: Optical Switch [4]

The high bandwidth capabilities of optical interconnects are due to the use of WDM. It statically allows the transfer of optical data using all wavelengths within a waveguide for the same source-destination pair's data stream. Optical switch with a smaller number of micro-ring resonator presents a better solution for hardware cost. The optical switch we used only required 12 micro-ring resonators. To implement a dynamic allocation(wavelengths of the same waveguide is divided among
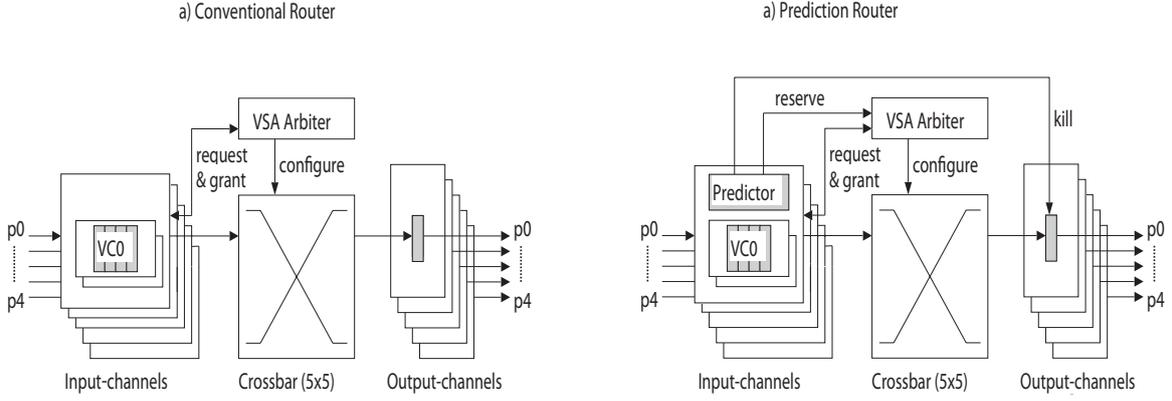
Figure 3: Electrical Routers

multiple data stream), however, the cost of the optical switch increases. The number of resonators is multiplied with the corresponding number of wavelengths used as each micro-ring resonator uses a unique resonance wavelength. The arrangement of the waveguides and micro-ring resonators made this optical switch suitable for mesh and torus networks that use dimension order routing (DOR). It removes unnecessary turns that are avoided in DOR.

2.2 Electrical Network

The electrical control network consists of electrical routers interconnected by electrical wires in a torus topology. We propose two path setup techniques to improve the performance of the control network by reducing the electrical network latency.

2.2.1 Predictive switching based path setup

For the predictive switching based path setup, we use prediction routers. The hardware area of the electrical network is increased by 4.8-12.0% as reported in our previous work [7]. Prediction routers speculatively forward the packets inside a router bypassing some pipeline stages. The prediction router is shown in Fig. 3(b). The differences from the conventional router shown in Fig. 3(a) are as follows:
1) A predictor is added at each input channel.
2) The arbitration unit for virtual-channel and switch allocations (VSA Arbiter) is modified to handle the tentative reservation from predictors.
3) And a kill signal is added at each output channel in order to remove miss-routed flits when the prediction fails [6].

The predictor in an input-channel forecasts which output channel will be used by the next packet transfer before it reaches the input-channel. Then it asserts the reserve signal to the arbiter in order to tentatively reserve a time-slot of the crossbar for the predicted output-channel. The VSA arbiter handles the request and reserve signals from each input-channel(configure). If the prediction fails, the kill signal is asserted to the miss-predicted output channel. The output-channel will mask all incoming data as dead flits (miss-routed flits) which never propagate to the outside of the router. With this technique, when the prediction hits, it is possible to complete the switch traversal (ST) within one router cycle and bypass the pipeline stages of routing computation (RC), virtual-channel allocation (VA), and switch allocation (SA) which are required in the conventional router [3]. When the prediction fails, the conventional packet processing is carried out. It is important to note here that there is no miss-penalty on the miss-routed latency.

Fig. 4, as an example, compares a timing diagram for sending a packet through 3 hops using a conventional router (Fig. 4(a)) and the prediction router for the electrical control network (Fig.

4(b)). With the prediction router, the $end-to-end-$ $latency$ is reduced by half from 12 router cycles, necessary in the conventional router, to only 6 cycles in the case of the predictions hit in two of the 3 hops.

By processing packets before they arrive at input buffers using look-ahead routing, only a single stage pipeline (ST) is necessary for packet transfer when prediction hits. The prediction mechanism, therefore, drastically reduces the packet processing latency per router. If a switching with high prediction hit rate is applied to the electrical control network of the HPNoC, it is possible to decrease the circuit setup latency and improves its overall performance.
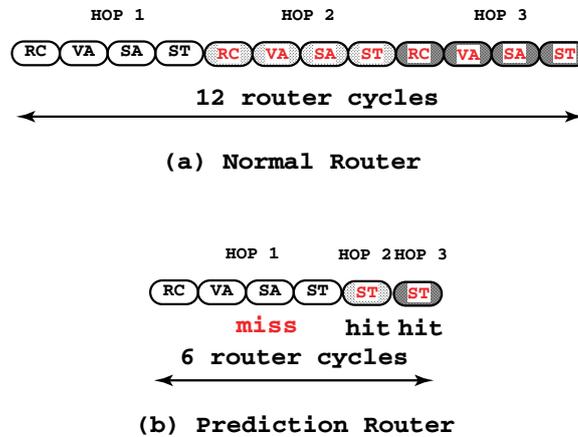
Figure 4: pipeline time diagram for normal and prediction router

Since some pipeline stages are skipped only when the prediction hits, the primary concern for reducing the communication latency is the prediction accuracy. We use the following two prediction algorithms.

- *Latest port matching (LP):* The LP strategy predicts in such a way that the next incoming packet will be forwarded to the same output-channel as that of the previous packet. The LP predictor requires only a single history record in each input-channel, leading to a lower hardware overhead cost.

- *Sampled pattern matching (SPM):* The SPM algorithm was originally proposed as a universal predictor [5]. It selects a value with the highest probability after a suffix sequence, called a marker, in a given data set. The predicted value is calculated by using the majority rule to all values appearing at positions just after the markers in the data. We can use it to predict an output-channel for the next incoming message of an input-channel by finding the most frequently used output-channel after the longest suffix sequence (marker) of the communication history. An example of prediction using the SPM prediction mechanism is shown in Fig. 5. In step 1 of the algorithm, the marker is determined by finding the longest repeated sequence from the history of past used output-channels used by an input-channel, in this example the marker is "0012" . Second, the values appearing at positions just after the markers in the history are recorded and counted (Step 2). Finally in Step 3, the predicted value is calculated by applying a majority rule to all values of Step 2. Here, since value "3" appears one time and value "2" appears two times, the predicted value is "2".

2.2.2   Reservation based path setup

A contention resolution mechanism is required when several path setup messages compete for the same path or a portion of a path. It directly affects the performance of the setup latency. For the prediction technique, we implement the simplest contention resolution mechanism, we called conventional path setup (CPS) shown in Fig. 6(a). In this case when two path setup messages for

```
History ──►  0 0 0 0 1 2 3 1 2 0 0 1 2 2 3 3 0 0 1 2 2 1 0 0 1 2 ?
```

Step 1. Find the longest suffix (marker) from the history

```
0 0 0 0 1 2 3 1 2 0 0 1 2 2 3 3 0 0 1 2 2 1 0 0 1 2 ?
```

result of step 1: the marker is *0 0 1 2*

Step 2. Record and count the outputs used after the marker

```
0 0 0 0 1 2 3 1 2 0 0 1 2 2 3 3 0 0 1 2 2 1 0 0 1 2 ?
```

result of step 2:   twice *2* and once *3*

Setp 3. Selected the most used port used after the marker.

result of step 3:   the predicted port is 2

Figure 5: Example of prediction using SPM scheme

the same portion of a path (path between node 14 and node 24), one of them is granted the path (communication between node 01 and node 24) and the other one is buffered until the path becomes available. The source-destination pair (11, 34) will set the path after its release by pair (01,24). The two source-destination communications finish at TIME 21.

In Fig. 6(b) we propose a reservation based path setup (RPS) mechanism. In this technique, the ungranted path setup message of the source-destination pair (11,34) instead of being buffered at node 14 where there is a path-conflict, it reserves the path and moves toward the destination. The release path message of the pair (01,24) sets the reserved path for communication at TIME 12. The two pair communications finish at TIME 19. Their latency for communication is reduced by two hop latencies. As shown in this example, the reservation mechanism also can reduce the path setup latency and improves the end-to-end communication latency in the HPNoC . To implement RPS, the electrical arbiter hardware of the conventional electrical router is slightly modified for handling path reservations. RPS only reduces path setup latency when contention occurs in the communication patterns. For traffic patterns such are neighbor in which node trends to communicate with their adjacent nodes, both CPS and RPS performs similarly. Table 1 summarizes the advantages and disadvantages of both path setup mechanisms.

Table 1: Advantages and Disadvantages of CPS and RPS

Path setup	Advantages	Disadvantages
CPS	-Simple arbitration scheme.	-Path setup messages are buffered when path conflicts occur.
RPS	-Reduction of latency when path conflict occurs.	-Extra arbitration required for handling reservation of paths.

3   Performance Evaluation

In this section First we compare the power consumption of an all-electrical NoC and a HPNoC, then we estimate the performance of our proposed path setup techniques for HPNoC.

(a) Conventional path setup mechanism (CPS)

```
TIME 0          TIME 1          TIME 2          TIME 3
[01]   ----->   [02]   ----->   [03]   ----->   [04]
SRC 01

TIME 1          TIME 2          TIME 3          TIME 4    TIME 4
[11]   ----->   [12]   ----->   [13]   ----->   [14]
SRC 11                                                    Conflict

[21]            [22]            [23]    TIME 13 [24]    TIME 5
                                                DEST 24

[31]            [32]            [33]    TIME 14 [34]
                                                DEST 34
```

TIME: one hop delay

SRC 01->DEST 24: TIME 5   Path Setup
                 TIME 6   Optical data transfer
                 TIME 12 Path Release

SRC 11->DEST 34: TIME 14 Path Setup
                 TIME 15 Optical data transfer
                 TIME 21 Path Release

(b) Reservation based path setup mechanism (RPS)

```
TIME 0          TIME 1          TIME 2          TIME 3
[01]   ----->   [02]   ----->   [03]   ----->   [04]
SRC 01

TIME 1          TIME 2          TIME 3          TIME 4    TIME 4
[11]   ----->   [12]   ----->   [13]   ----->   [14]
SRC 11                                                    Reserved

[21]            [22]            [23]    TIME 5   [24]    TIME 5
                                                DEST 24

[31]            [32]            [33]    TIME 6   [34]
                                                DEST 34
```

SRC 01->DEST 24: TIME 5   Path Setup
                 TIME 6   Optical data transfer
                 TIME 12 Path Release

SRC 11->DEST 34: TIME 12 Path Setup
                 TIME 13 Optical data transfer
                 TIME 19 Path Release

Figure 6: (a) Conventional vs (b) Reservation based path setup mechanisms

3.1 Power Consumption Estimation

The main motivation of using photonic NoC is its potential to reduce the high power consumption of an electrical NoC to provide the same performance for intra-chip communications. To offer the same performance of a photonic NoC, electrical NoC requires the use of many parallel links leading to a higher power dissipation of the network.

By scaling the power cost calculation method used in [9] to our 64 nodes torus network we

evaluate the power consumption of the electrical and HPNoC.

In the Electrical NoC, the total energy consumed by the network can be computed as :

$$E_{NETWORK-CYCLE} = (\sum_{j=1}^{N_L} U_{Lj} \times E_{FLIT-HOP}) \times f \tag{1}$$

where U_{Lj} is the average number of flits traversing link j per clock cycle, an estimate on the utilization of link j; $E_{FLIT-HOP}$ is the sum of energy spent by a flit in the different pipeline stages of flits processing; and f the clock frequency of the router.

For the HPNoC, the dissipated energy is estimated as the sum of the energy of two components: the photonic network, and the electrical control network.

- Since the electrical control network differs from the conventional electrical NoCs in terms of message size, the energy can be deduced from the electrical NoC's one using the same equation (1) scaled to the electrical control message size.

- The energy consumed by the photonic network consist of :
  1) The transmission energy which is calculated as :

$$P_{P-NoC,transmission} = NR_{ON-STATE} \times 0.5mW \tag{2}$$

Where $NR_{ON-STATE}$ is the number of micro-ring resonators in "ON" state, and $0.5mW$ is the assumed energy cost for a single micro-ring resonator in "ON" state [9]. No energy is consumed by an "OFF" state micro-ring resonator.
2) And the modulator/demodulator energy is estimated as:

$$P_{P-NoC,mod/demod} = 0.11pJ/bit.64.Bandwidth \tag{3}$$
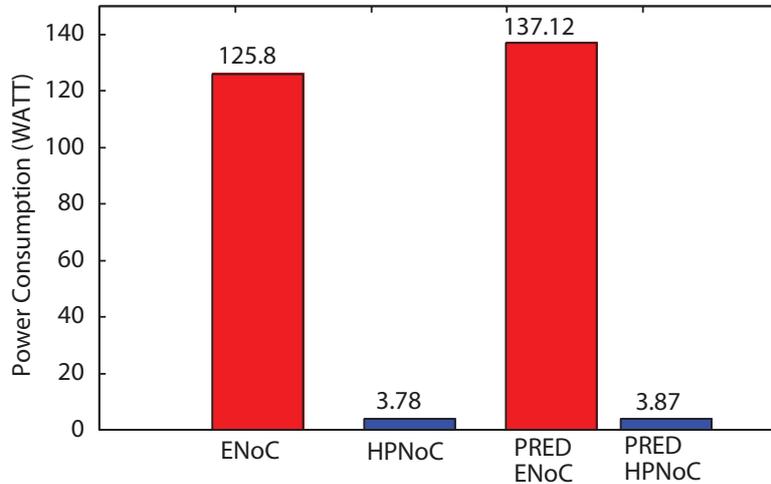
Figure 7: Power Consumption Cost

We compute the energy consumed by a HPNoC and a fully electrical NoC for a 32 nm node technology that uses a 5 GHz router clock frequency to provide the same performance bandwidth. By assuming an average link utilization of 50% for the 64 nodes torus of 800Gbps data transmission bandwidth, we estimated the energy consumed by the two networks. When using the prediction router the energy consumed is majorated by an extra 9% of the electrical network energy due to the extra overhead added by the prediction router [6].

Fig. 7 plots the power estimation results. It shows that the electrical NoC consumes a huge amount of power compared to the HPNoC to be able to deliver the same bandwidth performance. It further shows that the extra energy overhead required when using the prediction router is almost neglectable for the HPNoC.

3.2   Simulation Conditions

We evaluate the performance of the networks using a modified version of the booksim [3] cycle accurate simulator. For simulation, we use three probabilistic traffic patterns :

- *Uniform random :* Each node sends a packet to a randomly chosen node.

- *Neighbor :* Each node sends a packet to its neighboring nodes.

- *Bitreversal :* Each node sends a packet to a destination whose address is the bitreversal of the sending node address.

The Table 2, 3, and 4 summarize our simulation parameters.

Table 2: Simulation Parameters

Simulated Networks	$ENoC$ (w/wo prediction), $HPNoC$ (w/wo prediction), $HPNoC$ (CPS, RPS)
Topology	$2D\ Torus\ 64\ nodes$
Routing	DOR
Control message size	$4\ Bytes$
Data size	$20\ Bytes$
Prediction algorithms	$LP,\ SPM$
Traffic patterns	$Uniform,\ Neighbor,\ Bitreversal$

Table 3: Optical NoC Parameters

Number of wavelengths per waveguide	64
Data rate per wavelength	$12.5Gbps$
Total link bandwidth	$800\ Gbps$

Table 4: Electrical NoC Parameters

Router frequency	$5\ GHz$
Number of VC per physical channel	2
Channel width	$32\ bits$
Buffer size/VC/channel	$20\ Bytes$
Latency/hop without using prediction	$4\ router\ cycles$
Latency/hop when prediction is used and hit	$1\ router\ cycle$
Latency/hop when prediction is used and miss	$4\ router\ cycles$

3.3   Results and Discussion

The predictive switching and RPS are techniques to reduce the latency path setup messages spend in the electrical control network. By reducing the average path setup latency, the control network with this techniques can afford more messages before network saturation thus improving the overall performance of the HPNoC.

Fig. 8 (a), (b), and (c) show the simulation results for a fully electrical network under uniform, neighbor, and bitreversal traffic patterns, respectively. The results show that both LP and SPM prediction techniques improve the performance of the network for all traffic patterns. For instance,
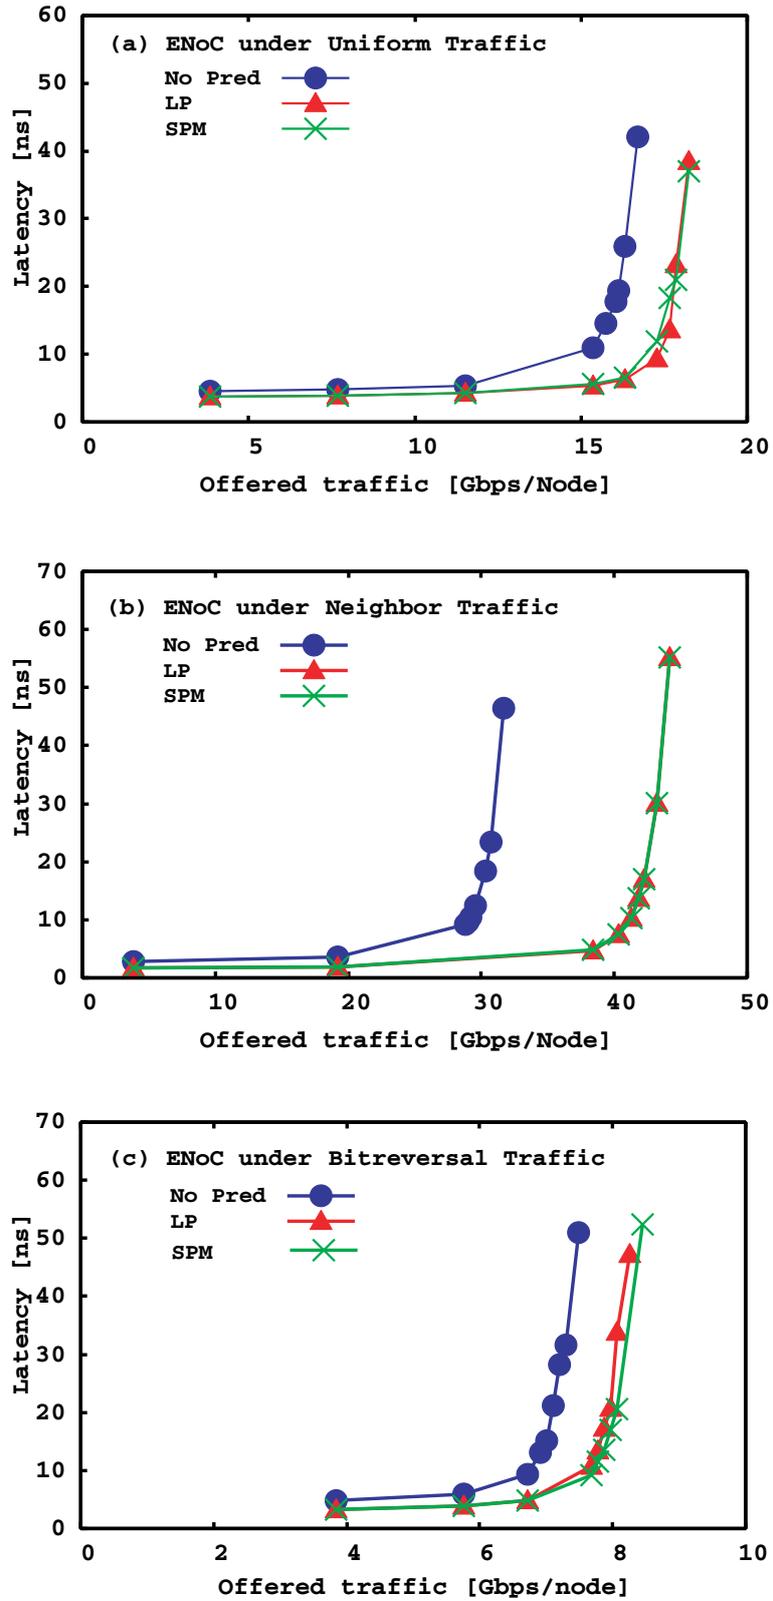
Figure 8: Electrical NoC under uniform (a), neighbor (b), and bitreversal traffic patterns (c)

using the prediction router, the electrical NoC can be loaded with nearly an extra 10 Gbps/node compared to the conventional electrical one for the neighbor traffic pattern as shown in Fig. 8 (b). In the case of uniform traffic patterns, due to the random communication pattern, LP and SPM schemes show nearly the same performance as seen in Fig. 8 (a). In the case of neighbor traffic pattern, due to the fact that nodes trend to communicate with their adjacent nodes, the LP scheme obtains nearly the same prediction hit rate as the SPM, leading to almost the same improvement of latency as shown in Fig. 8 (b). As seen in Fig. 8 (c), SPM prediction technique shows better performance than LP under bitreversal traffic pattern due to the analysis on the longer output history used by an input channel of SPM.

In Fig. 9 (a), (b), and (c), the HPNoC performance is evaluated for uniform, neighbor, and bitreversal, respectively with and without LP and SPM prediction mechanisms. The results show that both prediction techniques improve the network performance. In particular for neighbor traffic pattern shown in Fig. 9 (b), this performance is almost doubled with the prediction techniques. Furthermore, these results also show that even with the simplest LP prediction technique which requires only a single output history at each input-channel, we can achieve a considerable increase in performance.

Fig. 10 (a), (b), and (c) show a comparison of HPNoC against a fully electrical NoC under uniform, neighbor, and bitreversal traffic patterns, respectively with and without prediction technique. The results show that the HPNoC with the simplest LP predictive switching leads to better performance than all other simulated network configurations for all traffic patterns. Since the HPNoC uses a circuit switching flow control even for neighboring communication, a setup packet for establishing a path is necessary before communication can take place. The effect of path setup time for such communication pattern is particularly important in message delivery latency. That causes the packet switching ENoC without or with prediction outperforming the HPNoC without prediction as shown in in Fig. 10 (b). However, by reducing the effect of path setup time the HPNoC with prediction outperforms all other configurations.

In Fig. 11 (a), and (b), we compare the performance of the conventional path setup (CPS) mechanism and our proposed scheme (RPS) for uniform, and bitreversal traffic patterns, respectively. Results show an improvement in all cases. By reserving the path ahead a time instead of buffering the path setup message, the average path setup latency is considerably improved leading to a better overall performance of the HPNoC.

4   Conclusion and Future Works

Well designed optical interconnection has the potential to meet the high bandwidth and low power consumption required for future on-chip interconnection. In this paper, we have proposed path setup techniques to reduce the path setup latency for circuit switching HPNoC. The simulation results for probabilistic traffic patterns show that both techniques drastically improve the network performance of a conventional HPNoC. As crucial performance factor of the HPNoC is the setup time of the optical path, reducing the path setup latency in the electric NoC leads to a considerable gain in overall performance for HPNoC. An improvement of this study is to extend the simulation to some realistic traffic patterns. For the implementation of the RPS we slightly modify the arbitration scheme of a conventional electrical router, although no power consumption measurement were conducted yet, we assumed that the changes have no impact in the hardware and power consumption. A further analysis of the power consumption and hardware cost are left for future work. We will also investigate the use of dynamic wavelengths allocation which allows wavelengths of the same waveguide to be used for several message streams.

Acknowledgments
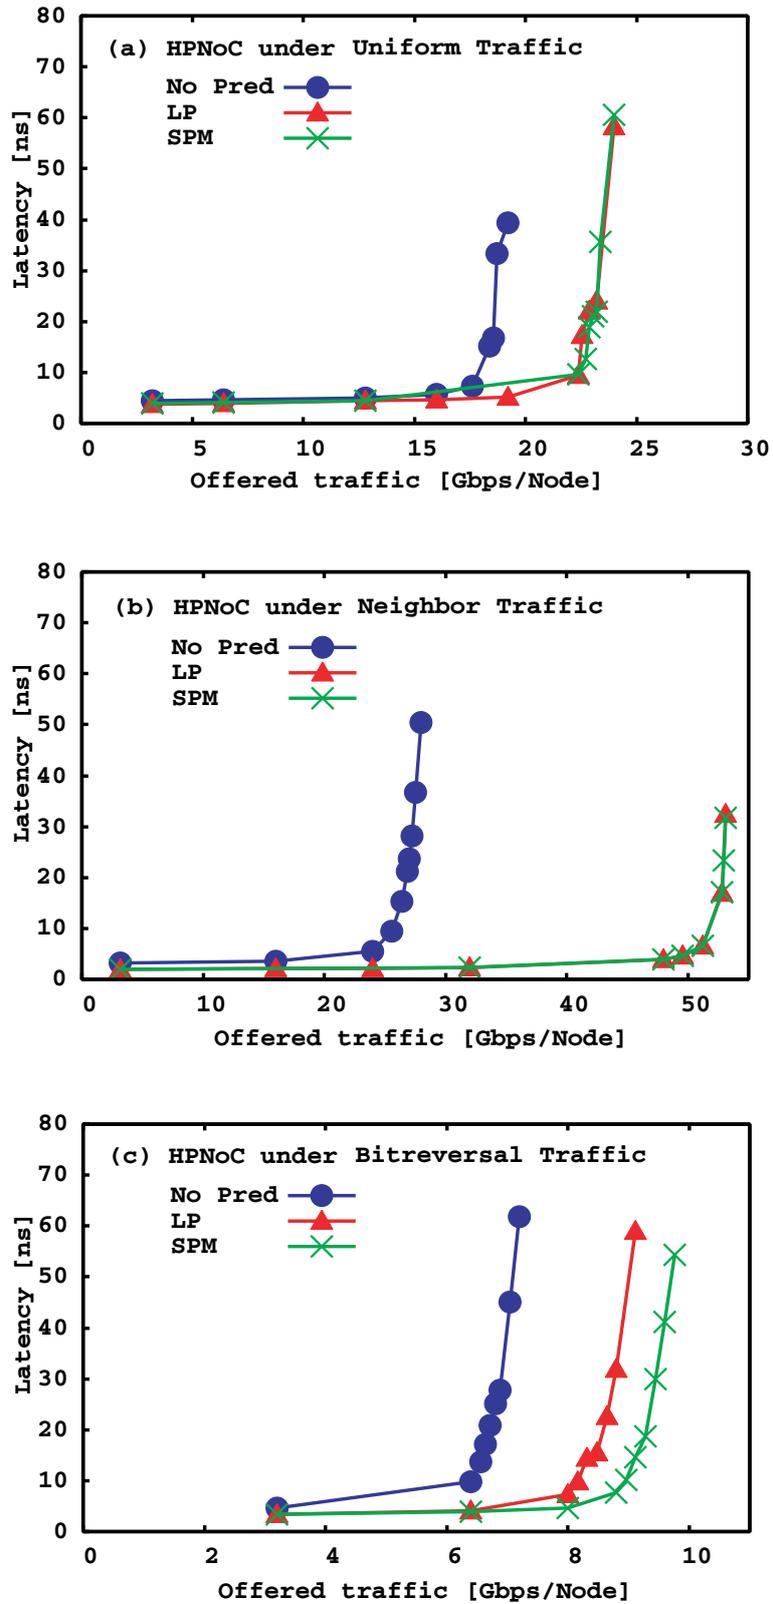
Figure 9: HPNoC under uniform (a), neighbor (b), and bitreversal traffic patterns (c)
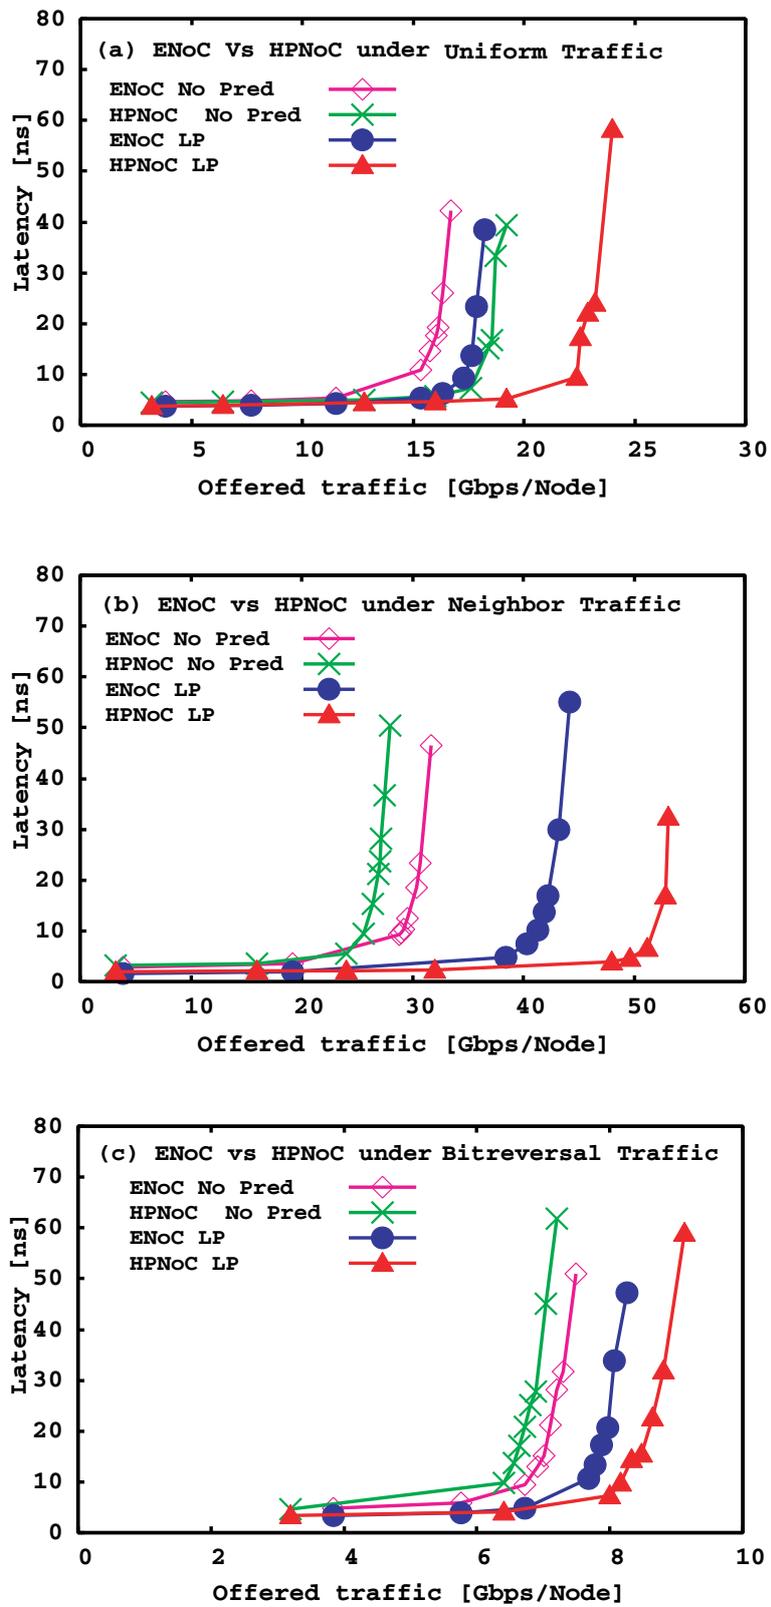
Figure 10: Electrical NoC vs HPNoC under uniform (a), neighbor (b), and bitreversal traffic patterns (c)
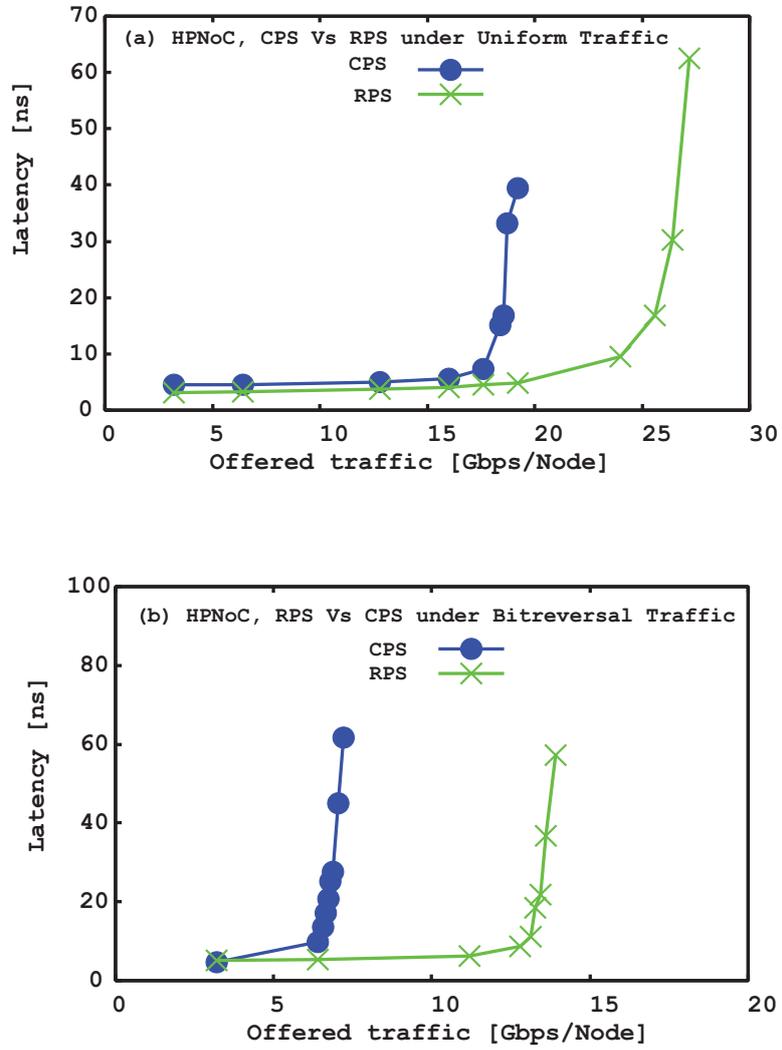
Figure 11: HPNoC, CPS vs RPS under uniform (a), and bitreversal traffic patterns (b)

References

[1] Christopher Batten, Ajay Joshi, Jason Orcutt, Anatoly Khilo, Benjamin Moss, Charles Holzwarth, Milos Popovic, Hanqing Li, Henry Smith, Judy Hoyt, Franz Kartner, Rajeev Ram, Vladimir Stojanovic, and Krste Asanovic. Building manycore processor-to-dram networks with monolithic silicon photonics. In *Proceedings of the 2008 16th IEEE Symposium on High Performance Interconnects*, pages 21–30, Washington, DC, USA, 2008. IEEE Computer Society.

[2] Mark J. Cianchetti, Joseph C. Kerekes, and David H. Albonesi. Phastlane: a rapid transit optical routing network. In *ISCA'09*, pages 441–450, 2009.

[3] William Dally and Brian Towles. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.

[4] Gilbert Hendry, Johnnie Chan, Shoaib Kamil, Lenny Oliker, John Shalf, Luca P. Carloni, and Keren Bergman. Silicon nanophotonic network-on-chip using tdm arbitration. *High-Performance Interconnects, Symposium on*, 0:88–95, 2010.

[5] P. Jacquet, W. Szpankowski, and I. Apostol. A universal predictor based on pattern matching. *Information Theory, IEEE Transactions on*, 48(6):1462 –1472, June 2002.

[6] H. Matsutani, M. Koibuchi, H. Amano, and T. Yoshinaga. Prediction router: Yet another low latency on-chip router architecture. In *High Performance Computer Architecture, 2009. HPCA 2009. IEEE 15th International Symposium on*, pages 367 –378, 2009.

[7] Hiroki Matsutani, Michihiro Koibuchi, Hideharu Amano, and Tsutomu Yoshinaga. Prediction router: A low-latency on-chip router architecture with multiple predictors. *Computers, IEEE Transactions on*, 60(6):783 –799, june 2011.

[8] Yan Pan, Prabhat Kumar, John Kim, Gokhan Memik, Yu Zhang, and Alok Choudhary. Firefly: illuminating future network-on-chip with nanophotonics. In *Proceedings of the 36th annual international symposium on Computer architecture*, ISCA '09, pages 429–440, New York, NY, USA, 2009. ACM.

[9] A. Shacham, K. Bergman, and L.P. Carloni. The case for low-power photonic networks on chip. In *Design Automation Conference, 2007. DAC '07. 44th ACM/IEEE*, pages 132 –135, 2007.

[10] A. Shacham, K. Bergman, and L.P. Carloni. On the design of a photonic network-on-chip. In *Networks-on-Chip, 2007. NOCS 2007. First International Symposium on*, pages 53 –64, May 2007.

[11] Assaf Shacham, Benjamin G. Lee, Aleksandr Biberman, Keren Bergman, and Luca P. Carloni. Photonic noc for dma communications in chip multiprocessors. In *Proceedings of the 15th Annual IEEE Symposium on High-Performance Interconnects*, HOTI '07, pages 29–38, Washington, DC, USA, 2007. IEEE Computer Society.

[12] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N.P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R.G. Beausoleil, and J.H. Ahn. Corona: System implications of emerging nanophotonic technology. In *Computer Architecture, 2008. ISCA '08. 35th International Symposium on*, pages 153 –164, 2008.