Theoretical and Applied Informatics
ISSN 1896–5334
Vol. 25 (2013), no. 1
pp. 41–48
DOI: 10.2478/thai-2013-0004

Optical interconnection networks with time slot routing

Ireneusz Szcześniak, Roman Wyrzykowski

Institute of Computer and Information Sciences
Częstochowa University of Technology
ul. Dąbrowskiego 73
42-200 Częstochowa
Poland

Received 10 January 2013, Revised 6 April 2013, Accepted 22 April 2013

Abstract: We propose the *time slot routing*, a novel routing scheme that allows for a simple design of interconnection networks. The simulative results show that the proposed scheme demonstrates optimal performance at the maximal uniform network load, and for uniform loads the network throughput is greater than for deflection routing.

Keywords: Interconnection networks, networks on chip, Beneš network, time slot, deflection, simulation.

1. Introduction

Interconnection networks not only connect components on a chip (networks on chip, NoC), but also separate chips on a motherboard along with computing and storage nodes. These networks are complicated, require large buffers, cause large latencies, drop packets and perform poorly under heavy loads – suffering the opposite of their desired properties.

A number of improvements have been recently proposed in the design of interconnection networks. In [3] bufferless routing is proposed, which is substantiated by extensive simulative performance evaluation. Small buffers and bufferless routing has been proposed a number of years ago [1].

In [4] authors on one hand recognize the bufferless advantage of deflection routing, and on the other hand notice the throughput degradation resulting from overloading the network, and so they propose methods of controlling congestion for bufferless deflection-routed networks.

We, in turn, introduce a novel idea of *time slot routing* for interconnection networks. Time slot routing is a form of time-division multiplexing (TDM) that is used in communication networks to share a transmission link, but we devised time slot routing as a simple and efficient means of routing packets in an interconnection network.

The article is organized as follows. In the next section, time slot routing is introduced. Then its performance is evaluated simulatively and compared to the performance of a network with deflection routing and with store-and-forward routing. The article ends with conclusions.

2. Time slot routing

There are n interconnected nodes that exchange data packets synchronously, i.e. according to a single clock shared by all nodes. The interconnection network does not have buffers, and its configuration, i.e. the configuration of its switching elements, determines the connections between inputs and outputs of the network. The delay incurred by the optical interconnection network is small and equals to the delay it takes light to traverse the network. The delay is below a nanosecond, and we assume it equals to one time slot.

A specific way of connecting inputs to outputs we call a permutation. In a single permutation every source node is connected to a destination node different from the source node. Permutations are changed every time slot, and so the time slot number determines the destination node for a packet sent from a given source node. There are $(n-1)$ permutations required, which are repeated periodically, so that each node can send packets to each of the other $(n-1)$ nodes. Tab. 1 reports three sample permutations P1, P2, and P3 required for a 4×4 network. The notation $1 \rightarrow 2$ used in the table means that node 1 is connected to node 2.

To demonstrate time slot routing scheme and to evaluate its performance, we chose the Beneš network for its small number of N required switching elements, as given by (1), in comparison to other networks such as the Banyan networks. Any other network type can be used with time slot routing as long as it is rearrangeable; the nonblocking property is not required as the network can be rearranged every time slot. If the constant latency is required, then every path between the source and destination nodes should be of equal length, which holds true for the Beneš network.

$$N = \frac{n(2\log_2(n) - 1)}{2} \tag{1}$$

Figures 1a, 1b, and 1c show the configurations of the 4×4 Beneš network that implement permutations P1, P2, and P3, respectively.

Depending on the network type, a permutation can be implemented by more than one network configuration. Having different configurations for the same permutation can allow for less frequent reconfigurations of the switching elements, which would relax the requirements on the properties of switching elements, such as reconfiguration time, reducing the cost of the network and increasing the network performance.

P1	P2	P3
$1 \to 2$	$1 \to 4$	$1 \to 3$
$2 \to 1$	$2 \to 3$	$2 \to 4$
$3 \to 4$	$3 \to 2$	$3 \to 1$
$4 \to 3$	$4 \to 1$	$4 \to 2$

Tab. 1. Sample required permutations for a 4×4 network

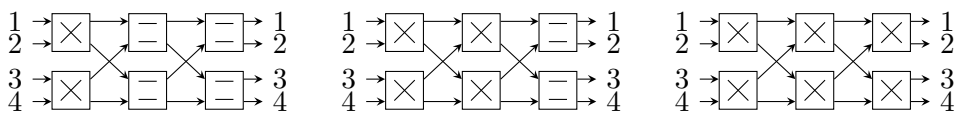

Fig. 1. Configurations of the 4×4 Beneš network

For instance, Fig. 2 shows configurations that are different from the configurations in Fig. 1, but implement the same permutations from Tab. 1. The sequence of configurations shown in Fig. 1a, 1b, 1c, 2a, 2b, 2c would produce the sequence of permutations P1, P2, P3, P3, P2, and P1, which would require the reconfiguration of every switching element at most every other time slot.

The advantages of time slot routing are as follows.

- *Simple design*. Time slot routing does not require buffers or packet header processing.
- *No packet loss*. Since no packets are lost in the network, there is no need for acknowledgments and retransmissions.
- *Optimal throughput* under the maximal uniform load.

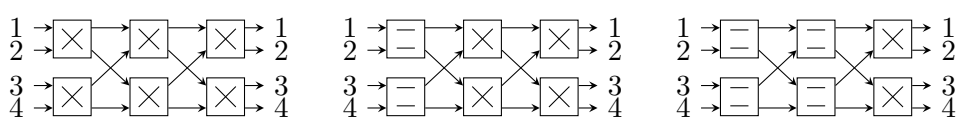

Fig. 2. Extra configurations of the 4×4 Beneš network

- *No overhead*, since no packet header is required. The network routes packets according to the time slot number.
- *Constant network latency*. Every packet travels in the network through the same number of nodes and links of equal length.
- *Fairness*. Every source node has a chance to send the same number of packets to the destination nodes as any other node.

3. Performance evaluation

The performance of time slot routing is compared to the performance of deflection routing and store-and-forward routing for three sizes of the Beneš network: 4×4 (6 switching elements), 16×16 (56 switching elements), and 64×64 (352 switching elements). The admission buffers at the nodes can grow without a limit, but the buffering at switching elements is subject to constraints as described below.

Deflection routing routes a packet along the shortest path, and misroutes it if the preferred output has already been taken by the other packet. If a packet prefers either output at a switching element (i.e. both outputs yield a path to the destination of the same length), one of them is chosen at random, which helps to improve the network performance and fairness. We used deflection routing results for comparison, because just like time slot routing, deflection routing does not require buffers.

Store-and-forward routing is used in traditional electronic interconnection networks, and requires buffers at every switching element. We evaluated the performance of storeand-forward routing for three sizes of buffers: 1, 3 and 5. Each output at a switching element has a buffer. When a packet arrives, it is enqueued in the buffer of its preferred output. If there is no space left in the buffer, the packet is dropped. If the packet prefers either output, the output is chosen at random, which helps to improve the network performance and fairness. If there is no space left in the randomly chosen buffer, the packet is enqueued in the buffer of the other output. If there is no space left in the other buffer, the packet is dropped. We assume that a packet stays at a switching element at least one time slot due to the need of electronic processing and buffering.

We carry out the performance evaluation simulatively. The analysis of a TDM channel with a Poisson input stream should apply to time slot routing, but it is rather complex as the discrete and continous time domains are intermixed [2], and so we reverted to simulations. Deflection routing could be evaluated analytically too, but its analysis is difficult and gives approximate results [5]. We evaluated the network throughput, the number of dropped packets, the admission queue delay, the total delay, and the admission queue size.

We assume that the network is uniformly loaded: each node sends packets to each of the other $(n-1)$ nodes, and so there are $n(n-1)$ flows in the network. The uniform load assumption is a common assumption for data center networks, and multi-threaded or distributed programs. At each source node there is one admission queue per flow that we allow to grow without a limit. Packets arrive to the buffer according to the Poisson distribution with intensity $\lambda = l$ packets per time slot, where $l \in (0, 1)$ is the network load. The network is maximally loaded when $l = 1$, i.e. every node has on average one message to send every time slot. Since every node sends packets uniformly to other $(n-1)$ nodes, the intensity of each flow is $\lambda/(n-1)$.

3.1. Simulations

The performance evaluation was carried out with 3000 simulation runs. There were five groups of simulations: one for time slot routing, one for deflection routing, and three for store-and-forward routing for the buffer sizes of 1, 3, and 5. There were three network sizes ($n = 4, 16, 64$) and twenty different loads ($l = 0.05, 0.1, \ldots, 1.0$), and so there were $5 \times 3 \times 20 = 300$ test cases, where each test case was simulated ten times with different pseudo-random values, totaling 3000 simulation runs.

We calculated the mean values of the metrics of interest for a test case based on the results of ten simulation runs for that test case. We also calculated the standard errors for the mean values, and found that they were below 1% of the mean value.

3.2. Comparison

Figures 3, 4, 5, 6, and 7 show the simulation results. In each of the figures, the subfigures (a) are for the 4×4 network, (b) for the 16×16 network, and (c) for the 64×64 network. In each of the figures the results for time slot routing are shown with the solid lines, and for deflection routing with the bullets (•). The results for store-and-forward routing are shown with the pluses (+) for buffers of size 1, with the crosses (×) for the buffers of size 3, and with the circles (◦) for the buffers of size 5. There are no error bars shown for the data points, because the standard errors were too small to be drawn.

The mean network throughput is shown in Fig. 3. Deflection routing performs poorly for high loads, because deflected packets saturate the network, thwarting packet admission and delivery. In contrast, time slot routing performs optimally in that it delivers all the offered load, reaching the maximum throughput at heavy loads: at most n packets per time slot for the $n \times n$ network. Also store-and-forward routing with buffers of size 5 performs worse than time slot routing.

The mean number of packets dropped per time slot is shown in Fig. 4. Since time slot routing and deflection routing do not lose packets, their results are not shown in

the figure. Store-and-forward routing, however, suffers a large number of dropped packets, even when buffers of size 5 are used.

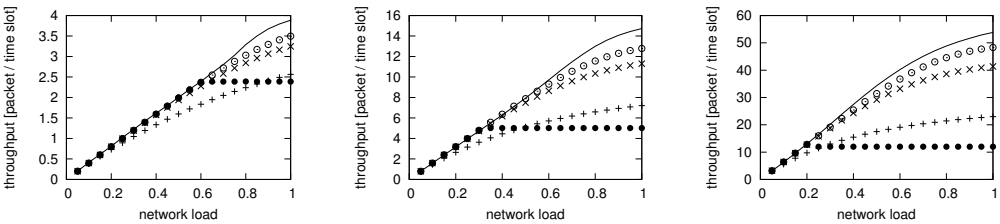

(a) (b) (c)

Fig. 3. Comparison of the mean network throughput

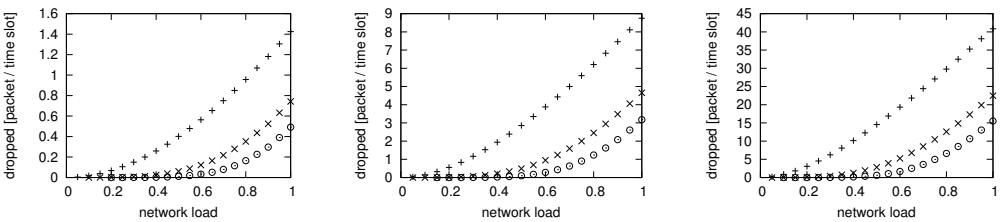

(a) (b) (c)

Fig. 4. Comparison of the mean number of dropped packets

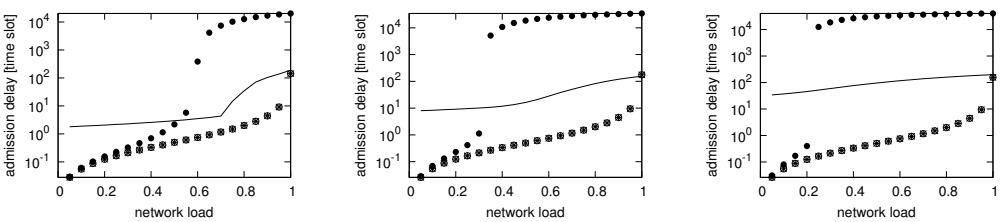

(a) (b) (c)

Fig. 5. Comparison of the mean admission delay

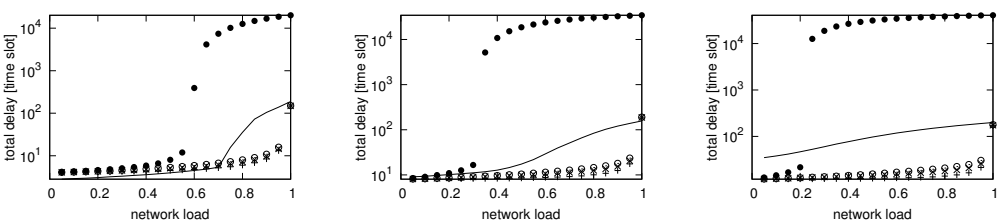

(a) (b) (c)

Fig. 6. Comparison of the mean total network delay

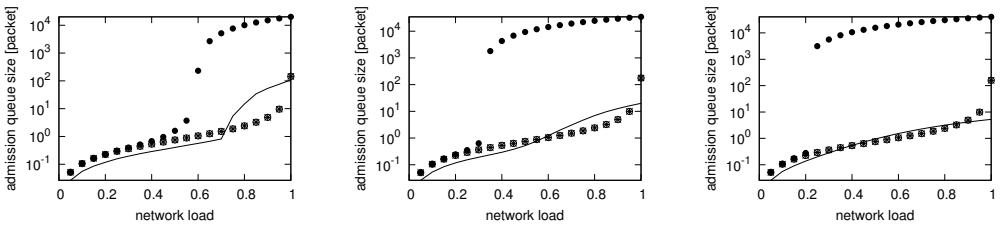

Fig. 7. Comparison of the mean admission queue size

The mean admission delay measured in time slots is shown in Fig. 5. Generally, deflection routing causes the largest delay because deflected packets live in the network long, blocking admission of new packets. The mean admission delay for time slot routing is at least $(n-1)/2$, and so it is larger than the admission delay for store-and-forward routing [2].

The mean total delay measured in time slots is shown in Fig. 6, which is the sum of the number of time slots spent waiting in the admission queue, and the number of time slots spent in the network. Again, time slot routing outperforms deflection routing. For and store-and-forward routing, the admission delay dominates, and so the time slot routing in general performs worse than store-and-forward routing.

The mean size of admission queues is shown in Fig. 7. For deflection routing, the long-lived deflected packets cause the admission queues to grow to very large sizes. Time slot routing and store-and-forward routing cause the admission queues to grow to comparable sizes.

4. Conclusion

We presented a novel concept of time slot routing, and evaluated its performance for the network of the Beneš type, but other network types can be used too. We report that time slot routing outperforms deflection routing, and compares favorably with store-and-forward routing. Time slot routing can be used not only for electronic NoC, but also optical and electronic interconnects used to connect computing and storage nodes. Time slot routing, unlike store-and-forward routing, allows for simple network design without buffers and header processing.

An interesting problem for future work is finding a network type, or algorithms of configuring existing network types, that would require only infrequent reconfiguration of switching elements.

References

1. M. Enachescu, Y. Ganjali, A. Goel, N. McKeown, T. Roughgarden, *Routers with very small buffers*, Proceedings of INFOCOM 2006, pages 1–11, April 2006.

2. King-Tim Ko and B. Davis, *Delay analysis for a TDMA channel with contiguous output and poisson message arrival*, IEEE Transactions on Communications, 32(6):707–709, June 1984.

3. T. Moscibroda, O. Mutlu, *A case for bufferless routing in on-chip networks,* SIGARCH Comput. Archit. News, 37:196–207, June 2009.

4. G. Nychis, Ch. Fallin, T. Moscibroda, O. Mutlu, *Next generation on-chip networks: what kind of congestion control do we need?* In Proceedings of the Ninth ACM SIGCOMM Workshop on Hot Topics in Networks, Hotnets '10, pages 12:1–12:6, New York, NY, USA, 2010. ACM.

5. I. Szcześniak, B. Mukherjee, T. Czachórski, *Approximate analytical performance evaluation of synchronous bufferless optical packet-switched networks*, IEEE/OSA Journal of Optical Communications and Networking, 3(10):806–815, October 2011.

Sieci połączeń z rutingiem ze szczelinami czasowymi

Streszczenie

W artykule opisujemy ruting ze szczelinami czasowymi, który wymaga jedynie prostej architektury sieci połączeń. Uzyskane wyniki symulacyjne pokazują, że ruting ze szczelinami czasowymi pozwala na osiągnięcie maksymalnej możliwej przepustowości sieci przy maksymalnym równomiernym obciążeniu sieci. Wyniki pokazują także, że ruting ze szczelinami czasowymi osiąga lepsze wyniki wydajnościowe niż ruting z odbiciami.