

Empirical Evaluation of Deterministic and Adaptive Routing with Constant-Area Routers

D. R. Miller

W. A. Najjar

Department of Computer Science, Colorado State University, Ft. Collins, CO 80523

{millerdi,najjar}@cs.colostate.edu

Abstract

This paper addresses the issue of how router complexity affects the overall performance in deterministic and adaptive routing under virtual cut-through switching in k -ary n -cube networks. First, the performance of various adaptive routers with constant area are compared. Second, the performance of adaptive and deterministic routers are compared under the same conditions. Finally, it is shown that, under certain conditions, deterministic routers can reach saturation points comparable to adaptive routers.

1 Introduction

The performance of the communication subsystem depends on several factors such as the network topology, the channel width, the routing algorithm and the router design and implementation. The routing algorithm determines the path taken by a message while traveling from its source to its destination. In dimension-order routing, a message is routed along decreasing dimensions with a dimension decrease occurring only when zero hops remain in all higher dimensions. Virtual Channels (VCs) are included in the router to avoid deadlock [11]. Deterministic routing algorithms can suffer from congestion since only a small subset of all possible paths between a source and destination are used.

In adaptive routing, messages are not restricted to a single path when traveling from source to destination: the choice of path can be made dynamically in response to current network conditions. Such schemes can minimize unnecessary waiting, and provide fault-tolerance. Several studies have shown that adaptive routing can achieve a lower latency, for the same load, than deterministic routing with the same clock cycle [19, 25].

The delay experienced by a message at each node can be broken down into: *switching* (or *routing*) de-

lay and *queuing* (or *buffering*) delay. The former is determined primarily by the complexity of the router. The latter is determined by the congestion at each node which in turn is determined by the degrees of freedom the routing algorithm allows a message. The main performance advantage of adaptive routing (besides its fault-tolerance) is that it reduces the queuing delay.

However, the clock cycle time of deterministic routers can be significantly lower than adaptive ones as shown in [7, 2]. Two main reasons explain this phenomenon:

- *Number of VCs:* Two VCs are sufficient to avoid deadlock in dimension ordered routing [11]; while adaptive routing (as in [16] and [4]) requires a minimum of three VCs in k -ary n -cube networks.
- *Output channel selection:* In dimension ordered routing, the output channel selection policy is very simple: it depends only on information contained in the message header itself whereas in adaptive routing the output channel selection policy depends also on the state of the router (i.e. the occupancy of various VCs) causing increased router complexity and thereby higher switching delays.

Because of these differences in complexity, the switch delays for adaptive routers can be much larger than those for deterministic routers. The results in [7, 2] show that the switch delays for the various adaptive routers are about half to more than twice as long as the dimension-order router for worm-hole routing. On the other hand, both deterministic and adaptive, require a variable amount of resources such as buffer area or physical channels between nodes. In [17], the advantage of adaptive routing in reducing queuing delays in the nodes between source and destination is accounted for in worm-hole routing.

In this paper we report on the performance of deterministic and adaptive routers for k -ary n -cube net-

works evaluated with a constant router area and only one physical channel per dimension per node using virtual cut-through switching. Router area, in this paper, is defined here as the buffer size *times* the number of VCs. The evaluation accounts for the increased delays due to varying buffer sizes in the router cost model.

The deterministic and adaptive routing algorithms on which this study is based are described in Section 2. The switch delay model, is based on the one used in [7, 2], is described in Section 3. The simulation results are discussed, in Section 4. Related work is discussed in Section 5 and concluding remarks in Section 6.

2 Routing and Switching Models

The network model used in this study is a k -ary n -cube using virtual cut-through switching [20]: message advancement is similar to worm-hole routing [29], except that the body of a message can continue to progress even while the message head is blocked, and the entire message can be buffered in a single node. Note that a header flit can progress to a next node only if the whole message can fit in the destination buffer. For simplicity all messages are assumed to have the same length. Three different traffic patterns are considered:

- Random Uniform: Source and destination nodes are uniformly distributed.
- Complement: Node $a_{n-1}a_{n-2}\dots a_1a_0$ communicates with node $\overline{a_{n-1}a_{n-2}\dots a_1a_0}$
- Perfect Shuffle: Node $a_{n-1}a_{n-2}\dots a_1a_0$ communicates with node $a_{n-2}a_{n-3}\dots a_0a_{n-1}$

2.1 Routing Models

The *deterministic routing* algorithm used is dimension-order routing [9, 11]. A message is routed along decreasing dimensions with a dimension decrease occurring only when zero hops remain in all higher dimensions. By assigning an order to the network dimensions, no cycle exists in the channel-dependency graph and the algorithm is deadlock-free.

The *adaptive routing* algorithm used is the one described in [15, 16, 4] (also known as the *-channels algorithm). Adaptive routing is obtained by using VCs along with dimension-order routing. A message can be routed, adaptively, in any dimension until it is blocked. Once a message is blocked, it is then routed using the dimension-order routing. Note that a message can still

return to adaptive routing at subsequent nodes. This algorithm has been proven to be deadlock-free with the following routing restrictions: when the message size is greater than the buffer size (i.e. size of the VC), deadlock is prevented by allowing the head flit of a message to advance to the next node only if the receiving queue at that node is empty. If the message size is less than the buffer size, deadlock is prevented by allowing a message to advance only when the whole message fits in the receiving queue at that node. This algorithm requires a minimum of three VCs per dimension per node for each physical unidirectional channel: the number of VCs grows linearly with the size of the network.

2.2 Switching Models

Both the deterministic and adaptive routing algorithms were implemented using one *physical channel* (PC) per dimension per node. Figure 1 shows a schematic for each of the routers simulated here for the 2D case. For both cases there is only one PC for the sink channel. Once this channel is assigned to a message, it is not released until the whole message has finished its transmission. All channels are unidirectional.

Note that the deterministic router uses storage buffers associated with output channels, while the adaptive router uses storage buffers associated with input channels. When using output buffers, the routing decision is made before buffering the message. This type of routing is ideal for deterministic routing because only one choice is available for an incoming message. When a message comes into a node, it can be immediately placed into the appropriate buffer.

When using input buffers, the routing decision is made after buffering the message in the buffer associated with the input channel. This strategy lacks the problem of early commitment of output channels. Since a message can usually be routed on several possible output channels in adaptive routing, this buffering strategy was used for the adaptive router.

The input/output selection policy used for adaptive routing is as follows: a round-robin policy is used for message selection first among all adaptive buffers and then among all deterministic buffers. Output channel selection is performed in each dimension with decreasing number of hops until a free channel is found. By using this output channel selection policy, the greatest amount of adaptivity for a message is retained which reduces blocking.

Note that because both the deterministic and adaptive routers are pipelined, buffers are needed on both

the input and output channels. In addition, the input buffer of the deterministic router ensures that no message is lost if a conflict arises.

3 Modeling Router Delay

The router delay models are based on the ones described in [7, 2, 17]. These models account for both the logic complexity of the routers as well as the size of the crossbar as determined by the number of VCs that are multiplexed on one physical channel. The models were modified to account for the varying buffer space in virtual cut-through switching as used in this paper. The model parameters are:

Symbol	Variable (delay)
T_{AD}	Address decoding
T_{ARB}	Routing arbitration
T_{CB}	Crossbar
T_{FC}	Flow control
T_{SEL}	Header selection
T_{VC}	Virtual channel controller
P	Max. no. of IP or OP ports in crossbar
F	Degrees of freedom (OP choices of a message)
C	No. of virtual channels
B	Buffer size (in number of flits)

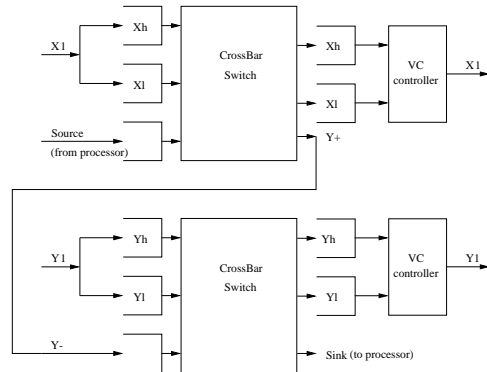
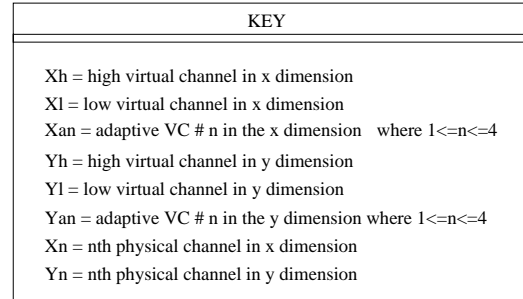
T_{AD} includes the time for examining the packet header and creating new packet headers for all possible routes. T_{ARB} is the time required for selecting among all possible routes. T_{CB} is the crossbar delay. The switch's crossbar is implemented as a tree of gates. T_{FC} includes the time for flow control between routers so that buffers do not overflow. T_{SEL} is the time for selecting the appropriate header. T_{VC} includes the time required for multiplexing VCs onto physical channels.

A deterministic router routes a message in either in the same dimension (either low or high channel) or routes it to the next dimension. Therefore the number of degrees of freedom (F) equals the number of switch crossbar ports (P).

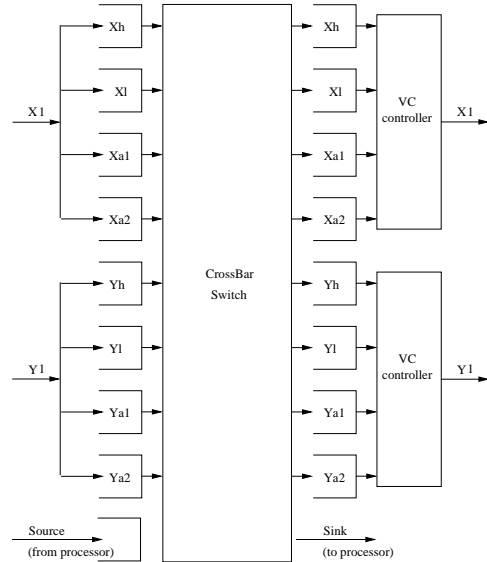
For adaptive routers, $F = P - 2(n - 1)$ (n is the network dimension) because adaptive routing can use the adaptive channels in all the dimensions while only two VCs per physical channel can be used in dimension-order (to avoid deadlock). Note that this expression includes the delivery port.

The timing parameters that are estimated by these models are:

- T_r : Time to route a message.



(a) Deterministic Router



(b) Adaptive Router

Figure 1: Schematics of 2D routers

B	T_r	T_s	T_c	CC Period
8	6.60	5.15	6.74	6.74
16	6.60	5.95	6.74	6.74
24	6.60	6.42	6.74	6.74
32	6.60	6.75	6.74	6.75
48	6.60	7.22	6.74	7.22
64	6.60	7.55	6.74	7.55
96	6.60	8.02	6.74	8.02

Table 1: Delays for Deterministic Routing Algorithm for k -ary 2-cube and 3-cube networks. All values in nanoseconds ($C = 2$ and $P = F = 3$)

- T_s : Transfer time of a flit to the corresponding output channel.
- T_c : Transfer time of a flit across a PC.

The delay equations are: $T_r = T_{AD} + T_{ARB} + T_{SEL}$

$$T_r = 2.7 + 0.6 + 0.6 * \log_2 F + 1.4 + 0.6 * \log_2 F$$

$$T_s = T_{FC} + T_{CB} + T_{Latch}$$

$$T_s = 0.8 + 0.6 * \log_2 B + 0.4 + 0.6 * \log_2 P + 0.8$$

$$T_c = 4.9 + T_{VC}$$

$$T_c = 4.9 + 1.24 + 0.6 * \log_2 C$$

The constants in these equations were obtained in [7] using router designs along with gate-level timing estimates based on a 0.8 micron CMOS gate array process. Using the above equations, the delay values were calculated for each of the router algorithms simulated and are shown in Tables 1, 2, and 3. To decrease the overall router delay, it is assumed that all three operations are overlapped through pipelining as described in [17], and therefore the clock period is determined by the longest delay:

$$T_{cc \text{ period}} = \text{Max}(T_r, T_s, T_c)$$

A number of observations can be made about the data in Tables 1, 2, and 3:

- In deterministic routers, increasing buffer size increases the overall router delay when moderate to large buffer sizes are used. For small buffer sizes the clock cycle is dominated by the transfer time T_c while for larger ones it is dominated by the switching time T_s .
- In adaptive routers, the clock cycle time is dominated by T_r . Increasing buffer size only increases the overall router delay when small number of VCs and large buffer sizes are used.

C	B	P	F	T_r	T_s	T_c	CC Period
3	8	7	5	7.49	5.88	7.09	7.49
	16			7.49	6.68	7.09	7.49
	24			7.49	7.15	7.09	7.49
	32			7.49	7.48	7.09	7.49
	48			7.49	7.95	7.09	7.95
	64			7.49	8.28	7.09	8.28
4	8	9	7	8.07	6.10	7.34	8.07
	16			8.07	6.90	7.34	8.07
	24			8.07	7.37	7.34	8.07
	32			8.07	7.70	7.34	8.07
	48			8.07	8.17	7.34	8.17
	64			8.07	8.50	7.34	8.50
5	8	11	9	8.50	6.28	7.53	8.50
	16			8.50	7.08	7.53	8.50
	24			8.50	7.54	7.53	8.50
	32			8.50	7.88	7.53	8.50
	48			8.50	8.34	7.53	8.50
	64			8.50	8.68	7.53	8.68
6	8	13	11	8.85	6.42	7.69	8.85
	16			8.85	7.22	7.69	8.85
	24			8.85	7.69	7.69	8.85
	32			8.85	8.02	7.69	8.85
	48			8.85	8.49	7.69	8.85
	64			8.85	8.82	7.69	8.85
96	8.85	9.29	7.69	9.29			

Table 2: Delays for and Adaptive Routing Algorithm for k -ary 2-cube networks

- Varying buffer size affects deterministic routers' clock cycles more than adaptive routers'.

The results show that are 12 to 40 % slower than deterministic routers. To offset this difference, adaptive routers must perform at least 12 to 40 % better than deterministic routers. These results are similar to the results in [2] where 15 % to 60 % improvement is required for f-flat routers with similar number of VCs and under worm-hole routing. The objective of this analysis is how much of this delay can be offset by lower queuing delay in virtual cut-through with a constant router area.

4 Experimental Results

Both deterministic and adaptive routing simulations were performed under all three different traffic conditions using a discrete-time simulator. Simulation

C	B	P	F	T_r	T_s	T_c	CC Period
3	8	10	6	7.80	6.19	7.09	7.80
	16			7.80	6.99	7.09	7.80
	24			7.80	7.46	7.09	7.80
	32			7.80	7.79	7.09	7.80
	48			7.80	8.26	7.09	8.26
	64			7.80	8.59	7.09	8.59
	96			7.80	9.06	7.09	9.06
4	8	13	9	8.50	6.42	7.34	8.50
	16			8.50	7.22	7.34	8.50
	24			8.50	7.69	7.34	8.50
	32			8.50	8.02	7.34	8.50
	48			8.50	8.49	7.34	8.50
	64			8.50	8.82	7.34	8.82
	96			8.50	9.29	7.34	9.29
5	8	16	12	9.00	6.60	7.53	9.00
	16			9.00	7.40	7.53	9.00
	24			9.00	7.87	7.53	9.00
	32			9.00	8.20	7.53	9.00
	48			9.00	8.67	7.53	9.00
	64			9.00	9.00	7.53	9.00
	96			9.00	9.47	7.53	9.47
6	8	19	15	9.39	6.75	7.69	9.39
	16			9.39	7.55	7.69	9.39
	24			9.39	8.02	7.69	9.39
	32			9.39	8.35	7.69	9.39
	48			9.39	8.82	7.69	9.39
	64			9.39	9.15	7.69	9.39
	96			9.39	9.62	7.69	9.62

Table 3: Delays for and Adaptive Routing Algorithm for k -ary 3-cube networks

results were obtained for various k -ary 3-cube and k -ary 2-cube networks. The simulation uses a stabilization threshold of a 0.005 difference between throughput 1000 clock cycles apart to determine steady state. Message sizes varied from 8 to 32 flits and throughput from 0.1 until saturation was reached in 0.1 increments. The buffer sizes simulated for adaptive and deterministic routing are one, two and three times the message size. The number of VCs used by the adaptive router varies from three to six VCs per dimension per node, while the number used by the deterministic router is always two. The simulator implements a back-pressure mechanism which results in a negative slope of the latency versus throughput plots at higher loads. The following notation is used throughout this section: L: message size in flits, B: buffer size in flits, C: number of VCs per PC, D: deterministic routing, and A: adaptive routing.

The results in this section are grouped as follows: Section 4.1 discusses the trade-offs between the num-

Area	L	B	C	Ran	Bit	Per
48	8	8	6	0.320	0.290	0.275
		16	3	0.330	0.289	0.258
96	16	16	6	0.319	0.302	0.277
		32	3	0.345	0.300	0.258
192	32	32	6	0.320	0.301	0.280
		64	3	0.313	0.275	0.241

Table 4: Throughput saturation point (flits/ns/node) as function of area for a 10-ary 3-cube network

ber and size of the VCs in adaptive routers. Section 4.2 reports on the comparison of dimension-order and adaptive routing with fixed size buffers. Section 4.3 describes the characteristics of both types of routers that achieve similar saturation points for small message sizes.

Type	Area	L	B	C	Ran	Bit	Per
Deter	128	32	64	2	0.236	0.219	0.129
Adapt			32	4	0.316	0.278	0.264
Deter	192	32	96	2	0.245	0.233	0.130
Adapt			32	6	0.320	0.301	0.280
			64	3	0.313	0.275	0.241

Table 5: Constant area throughput saturation point (flits/ns/node) for deterministic and adaptive routing in a 10-ary 3-cube network

4.1 Constant Area Adaptive Routers

Constant area is defined here as keeping the sum of all VC buffer sizes constant. The constant area results for various message lengths and the three types of traffic are shown in Figures 2, 3, and 4, and the corresponding saturation points are shown in Table 4. In this analysis the saturation point is defined as

Type	Area	L	B	C	Ran	Bit	Per
Deter	48	8	24	2	0.272	0.257	0.143
Adapt	24		8	3	0.276	0.258	0.219
Deter	96	16	48	2	0.264	0.245	0.144
Adapt	48		16	3	0.289	0.258	0.229
Deter	192	32	96	2	0.245	0.233	0.130
Adapt	96		32	3	0.288	0.259	0.231

Table 6: Throughput saturation points (flits/ns/node) in a 10-ary 3-cube network with deterministic routing having twice the area as adaptive

the last throughput at which the *accepted* load in the network was $\pm 1.5\%$ of the *offered* load.

The effects of traffic pattern on the saturation points shows that for all message sizes, the highest saturation points obtained are always for uniform random traffic patterns while the lowest saturation points are always obtained under perfect shuffle traffic. While the simulations performed under complement traffic have 5 % to 13 % lower saturation points than the uniform random traffic simulations, perfect shuffle traffic exhibits 13 % to 26 % lower saturation points.


In addition, the results show that under random uniform traffic having fewer VCs always produces better performance. However, for less uniform traffic, a tradeoff exists between the size and number of VCs: at low throughput, fewer VCs result in a lower message latency. At high throughput, fewer VCs result in a lower saturation point. This phenomenon can be explained by the fact that at low throughput the interleaving of the messages among many VCs results in a higher latency.

However, at high throughput, message blocking becomes more frequent and the additional adaptive VCs can alleviate the congestion and reduce the blocking time of messages. Therefore, the larger clock cycle times required for numerous VCs pays off in the form of higher saturation points. Under random uniform traffic much interleaving of messages is not required for high saturation points. For less uniform traffic, the greater number of VCs (more interleaving) is crucial for higher saturation points.


In [13] Dally noted a similar phenomenon for dimension-order worm-hole routing in k -ary n -cube networks. However, by accounting for the actual clock cycle times, the difference between a small and a large number of VCs at low throughput is more pronounced. This difference is also obviously accentuated by the message size.

4.2 Constant Area Deterministic vs. Adaptive Routers


The constant area results for various message lengths and different traffic patterns are shown in Figures 5, 6, and 7, and the corresponding saturation points are shown in Table 5. These results show the tradeoff between deterministic and adaptive routers under all three types of traffic: at low throughput, deterministic routing almost always results in lower message latency. At high throughput, adaptive routing always achieves higher saturation points. Note that deterministic routing has a particularly poor performance under perfect shuffle traffic due to its very



(a) Buffer area = 48 flits and $L = 8$ flits




(b) Buffer area = 96 flits and $L = 16$ flits




(c) Buffer area = 192 flits and $L = 32$ flits


Figure 2: Adaptive routing latency (10-ary 3-cube) under random uniform traffic




(a) Buffer area = 48 flits and $L = 8$ flits




(a) Buffer area = 48 flits and $L = 8$ flits




(b) Buffer area = 96 flits and $L = 16$ flits



(b) Buffer area = 96 flits and $L = 16$ flits



(c) Buffer area = 192 flits and $L = 32$ flits




(c) Buffer area = 192 flits and $L = 32$ flits

Figure 3: Adaptive routing latency (10-ary 3-cube) under complement traffic


Figure 4: Adaptive routing latency (10-ary 3-cube) under perfect shuffle traffic

localized traffic pattern.

Deterministic routing almost always results in lower message latency at low channel throughput because the router cycle time dominates the performance. At high throughput, adaptive routing results in a higher saturation point because the performance is now dominated by routing freedom which results in lower queuing delays. The only exception to this behavior is when buffer area equals 192. For each traffic type, adaptive routing performs either better than or comparable to deterministic routing when $B=64$ and $C=3$ even at low throughput. Under these circumstances, the additional delay due to increased buffer sizes for the deterministic router offsets its previous good performance at low throughput.




(a) Buffer area = 48 flits and $L = 8$ flits




(b) Buffer area = 192 flits and $L = 32$ flits

Figure 5: Adaptive and Dimension-Order routing latency (10-ary 3-cube) under constant buffer area and random uniform traffic



(a) Buffer area = 48 flits and $L = 8$ flits




(b) Buffer area = 192 flits and $L = 32$ flits


Figure 6: Adaptive and Dimension-Order routing latency (10-ary 3-cube) under constant buffer area and complement traffic

4.3 Comparable Performance Deterministic and Adaptive Routers


In this section, the performance of deterministic and adaptive routers with different total buffer areas are compared to determine if dimension-order routing can ever give similar performance to adaptive routing. The results are shown in Figures 8, 9, and 10 for various message lengths and different traffic patterns. The corresponding saturation points are shown in Table 6. The results for all three traffic types show that a deterministic router with $B_d = 3L$ and $C_d = 2$ gives about the same saturation point as an adaptive router with $B_a = L$ and $C_a = 3$ when message length is small




(a) Buffer area = 48 flits and $L = 8$ flits



(a) $L = 8$ flits



(b) Buffer area = 192 flits and $L = 32$ flits




(b) $L = 16$ flits

Figure 7: Adaptive and Dimension-Order routing latency (10-ary 3-cube) under constant buffer area and perfect shuffle traffic

and traffic is relatively uniformly random. As message size increases and traffic patterns become less random, adaptive routing becomes better at all throughput values. Once again this is because the increased buffer space offsets the previous low delay associated with deterministic routers.


4.4 Discussion of Results

The results in this paper show that although adaptive routers have longer cycle times than deterministic routers, the added delay is offset at high loads by the adaptive router's ability to provide more degrees of freedom in routing. However, a tradeoff does exist be-




(c) $L = 32$ flits


Figure 8: Latency performance for dimension-order and adaptive routing on a 10-ary 3-cube network under random uniform traffic and when deterministic router buffer area is twice as great as the adaptive router buffer area




(a) $L = 8$ flits




(a) $L = 8$ flits




(b) $L = 16$ flits



(b) $L = 16$ flits



(c) $L = 32$ flits



(c) $L = 32$ flits

Figure 9: Latency performance for dimension-order and adaptive routing on a 10-ary 3-cube network under complement traffic when deterministic router buffer area is twice as great as the adaptive router buffer area

Figure 10: Latency performance for dimension-order and adaptive routing on a 10-ary 3-cube network under perfect shuffle traffic when deterministic router buffer area is twice as great as the adaptive router buffer area

tween the two router models. While adaptive routers give higher saturation points deterministic routers give lower message latency at low throughput under uniform random traffic. The cross-over point, where the two latencies are equal, depends, to a large extent, on the number of VCs used in the adaptive routing, buffer size, and traffic pattern.

Among adaptive routers with various total buffer area, a tradeoff in performance exists as well. At low load, fewer VCs result in a lower message latency because the benefit of message interleaving with many VCs is not necessary at low loads and the extra delay required by many VCs is not offset by its performance. At high load, fewer VCs result in a lower saturation point.

The empirical comparison of the performance of deterministic and adaptive routers with different areas, shows that a deterministic router with twice the total buffer area of an adaptive one achieves about the same saturation point for short messages and under relatively random uniform traffic. All of these results have shown that adaptive routers can be worth their extra complexity especially with large message sizes and under non-uniform traffic.

5 Related Work

Much of the research on adaptive routing has been done in the context of developing deadlock-free routing algorithms. Some of the strategies for preventing deadlock that have been proposed include variations on the use of VCs [11, 27, 5, 16, 31], and other strategies involving restrictions on the choice of paths or involving the use of additional hardware [19, 8, 24]. [29, 26, 21] involve deadlock freedom for worm-hole switching, randomization techniques are used in [23], and deadlock recovery is performed in [32]. A universal proof technique for deadlock freedom is given in [30]. Fault tolerant adaptive routing techniques have also been discussed and include [21, 14, 27]. Some studies have addressed the effects of other design parameters on adaptive routing performance, including [3, 22, 6, 25].

In most of these studies, however, the emphasis is on the algorithm design and implementation of the router algorithm is secondary. [12] uses implementation techniques for deadlock-free dimension-order worm-hole routing based on full crossbar architectures. Later optimizations were done, such as partitioning crossbars [10, 18]. [1] analyzes the effect of message length and locality on the performance of networks of varying sizes and dimensions for dimension-

order worm-hole routing when both switch and wire delays are accounted for.

Adaptive router implementations in [2, 7] are compared to deterministic routers. However, the comparisons do not account for the advantage of adaptive routers in reducing queuing delays in the nodes between source and destination. In addition, the various routing algorithms evaluated, both deterministic and adaptive, require a variable amount of resources such as buffer area or physical channels between nodes. In [17], this advantage is taken into account but only for worm-hole routing.

6 Conclusions

This paper reports on the latencies of adaptive and deterministic routers under constant area constraints. The results indicate that a tradeoff does exist between constant area adaptive routers under non-uniform traffic patterns. At low throughput, fewer VCs result in a lower message latency. At high throughput, fewer VCs result in a lower saturation point. However, for uniform traffic, adaptive routers with fewer VCs always perform better.

Deterministic routers have a lower message latency than adaptive routers with the same area under uniform traffic, small message and buffer sizes, and low throughput. They also have a higher saturation point under the same conditions. However, for large message and buffer sizes, adaptive routing becomes performs better at all throughput values for all types of traffic.

Finally, it was shown that a deterministic router can achieve a latency and saturation point similar to an adaptive one with half the total area for small message size and with random or complement traffic.

Because the adaptive routing technique used here does not perform as well at low throughput as the deterministic routing technique under numerous circumstances, future work could include improving this aspect. One method in which this could be done is implementing both deterministic and adaptive routing in such a way that good performance is obtained at low as well as high throughput. One type of routing strategy which does just this is presented in [28].

References

- [1] A. Agarwal. Limits on interconnection network performance. *IEEE Trans. on Parallel and Distributed Systems*, 2(4):398–412, October 1991.

- [2] K. Aoyama and A. Chien. The cost of adaptivity and virtual lanes in wormhole router. *J. of VLSI Design*, 2(4), 1995.
- [3] D. Badouel, C. A. Wuthrich, and E. L. Fiume. Routing strategies and message contention on low-dimensional interconnection networks. Technical report, Computer System Research Institute, University of Toronto, 1991.
- [4] P. Berman, L. Gravano, G. Pifarre, and J. Sanz. Adaptive deadlock and livelock free routing with all minimal paths in torus networks. In *Proc. of the Symp. on Parallel Algorithms and Architectures*, pages 3–12, 1992.
- [5] R. Boppana, S. Chalasani, and C.S. Raghavendra. Adaptive packet-routing algorithms for k -ary n -cubes. 1992.
- [6] R. V. Boppana and S. Chalasani. A comparison of adaptive wormhole routing algorithms. In *Int. Symp. on Computer Architecture*, pages 351–360, 1993.
- [7] A. Chien. A cost and speed model for k -ary n -cube wormhole routers. In *IEEE Proc. of Hot Interconnects*, Aug. 1993.
- [8] A. A. Chien and J. H. Kim. Planar-adaptive routing: low cost adaptive networks for multiprocessors. *Int. Symp. on Computer Architecture*, pages 268–277, May 1992.
- [9] W. Dally, A. Chien, and et al. The J-Machine: a fine-grain concurrent computer. In *Proc. of the IFIP Congress*, pages 1147–1153, Aug. 1989.
- [10] W. Dally and P. Song. Design of a self-timed VLSI multicomputer communication controller. In *Proc. of the Int. Conf. on Computer Design*, pages 230–40, 1987.
- [11] W. J. Dally. Virtual-channel flow control. *IEEE Trans. on Computers*, 3(2):194–205, March 1992.
- [12] W. J. Dally and C. L. Seitz. The torus routing chip. *J. Dist. Computing*, 1(3):187–196, 1986.
- [13] W. J. Dally and C. L. Seitz. Deadlock-free message routing in multiprocessor interconnection networks. *IEEE Trans. on Computers*, C-36(5):547–553, May 1987.
- [14] B. Dao, J. Duato, and S. Yalamanchili. Configurable flow control mechanisms for fault-tolerant routing. In *Int. Symp. on Computer Architecture*, pages 220–9, 1995.
- [15] J. Duato. Deadlock-free adaptive routing algorithms for multicomputers: Evaluation of a new algorithm. In *Proc. of the 3rd IEEE Symp. on Parallel and Distributed Processing*, Dec. 1991.
- [16] J. Duato. A new theory of deadlock-free adaptive routing in wormhole networks. *IEEE Trans. on Parallel and Distributed Systems*, 4(12):1320–1331, December 1993.
- [17] J. Duato and P. Lopez. Performance evaluation of adaptive routing algorithms for k -ary n -cubes. In *Parallel Computer Routing and Communication*, pages 45–59, 1994.
- [18] C. Flaig. VLSI mesh routing systems. Master's thesis, California Institute of Technology, May 1987.
- [19] C. L. Glass and L. M. Ni. The turn model for adaptive routing. In *Int. Symp. on Computer Architecture*, pages 278–287, 1992.
- [20] P. Kermani and L. Kleinrock. Virtual cut-through: a new computer communication switching technique. *Computer Networks*, 3:267 – 286, 1979.
- [21] J. Kim, Z. Liu, and A. Chien. Compressionless routing. In *Proc. of ISCA*, Apr. 1994.
- [22] J. H. Kim and A. A. Chien. The impact of packetization in wormhole-routed networks. In *PARLE '93*, pages 242–253, 1993.
- [23] S. Konstantinidou and L. Snyder. Chaos router: architecture and performance. *Int. Symp. on Computer Architecture*, pages 212–221, 1991.
- [24] G. Denicolay L. Gravano, G. D. Pifarre and J. L. C. Sanz. Adaptive deadlock-free wormhole routing in hypercubes. In *Fifth Int'l. Parallel Processing Symp.*, pages 512–517, Beverly Hills, CA, 1992.
- [25] Annette Lagman. *Modelling, Analysis and Evaluation of Adaptive Routing Strategies*. PhD thesis, Colorado State University, Computer Science Department, November 1994.
- [26] X. Lin and L. M. Ni. Deadlock-free multicast wormhole routing in multicomputer networks. In *Int. Symp. on Computer Architecture*, pages 116–125, 1991.
- [27] D. H. Linder and J. C. Harden. An adaptive and fault tolerant wormhole routing strategy for k -ary n -cubes. *IEEE Trans. on Computers*, 40(1):2–12, January 1991.
- [28] D. Miller and W. Najjar. Preliminary evaluation of a hybrid deterministic/adaptive router. In *Parallel Computer Routing and Communication*, 1997.
- [29] L. M. Ni and P. K. McKinley. A survey of wormhole routing techniques in direct networks. *IEEE Computer*, pages 62–76, 1993.
- [30] L. Schwiebert and D. Jayasimha. A universal proof technique for deadlock-free routing in interconnection networks. In *7th Annual Symp. on Parallel Algorithms and Architecture*, pages 175–84, 1995.
- [31] C-C. Su and K. G. Shin. Adaptive deadlock-free routing in multicomputers using only one extra virtual channel. In *Int. Conf. on Parallel Processing*, 1993.
- [32] Anjan K. V. and T. Pinkston. An efficient, fully adaptive deadlock recovery scheme: DISHA. In *Int. Symp. on Computer Architecture*, pages 201–10, 1995.