

Evaluating Elastic Buffer and Wormhole Flow Control

George Michelogiannakis, *Student Member, IEEE*,
Daniel U. Becker, and
William J. Dally, *Fellow, IEEE*

Abstract—With the emergence of on-chip networks, router buffer power has become a primary concern. Elastic buffer (EB) flow control utilizes existing pipeline flip-flops in the channels to implement distributed FIFOs, eliminating the need for input buffers at the routers. EB routers have been shown to be more efficient than virtual channel routers, as they do not require input buffers or complex logic for managing virtual channels and tracking credits. Wormhole routers are more comparable in terms of complexity because they also lack virtual channels. This paper compares EB and wormhole routers and explores novel hybrid designs to more closely examine the effect of design simplicity and input buffer cost. Our results show that EB routers have up to 25 percent smaller cycle time compared to wormhole and hybrid routers. Moreover, EB flow control requires 10 percent less energy to transfer a single bit through a router and offers three percent more throughput per unit energy as well as 62 percent more throughput per unit area. The main contributor to these results is the cost and delay overhead of the input buffer.

Index Terms—On-chip interconnection networks, interconnection architectures.

1 INTRODUCTION

SCALABLE networks-on-chip [1] have been developed to serve the communication requirements of chips with numerous processors, memory elements or other logic blocks. Past implementations have accounted for up to 30 percent of the total power consumption of the Intel 80-core Terascale chip [2] and 40 percent of the MIT RAW chip [3]. Compared to off-chip networks, wires are available in abundance in the on-chip environment, while buffer cost is more significant [1]. Buffers are commonly implemented in network routers and used by the flow control scheme to enqueue contenting packets or flits [4]. Buffers can occupy as much as 75 percent of the network area, as in the TRIPS prototype [5]. Technology constraints can significantly affect buffer implementation and cost.

There have been multiple proposals to eliminate these buffers in order to reduce area and power consumption. In such bufferless flow control schemes, contenting packets or flits are either dropped and retransmitted by their source [6] or deflected to a free output port [7]. Since this penalizes network performance and energy consumption under high loads, these schemes are aimed at lightly loaded networks. However, these networks can still be designed to provide guaranteed throughput and multicast support [8].

In addition, hybrid schemes combining wormhole and circuit-switching flow controls have been proposed [9], [10]. They use circuit switching to establish connections between frequent communication pairs and bypass intermediate router pipelines and buffers. Express virtual channel (VCs) [11] accomplish the same goal by defining multihop express paths in the network.

Elastic buffer (EB) flow control was recently proposed to eliminate router buffers while preserving buffering in the

network [12]. Pipeline flip-flop (FFs) become EBs with two storage locations through the addition of a small logic block which controls their master and slave latch enable inputs independently. The two inputs are still gated by the clock as in master-slave FFs. EBs can be implemented as custom cells to increase efficiency and guarantee that the two latches can never be enabled simultaneously. Having EBs in sequence enables channels to act as distributed FIFOs. Flits progress among EBs using a ready-valid handshake similar to the valid-stall handshake of [13]. Therefore, channels are used for buffering and router buffers are removed, eliminating the associated area and energy costs. EB networks trade off these savings for a wider datapath that increases their throughput and makes them more energy and area efficient than VC networks [12], [14]. Fig. 1 shows an illustration of an EB.

EB routers have the area and energy benefits of circuit-switched routers, without the latency and energy overhead of setting up and tearing down circuits. They do not require credits, and only need simple per-output arbiters instead of full-fledged allocators, as each input may only request a single output. For 5×5 2D mesh routers, the baseline two-stage router presented in [12] reduces cycle time by 18 percent, area by 77 percent, and dynamic power by 95 percent compared to a similar VC router with two VCs, eight buffer slots statically assigned to each, and buffers implemented as FF arrays. Assuming equal clock frequencies for both and efficient custom SRAM input buffers, EB networks provide up to 12 percent more throughput per unit power.

Further work has proposed improved EB router designs in the form of the enhanced two-stage and the single-stage routers [14]. The enhanced two-stage router, shown in Fig. 2, improves the baseline design of [12] by reducing the cycle time by 42 percent. This is accomplished by using a two-slot output EB instead of a three-slot output EB which has to be implemented as a FIFO, and by using look-ahead routing [15] to precompute routing decisions. Therefore, routing computation is performed in parallel with arbitration, removing routing from the critical path of the first pipeline stage. On the other hand, the single-stage router merges the two pipeline stages of the enhanced two-stage router in order to avoid pipelining overhead and to optimize latency; it is illustrated in Fig. 3. The single-stage router reduces the required energy per bit by 29 percent and the zero-load latency in clock cycles by 24 percent compared to the enhanced two-stage router.

This paper compares EB and wormhole routers [16], [17]. Although wormhole routers still feature input buffers and credit-based flow control, they lack VCs, and thus do not require VC and switch allocators. Instead, like EB routers, they use simple output arbiters. Their credit handling logic is also simpler than that of VC routers. Consequently, their design complexity is more comparable to EB routers. In this paper, we investigate whether the gains reported in [12] are mostly attributed to design simplicity due to removing VCs and allocation, or input buffer cost. Because accurately modeling timing paths and buffer cost is crucial, we base our analysis on fully placed and routed standard-cell implementations.

Furthermore, we investigate and evaluate novel hybrid EB-wormhole router designs, which add input buffers to the enhanced two-stage and single-stage EB routers. These routers still use the channels as distributed FIFOs, but provide additional buffering. We examine the trade-off between the cost of the added input buffer and the associated increase in throughput in a simple router design without VCs. Our results demonstrate how adding extra input buffering in EB networks affects area and power efficiency.

Our results for an 8×8 2D mesh with dimension-order routing (DOR) show that the enhanced two-stage router on average has a 25 percent smaller cycle time. The single-stage EB router offers three percent more throughput per unit energy, 62 percent more throughput per unit router area, and has a 10 percent lower zero-load latency per unit throughput compared to the most efficient

• The authors are with the Electrical Engineering Department, Stanford University, Stanford, CA 94305.
E-mail: {mihelog, dub, dally}@cs.stanford.edu.

Manuscript received 16 Feb. 2010; revised 17 June 2010; accepted 14 Oct. 2010; published online 6 Dec. 2010.

Recommended for acceptance by L. John.

For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number TC-2010-02-0114. Digital Object Identifier no. 10.1109/TC.2010.243.

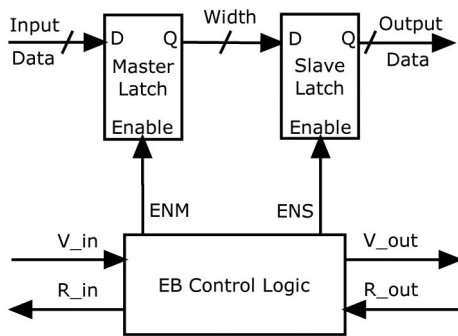


Fig. 1. An elastic buffer.

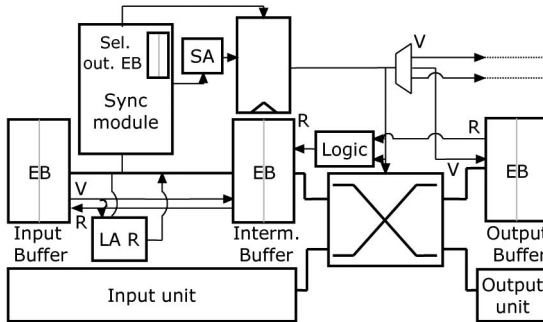


Fig. 2. The enhanced two-stage EB router.

FIFO-based router (i.e., wormhole or hybrid) for each case; results are averaged across all datapath widths we evaluated.

The FIFO is the dominant factor for all results. Buffer read was part of the critical path in most configurations. Moreover, for a 20 percent flit injection rate and by averaging among the datapath widths we evaluated, the FIFO contributes 56 and 60 percent of the area and 21 and 17 percent of the power of the wormhole and hybrid routers, respectively. Furthermore, the credit handling logic adds complexity and has a higher processing delay in the upstream router, which causes flits to wait an extra cycle before departing from a buffer after congestion is alleviated at the next downstream router.

EB routers are more efficient than FIFO-based routers in all aspects. Among the EB routers, the single-stage router is preferable in terms of throughput efficiency, router area, energy, and zero-load latency. Trading off energy and area savings for an increased datapath width provides greater benefits than the enhanced two-stage router's smaller cycle time for the above aspects. Therefore, we show that EB flow control is more efficient than wormhole flow control due to the FIFO overhead.

In summary, the main contributions of this paper are:

- We compare the EB routers of [14] against the wormhole router. We show that the EB routers are more efficient than the wormhole router for all comparison aspects. Thus, the dominant factor for EB network gains is the cost and delay overhead associated with the input buffer.
- We explore hybrid EB-wormhole router designs, and show that adding extra input buffering in EB networks reduces efficiency.
- For all conclusions, we offer insight and isolate the underlying causes. Specifically, the primary contributor to the above results is the input FIFO complexity and cost.

2 HYBRID EB-WORMHOLE ROUTERS

EB routers already feature an EB at their inputs for pipelining [12], [14]. That input EB essentially is a two-slot input FIFO.

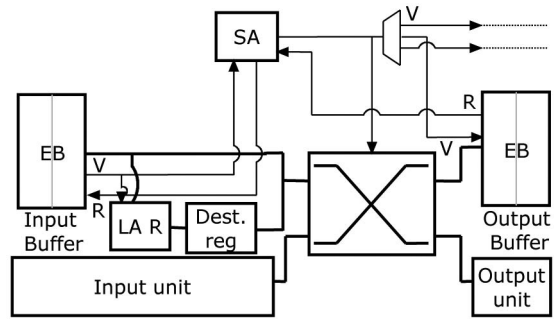


Fig. 3. The single-stage EB router.

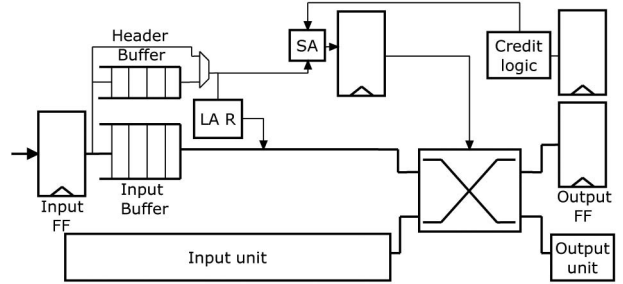


Fig. 4. The wormhole router.

Wormhole routers [17], on the other hand, typically have significantly deeper input FIFOs; in particular, with credit-based flow control, the FIFOs need to be at least deep enough to cover the credit round-trip and processing delay in order to avoid channel underutilization.

The wormhole router implemented for this study is shown in Fig. 4; we use a two-stage pipeline in order to avoid excessive critical path delay. The credit handling logic is placed at each output port and drives a single-bit signal to the output's arbiter, inhibiting grants if no credits are available. For energy efficiency reasons, our router design does not have an intermediate pipeline register between the input buffer and the crossbar.¹ Switch arbitration in a given cycle sets up the crossbar control signals for the next cycle; hence, when a tail flit is at the head of the buffer and begins crossbar traversal, the control signals driving the switch arbitration logic must be generated from the head flit behind it. In order to be able to handle this situation without adding either a second read port to the input buffer or a stall cycle between successive packets, we track routing information for each packet in a separate header buffer. This buffer's small width allows us to choose an implementation that is optimized for fast read-out at the cost of slightly increased energy, and to thus minimize the delay going into the switch arbiter logic. As with the baseline EB designs, we include input- and output-side register stages to ensure that the full clock cycle is available to the adjacent channel segments for signal propagation. Without such registers, the FIFO write delay caused by address decoding, internal fan out to the individual storage elements, and their respective setup time requirements-would have to be borrowed from the preceding channel segment, and both the FIFO's read delay and the crossbar traversal delay would have to be borrowed from the subsequent one; this reduces the time available for signal propagation and thus the maximum channel length. Since the transfer of flits from the input register to the FIFO is overlapped with switch allocation, it does not incur additional pipeline delay.

1. Note that even with such a register, allowing the flit buffer to be fully bypassed adversely affects timing, as the bypass logic depends on the outcome of switch arbitration, which in turn is not available until late in the cycle.

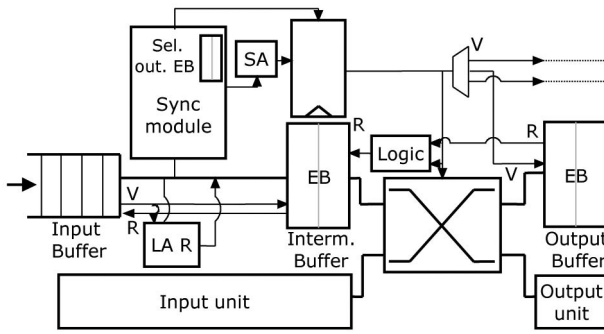


Fig. 5. Two-stage hybrid EB-wormhole router.

The hybrid EB-wormhole router designs we propose in this study replace the input EB with a FIFO. This two-stage hybrid router design is illustrated in Fig. 5. Credit-based flow control is replaced with the ready-valid handshake of EB channels. Therefore, the input FIFO also interfaces using the ready-valid handshake. The illustrated design has an intermediate EB and synchronization logic functioning in the same way as the enhanced two-stage router [14]. While the intermediate EB is not required, removing it would require the router to read from FIFO locations other than the head or have a separate structure as described above. Therefore, the total number of flits that the hybrid router can store at each input is given by the number of entries in the FIFO plus two for the intermediate EB. Unlike the baseline EB and wormhole designs, this hybrid design writes directly from the channel into the input buffer, reducing energy consumption at the cost of additional timing pressure on the preceding channel segment. While this favors the hybrid router in terms of energy and cost because we ignore the extra time borrowed from the channel, we show in Section 3.2 that at least one of the two baseline EB routers is still preferable to it in each case.

Further hybrid designs are possible. For example, the input EB can be preserved in addition to the FIFO, as illustrated in Fig. 6. This makes the FIFO smaller for the same amount of total input buffering. It also enables flits that do not face contention to bypass the FIFO and be stored directly to the intermediate EB if the FIFO is empty. However, if the FIFO cannot be bypassed, additional energy is expended for traversing the extra input EB and multiplexing logic. On the other hand, if the FIFO is bypassed, it may be small enough to have a comparable energy overhead with the input EB and the bypassing logic complexity. Therefore, we do not use buffer bypassing in this work for the wormhole and hybrid routers. However, buffer bypassing may be of interest to lightly loaded networks depending on their buffer sizes.

To fully explore the design space, we also briefly consider hybrid routers based on the single-stage router.

3 EVALUATION

This section begins by outlining our evaluation methodology in Section 3.1, and then presents our results in Section 3.2.

3.1 Methodology

Implementation results were obtained by synthesizing a single instance of each router design using Synopsys Design Compiler and performing place and route (PnR) using Cadence Silicon Encounter. All data points were extracted for routers placed and routed for their minimum clock cycles. This is a design choice which prioritizes cycle time, instead of energy or area. The minimum clock cycle for each router was determined after performing static timing analysis with post-PnR parasitics. Similarly, energy per transferred bit was calculated by driving the post-PnR netlists with pseudorandom input traffic using a test environment separate from the cycle-accurate network simulator that we discuss below.

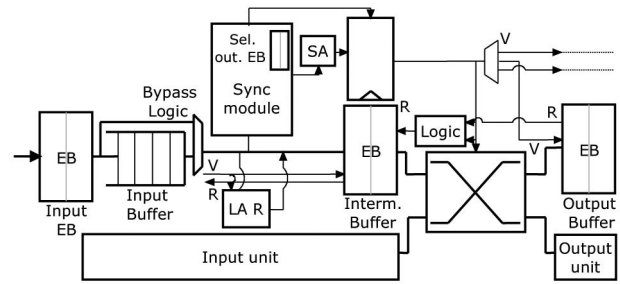


Fig. 6. Two-stage hybrid EB-wormhole router with input EB.

Low-power optimizations, such as clock gating the FIFO FFs, were automatically applied by the synthesis and PnR tools.

We used a commercial 45 nm low-power technology library under worst-case conditions. The initial floorplan utilization was set to 70 percent. Primary input and output driving strengths, loads, and timing constraints were specified to realistically model network channels at the router ports. Due to technology usage constraints, the FIFOs were implemented from FF arrays. We assumed an 8×8 2D mesh network using radix-5 routers with a single network terminal attached to each. Router ports were placed in the floorplan according to the inter-router connections of the assumed network. Deterministic DOR was used. Round-robin arbiters were used for switch arbitration. The switch was implemented using multiplexers.

The network throughput and latency data points were generated using a modified version of Booksim [18], a cycle-accurate network simulator. The packet size was held constant at 512 bits. No communication protocol was assumed; therefore, we used a single physical network defining a single traffic class. Unless otherwise specified, for each data point we used the maximum achievable clock frequency and corresponding area and energy results from PnR for each router and datapath width. However, we include comparisons for equal clock frequencies to account for scenarios where the frequency is limited by external factors. In these cases, we use the same netlists optimized for minimum cycle time.

All network channels had a propagation delay of two cycles and were clocked at the same frequency as each router; i.e., each channel had one FF or EB midway between the two connected routers. Thus, EB networks with the enhanced two-stage router had four EBs (eight buffer slots) between consecutive router switches, while networks with the single stage router had three. The wormhole router input FIFO also had eight buffer slots; seven of these were necessary to cover the buffer turnaround delay, and adding an extra entry to bring the buffer size to a power of two allowed us to simplify buffer control logic. The hybrid router was configured with a six-entry FIFO in addition to its intermediate EB to also provide a total of eight buffer slots. In order to maintain the aforementioned properties, we kept the number of buffer slots constant when adjusting the datapath width.

For comparisons with equalized throughput, area or energy, Pareto-optimal curves were generated. Area and energy were calculated for a single router of each type, and thus do not account for network channels. Since the number of wires in a channel is the same for ready-valid handshake and credit-based handshake, the difference between these two cases in the place and route flow would be minimal. The datapath width was swept from 29 to 171 bits such that packets consisted of 3-18 flits; widths were chosen such that fragmentation was avoided. Each flit consisted of a single phit. We used a set of traffic patterns comprising uniform random, random permutation, shuffle, bit complement, tornado, and neighbor traffic [18]. The maximum throughput is averaged across all traffic patterns. Percentage summaries were calculated by measuring the average distance between corresponding data points.

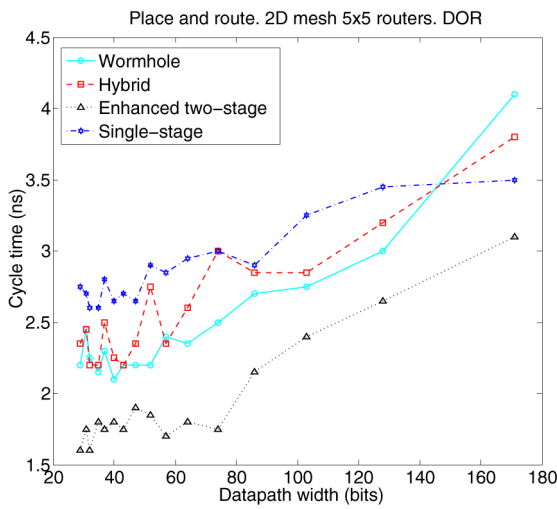


Fig. 7. Router cycle time after PnR.

3.2 Results

Figs. 7, 8, and 9 illustrate the PnR implementation results. The curves are not smooth as a result of the heuristic algorithms that EDA tools use to perform optimizations on discrete values, e.g., for performing cell sizing.

Comparison results of the single and enhanced two-stage EB routers are consistent with [14]. The wormhole and hybrid routers have a 27 and 34 percent larger average cycle time compared to the enhanced two-stage router, respectively. Likewise, they require 1 and 21 percent more energy per bit and occupy 2.3 and 1.5 times the area compared to the single-stage router. These comparison results are averaged across all datapath widths.

The enhanced two-stage router requires the most energy because it was placed and routed to meet its small cycle time. This increases the sizing of cells in the router. The critical paths of the wormhole and hybrid routers begin at the FIFO or EB read, go through the switch, and terminate at the output EB or register. The hybrid router occupies less area than the wormhole router because its FIFO, which is the dominant factor, is smaller by two slots. However, it requires more energy per bit compared to the wormhole router.

To further illustrate the FIFO overhead, Figs. 12 and 13 show the area and power breakdown for the hybrid and wormhole routers. The single-stage router is included for comparison. The

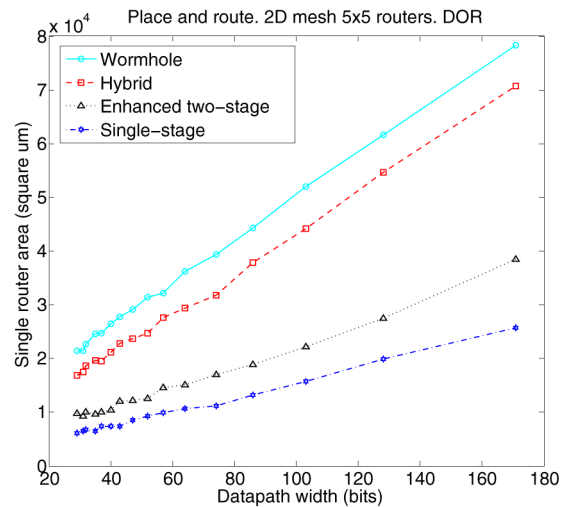


Fig. 9. Router area after PnR.

input module bars include routing computation, credit handling as well as the intermediate EB, and associated logic for the hybrid router. The output module bars include the FFs or EBs between the crossbar and the channel. They also include credit handling logic located at the output side. All other control logic, including arbitration, is included in the “other” bars.

As shown, 56 percent of the area and 21 percent of the power of the wormhole router were in the FIFO. The FIFO in a similar hybrid router constituted 60 percent of the area and 17 percent of the power. The credit logic was the primary contributor for the increase in input and output area for the wormhole router. On the other hand, using two library latch cells instead of a single FF cell causes the increase in input and output power for the hybrid and EB routers.

Table 1 contains the cell, gate, and net counts of the four routers with a 64 bit datapath, and shows the relative differences compared to the single-stage router. The input FIFOs lead to significant increases compared to the EB routers. The difference between the wormhole and hybrid routers is not significant. However, this is not true for the number of gates because of the two extra FIFO slots and the credit handling logic in the wormhole router.

The PnR results highlight the adverse effects of the input FIFO in terms of area, energy, and cycle time. This is especially apparent when comparing the enhanced two-stage and hybrid routers because they differ only in that the input EB is replaced by a FIFO. This demonstrates that explicitly adding FIFOs in the wormhole routers carries a significant additional cost compared to using existing channel FFs as EBs.

Figs. 10 and 11 plot injection rate versus latency for uniform traffic. They assume an equal datapath width for all routers. For equal clock frequencies, the single-stage router has a 19 percent reduced zero-load latency compared to the other three routers. However, assuming that each router operates at its maximum clock frequency, the enhanced two-stage router has a 25 percent reduced zero-load latency—measured in absolute time—compared

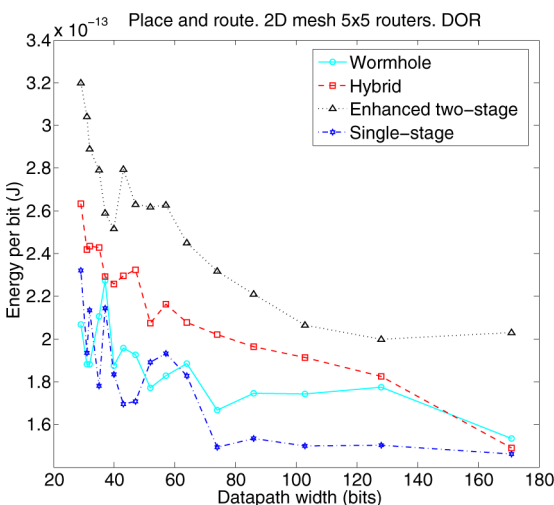


Fig. 8. Energy per transferred bit after PnR.

TABLE 1
PnR Metrics

	1-stage	Enh. 2-stage	Hybrid	Wormhole
Cells	4499	6353 (+41%)	8977 (+100%)	8813 (+96%)
Gates	10073	14247 (+41%)	27710 (+175%)	34176 (+239%)
Nets	4080	5979 (+47%)	9249 (+127%)	8550 (+110%)

64 bit datapath. 5 × 5 mesh routers with DOR.

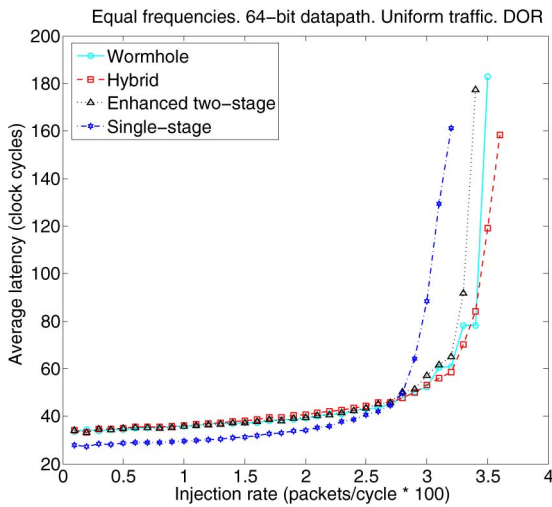


Fig. 10. Injection rate versus latency (equal frequencies).

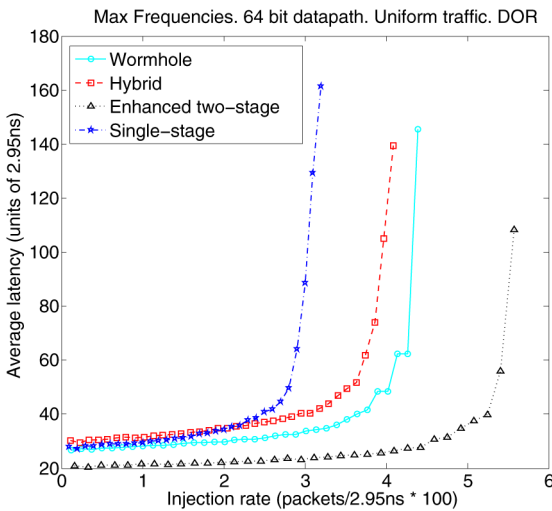


Fig. 11. Injection rate versus latency (max. frequencies).

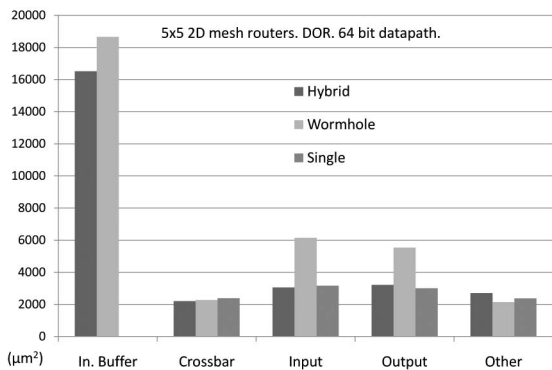


Fig. 12. Area breakdown after PnR.

to the single-stage router. The wormhole and single-stage routers have comparable zero-load latencies.

The maximum throughput of each router for a given datapath width and clock frequency is primarily affected by the number of available buffer slots. This accounts for the 13 and 6 percent increased throughput of the wormhole router compared to the single-stage and enhanced two-stage routers, respectively, when running at the same clock frequency. On the other hand, while both credits and ready signals have the same propagation delay

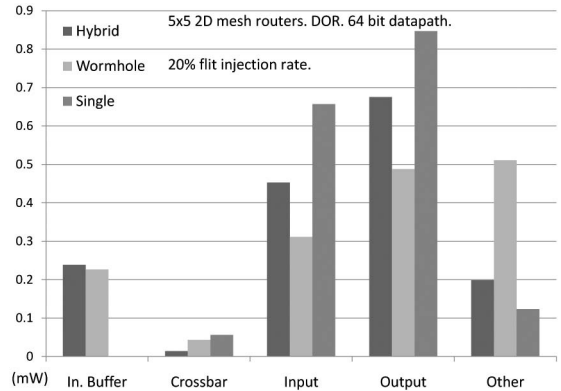


Fig. 13. Power breakdown after PnR.




Fig. 14. Maximum throughput by traffic pattern.

in cycles, the effective buffer turnaround time is one cycle higher for the wormhole router as credits are consumed during the switch arbitration stage one cycle before the flit actually leaves the buffer; this effectively increases buffer occupancy. Furthermore, the EB routers effectively provide additional buffer capacity in the form of the output EB. Together, these factors increase the hybrid router’s maximum throughput beyond that of the wormhole router. If routers operate at their maximum frequencies, the enhanced two-stage router has a 27 and 36 percent higher maximum throughput—measured in absolute time—compared to the wormhole and hybrid router, respectively. This illustrates that the difference in cycle time has a significant impact. However, this comparison does not take into account the impact on area and power.

To ensure that the four routers behave consistently under traffic patterns other than uniform random, we have repeated the above experiment for each of the other five traffic patterns in our set. The results for maximum throughput are summarized in Fig. 14. As shown, performance remains consistent with previous observations. The same is true for latency as a function of the injection rate. Consequently, we can safely average between traffic patterns for the rest of this section.

In order to perform a fair comparison of the efficiency of the four routers, we equalize throughput, area or energy by modifying datapath width. Therefore, the four networks will have different datapath widths, and consequently each packet will consist of a different number of flits.

For our first comparison, we equalize the maximum throughput. With equal maximum throughput and routers operating at their maximum frequencies, the single-stage router requires the same amount of energy to transfer a single bit from an input to an output as the wormhole router, and 17 percent less compared to the hybrid router. Due to the single-stage router’s low energy overhead, it has a wider datapath but the wormhole and hybrid routers have more buffering slots. Therefore, even though FIFOs are more costly, in our network they provide twice as much

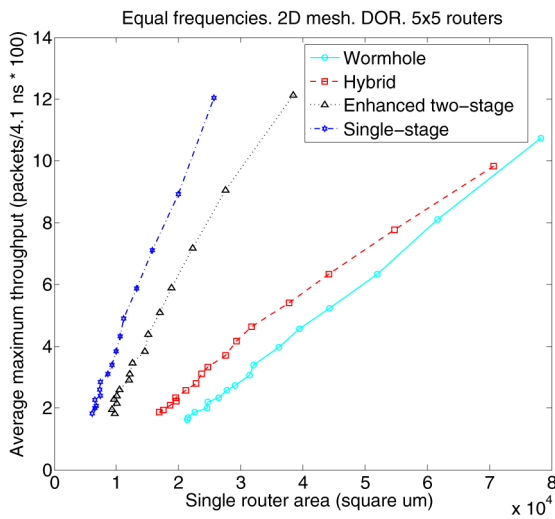


Fig. 15. Throughput versus area (equal frequencies).

buffering compared to the single-stage router. This increases the maximum throughput of the wormhole and hybrid routers.

The trends for router area are similar to those for energy. The Pareto-optimal curves relating maximum throughput and area are shown in Figs. 15 and 16. When operating all routers at their maximum frequencies, time is given in units of the largest cycle time among all data points (4.1ns). In this case, the single-stage router requires 66 percent less area for the same throughput compared to the wormhole router, and 65 percent compared to the hybrid router. If we equalize for router area or energy, the datapath of the single-stage router is wider than those of the wormhole and hybrid routers. As a result, the single-stage router offers three percent more throughput per unit energy compared to the wormhole router and 18 percent more compared to the hybrid router. The single-stage router also offers 67 percent more throughput per unit area compared to the wormhole router, and 62 percent more compared to the hybrid router. The trends hold when operating all routers at the same clock frequency. In this case, however, the single-stage router provides more throughput per unit area because the enhanced two-stage router no longer benefits from its higher maximum operating frequency.

In a Pareto-optimal point comparison with routers operating at their maximum frequencies, the single-stage router has a

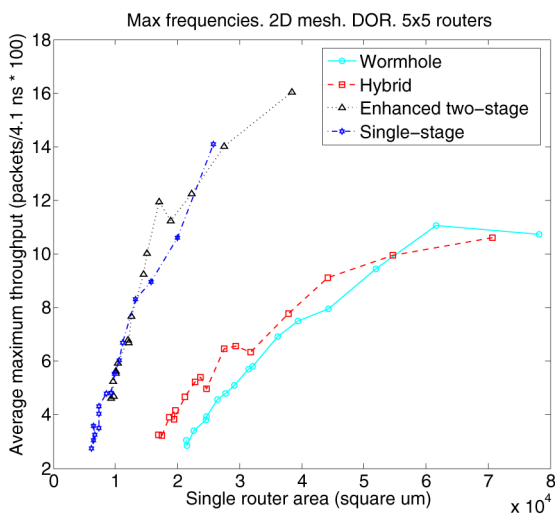


Fig. 16. Throughput versus area (max. frequencies).

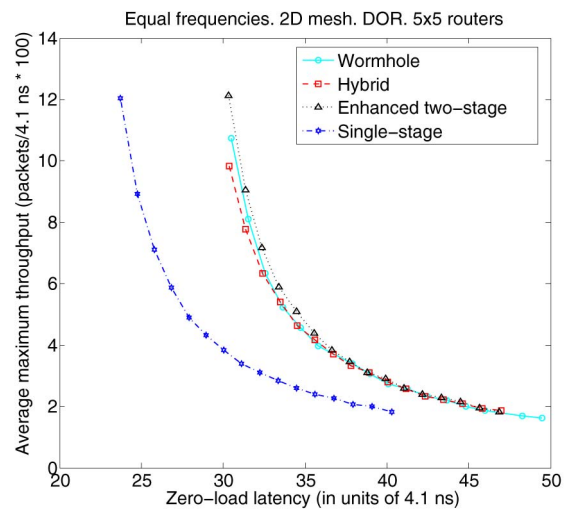


Fig. 17. Throughput versus latency (equal frequencies).

three percent increased zero-load latency for the same maximum throughput compared to the enhanced two-stage router. This result differs from [14] because our network for this comparison has two-cycle channels that are clocked at the same frequency as the routers, and thus favors the enhanced two-stage router. However, compared to the single-stage router, the wormhole and the hybrid routers still have a 10 percent higher zero-load latency for the same saturation throughput. This is because in our 2D mesh with a large measured average hop count of 6.2, the difference in cycle time between the wormhole and hybrid routers on the one hand and the single-stage router on the other hand is not large enough to outweigh the latter's smaller pipeline depth. Therefore, while the enhanced two-stage router has a lower zero-load latency than the three other routers, the single-stage router is still preferable to the FIFO-based routers. The Pareto-optimal curves are shown in Figs. 17 and 18.

With routers operating at the same clock frequency, the single-stage router leverages its reduced pipeline depth and offers lower zero-load latency for the same maximum throughput. The other three routers are comparable because they all have two pipeline stages; the slight differences in performance are due to different amounts of buffering provided by each and the resulting differences in maximum throughput.

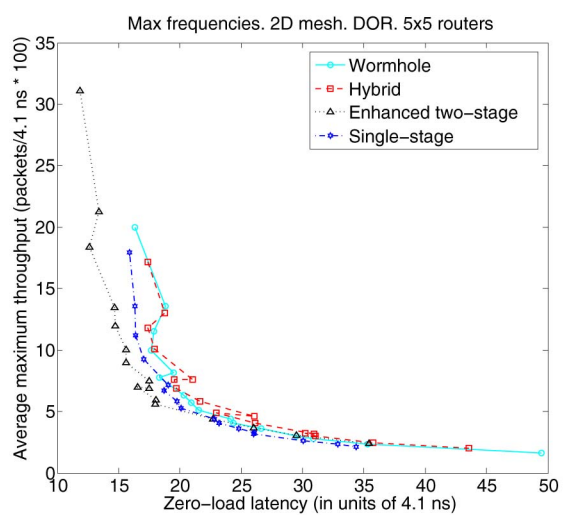


Fig. 18. Throughput versus latency (max. frequencies).

We also compared a hybrid EB-wormhole router based on the single-stage router. However, because of the significant timing overhead introduced by the FIFO and because of the lack of pipelining to split the critical path, this design had a 65 percent increased cycle time compared to the two-stage hybrid design we evaluate in this section. Moreover, the addition of a FIFO defeats the primary advantage of the single-stage router: its design simplicity. Consequently, we did not investigate this particular design point further.

Networks with different channel lengths will affect the number of buffer slots available to EB routers and thus the maximum throughput. However, this is a minor effect [12] as maximum throughput is mostly affected by contention in the routers. We have shown that even in a topology with short channels such as our 2D mesh, EB networks are more efficient, and that adding extra buffering in the form of input FIFOs decreases network efficiency. Topologies with longer channels provide more buffering opportunities for EB networks, while the associated impact on cost is comparable for EB and VC or wormhole networks. Also, longer channels require wormhole routers to use bigger input buffers in order to ensure that the longer credit roundtrip time can be covered. Hybrid routers, on the other hand, use the ready-valid handshake, and are thus not subject to this requirement.

Changing the network diameter or router radix will affect the fraction of the power and area that is consumed by the FIFOs, and will thus affect our comparison results. However, as long as FIFOs remain more expensive than EBs, EB routers will likely continue to be more cost efficient; otherwise, hybrid EB-wormhole designs will prove beneficial since they combine increased buffering with credit-less flow control and still take advantage of the FFs already present in the channel as distributed storage. The use of SRAM-based—rather than FF-based—FIFOs could also have impacted our results; however, for the FIFO sizes and technology library considered here, FF-based FIFOs are more area and energy efficient than compiler-generated SRAMs. Even for larger sizes, SRAM-based buffers may have larger power overheads than our results of 21 percent for the wormhole router and 17 percent for the hybrid router [19]. The exact SRAM implementation, e.g., compiler-generated versus custom-designed, may also significantly affect the SRAM cost. However, we did not investigate the implementation of efficient custom SRAMs in this study.

Overall, our results show that the single-stage router is preferable to the three other routers in terms of throughput efficiency, area, and energy. The enhanced two-stage router provides slightly better zero-load latency; however, the single-stage router remains preferable to the wormhole and hybrid routers in all aspects. This is true regardless of whether the routers operate at maximum or equal clock frequencies. Despite the single-stage router's larger cycle time, its area and energy savings can be traded for a wider datapath, enabling higher throughput and lower serialization latency. This highlights that design simplicity can result in decreased overhead and higher performance efficiency. Design simplicity also offers other advantages, such as reduced PnR flow turnaround time.

Thus, we have illustrated that EB flow control is beneficial compared to wormhole flow control, although the optimal type of EB router may differ depending on the network configuration. A previous study [12] came to similar conclusions when comparing EB and VC flow control.

4 CONCLUSIONS

In this work, we presented a comparison of EB and wormhole flow control techniques. Like VC routers, wormhole routers use FIFOs

at their inputs for buffering instead of using channels as distributed FIFOs like EB routers. However, they do not use VCs and therefore do not require VC allocators. Also, like EB routers, they use output arbiters instead of a full switch allocator. While they use credits, the associated control logic is much simpler due to the lack of VCs. A wormhole router's buffer has to at least be large enough to cover the buffer turnaround time; furthermore, as credits are consumed during the switch arbitration stage one cycle before the flit actually leaves the buffer, buffer occupancy is higher than for a router with the same amount of buffer space that uses the ready-valid handshake.

We also explore novel hybrid EB-wormhole router designs. Such designs replace the input EB with a FIFO. Hybrid routers still use the channels for buffering, but the additional FIFOs enable higher throughput at the cost of increased area and power. Hybrid routers use the ready-valid handshake instead of credits. Due to the complexity of the FIFO logic, a hybrid router based on the single-stage router has a 65 percent increased cycle time compared to a hybrid router based on the enhanced two-stage router. Therefore, we only consider the latter design in this study.

With the EB-wormhole comparison and the hybrid design exploration we closely examine the effect of design simplicity and input buffer cost. Moreover, we investigate the trade-off between input buffer cost and increased maximum throughput, as well as the trade-off of gaining buffering through network channels on the one hand and through additional input buffers on the other hand. As a result, we gain a better understanding of where the advantages of EB flow-control over VC flow-control reported in [12] originate.

As our results show for an 8×8 2D mesh with DOR, the enhanced two-stage router has a 25 percent smaller cycle time compared to the wormhole router, the fastest among the FIFO-based routers (i.e., wormhole or hybrid). However, the single-stage router is able to trade off area and power savings from its significantly simpler design for a wider datapath. That way, it offers three percent more throughput per unit energy, 62 percent more throughput per unit router area, and has a 10 percent lower zero-load latency for equal saturation throughput when compared to the most efficient FIFO-based router for each case. The single-stage router is preferable to the FIFO-based routers in all aspects.

The dominant factor in all results is the input FIFO. It is part of the router's critical path in all configurations. Moreover, it contributes 56 and 60 percent of the area and 21 and 17 percent of the power of the wormhole and hybrid routers, respectively, under a 20 percent flit injection rate. This highlights that using the channels for buffering instead of FIFOs provides significant gains, even in a topology with short channels such as our 2D mesh. Topologies with longer channels will provide more buffering in the EB channel, while affecting the cost of the EB and VC or wormhole networks similarly.

In summary, in our network configuration, EB flow control using the single-stage EB router is more efficient than either wormhole flow control or the hybrid designs. Since EB flow control has been shown to be more efficient than VC flow control, EBs represent a preferable choice over a wide range of buffered networks-on-chip, which are currently the dominant design choices.

ACKNOWLEDGMENTS

This work was supported in part by the US National Science Foundation (NSF) under Grant CCF-0702341, in part by the National Security Agency under Contract H98230-08-C-0272-P-3, in part by a Robert Bosch Stanford Graduate Fellowship, and in part by a Prof. Michael J. Flynn Stanford Graduate Fellowship.

REFERENCES

- [1] W.J. Dally and B. Towles, "Route Packets, Not Wires: On-Chip Interconnection Networks," *Proc. 38th Ann. Conf. Design Automation*, 2001.
- [2] Y. Hoskote, S. Vangal, A. Singh, N. Borkar, and S. Borkar, "A 5-GHz Mesh Interconnect for a Teraflops Processor," *IEEE Micro*, vol. 27, no. 5, pp. 51-61, Sept./Oct. 2007.
- [3] M.B. Taylor, W. Lee, J. Miller, D. Wentzlaff, I. Bratt, B. Greenwald, H. Hoffmann, P. Johnson, J. Kim, J. Psota, A. Saraf, N. Shnidman, V. Strumpfen, M. Frank, S. Amarasinghe, and A. Agarwal, "Evaluation of the Raw Microprocessor: An Exposed-Wire-Delay Architecture for Ilp and Streams," *ISCA '04: Proc. 31st Ann. Int'l Symp. Computer Architecture*, pp. 2-13, 2004.
- [4] W.J. Dally, "Virtual-Channel Flow Control," *IEEE Trans. Parallel and Distributed Systems*, vol. 3, no. 2, pp. 194-205, Mar. 1992.
- [5] P. Gratz, C. Kim, R. McDonald, S. Keckler, and D. Burger, "Implementation and Evaluation of On-Chip Network Architectures," *Proc. Int'l Conf. Computer Design (ICCD '06)*, pp. 477-484, 2006.
- [6] C. Gómez, M.E. Gómez, P. López, and J. Duato, "Reducing Packet Dropping in a Bufferless NoC," *Proc. 14th Int'l Euro-Par Conf. Parallel Processing*, 2008.
- [7] T. Moscibroda and O. Mutlu, "A Case for Bufferless Routing in On-Chip Networks," *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3, pp. 196-207, 2009.
- [8] J. Liu, L.-R. Zheng, and H. Tenhunen, "A Guaranteed-Throughput Switch for Network-on-Chip," *Proc. Int'l Symp. System-on-Chip*, pp. 31-34, 2003.
- [9] J. Duato, P. Lopez, F. Silla, and S. Yalamanchili, "A High Performance Router Architecture for Interconnection Networks," *Proc. Int'l Conf. Parallel Processing*, vol. 1, pp. 61-68, citeseer.ist.psu.edu/duato96high.html, 1996.
- [10] G. Michelogiannakis, D. Pnevmatikatos, and M. Katevenis, "Approaching Ideal NoC Latency with Pre-Configured Routes," *NOCS '07: Proc. First Int'l Symp. Networks-on-Chip*, pp. 153-162, 2007.
- [11] A. Kumar, L.-S. Peh, P. Kundu, and N.K. Jha, "Express Virtual Channels: Towards the Ideal Interconnection Fabric," *Proc. 34th Ann. Int'l Symp. Computer Architecture*, pp. 150-161, 2007.
- [12] G. Michelogiannakis, J. Balfour, and W.J. Dally, "Elastic Buffer Flow Control for On-Chip Networks," *HPCA '09: Proc. 15th Int'l Symp. High-Performance Computer Architecture*, pp. 151-162, 2009.
- [13] A. Pullini, F. Angiolini, D. Bertozzi, and L. Benini, "Fault Tolerance Overhead in Network-on-Chip Flow Control Schemes," *SBCCI '05: Proc. 18th Ann. Symp. Integrated Circuits and System Design*, pp. 224-229, 2005.
- [14] G. Michelogiannakis and W.J. Dally, "Router Designs for Elastic Buffer On-Chip Networks," *SC '09: Proc. Conf. High Performance Computing Networking, Storage and Analysis*, pp. 1-10, 2009.
- [15] M. Galles, "Spider: A High-Speed Network Interconnect," *IEEE Micro*, vol. 17, no. 1, pp. 34-39, Jan./Feb. 1997.
- [16] L. Ni and P. McKinley, "A Survey of Wormhole Routing Techniques in Direct Networks," pp. 492-506, 2000.
- [17] A. Banerjee, R. Mullins, and S. Moore, "A Power and Energy Exploration of Network-on-Chip Architectures," *NOCS '07: Proc. First Int'l Symp. Networks-on-Chip*, pp. 163-172, 2007.
- [18] W.J. Dally and B. Towles, *Principles and Practices of Interconnection Networks*. Morgan Kaufmann Publishers, Inc., 2003.
- [19] A.B. Kahng, B. Li, L.-S. Peh, and K. Samadi, "Orion 2.0: A Fast and Accurate NoC Power and Area Model for Early-Stage Design Space Exploration," *Proc. 46th Ann. Conf. Design Automation*, pp. 423-428, 2009.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.