MinBD: Minimally-Buffered Deflection Routing for Energy-Efficient Interconnect

Chris Fallin, Greg Nazario, Xiangyao Yu[†], Kevin Chang, Rachata Ausavarungnirun, Onur Mutlu

Carnegie Mellon University
{cfallin,gnazario,kevincha,rachata,onur}@cmu.edu

[†]Tsinghua University & Carnegie Mellon University
yxythu@gmail.com

A conventional Network-on-Chip (NoC) router uses input buffers to store in-flight packets. These buffers improve performance, but consume significant power. It is possible to bypass these buffers when they are empty, reducing dynamic power, but static buffer power, and dynamic power when buffers are utilized, remain. To improve energy efficiency, *bufferless deflection routing* removes input buffers, and instead uses deflection (misrouting) to resolve contention. However, at high network load, deflections cause unnecessary network hops, wasting power and reducing performance.

In this work, we propose a new NoC router design called the *minimally-buffered deflection* (MinBD) router. This router combines deflection routing with a small "side buffer," which is much smaller than conventional input buffers. A MinBD router places some network traffic that would have otherwise been deflected in this side buffer, reducing deflections significantly. The router buffers only a fraction of traffic, thus making more efficient use of buffer space than a router that holds every flit in its input buffers. We evaluate MinBD against input-buffered routers of various sizes that implement buffer bypassing, a bufferless router, and a hybrid design, and show that MinBD is more energy-efficient than all prior designs, and has performance that approaches the conventional input-buffered router with area and power close to the bufferless router.

I. INTRODUCTION

A network-on-chip is a first-order component of current and future multicore and manycore CMPs (Chip Multiprocessors) [10], and its design can be critical for system performance. As core counts continue to rise, NoCs with designs such as 2D-mesh (e.g., Tilera [40] and Intel Terascale [19]) are expected to become more common to provide adequate performance scaling. Unfortunately, packet-switched NoCs are projected to consume significant power. In the Intel Terascale 80-core chip, 28% of chip power is consumed by the NoC [19]; for MIT RAW, 36% [35]; for the Intel 48-core SCC, 10% [3]. NoC energy efficiency is thus an important design goal [4], [5].

Mechanisms have been proposed to make conventional input-buffered NoC routers more energy-efficient. For example, bypassing empty input buffers [39], [27] reduces some dynamic buffer power, but static power remains.[1] Such bypassing is also less effective when buffers are not frequently empty. *Bufferless deflection routers* [12], [28] remove router input buffers completely (hence eliminating their static and dynamic power) to reduce router power. When two flits[2] contend for a single router output, one must be deflected to another output. Thus, a flit never requires a buffer in a router. By controlling

which flits are deflected, a bufferless deflection router can ensure that all traffic is eventually delivered. Removing buffers yields simpler and more energy-efficient NoC designs: e.g., CHIPPER [12] reduces average network power by 54.9% in a 64-node system compared to a conventional buffered router.

Unfortunately, at high network load, deflection routing reduces performance and efficiency. This is because deflections occur more frequently when many flits contend in the network. Each deflection sends a flit further from its destination, causing unnecessary link and router traversals. Relative to a buffered network, a bufferless network with a high deflection rate wastes energy, and suffers worse congestion, because of these unproductive network hops. In contrast, a buffered router is able to hold flits (or packets) in its input buffers until the required output port is available, incurring no unnecessary hops. Thus, a buffered network can sustain higher performance at peak load.

Our goal is to obtain the energy efficiency of the bufferless approach with the high performance of the buffered approach. One prior work, AFC (Adaptive Flow Control), proposes a hybrid design that switches each router between a conventional input-buffered mode and a bufferless deflection mode [21]. However, switching to a conventional buffered design at high load incurs the energy penalty for buffering every flit: in other words, the efficiency gain over the baseline input-buffered router disappears once load rises past a threshold. AFC also requires the control logic for both forms of routing to be present at each network node, and requires power gating to turn off the input buffers and associated logic at low load. Another prior design, the Chaos router [23], also combines buffering with deflection, but still uses its input buffers for every flit that arrives. Ideally, a router would contain only a small amount of buffering, and would use this buffer space only for those flits that actually require it, rather than all flits that arrive.

We propose *minimally-buffered deflection* routing (MinBD) as a new NoC router design that combines both bufferless and buffered paradigms in a more fine-grained and efficient way. MinBD uses deflection routing, but also incorporates a small buffer. The router does not switch between modes, but instead, always operates in a minimally-buffered deflection mode, and can buffer or deflect any given flit. When a flit first arrives, it does not enter a buffer, but travels straight to the routing logic. If two flits contend for the same output, the routing logic chooses one to deflect, as in a bufferless router. However, the router can choose to buffer up to one deflected flit per cycle rather than deflecting it. This fine-grained buffering-deflection hybrid approach significantly reduces deflection rate (by 54% in our evaluations), and improves performance, as we show. It also incurs only a fraction of the energy cost of a conventional

---

[1]One recent estimate indicates that static power (of buffers and links) could constitute 80–90% of interconnect power in future systems [6].

[2]In a conventional bufferless deflection network, *flits* (several of which make up one packet) are independently routed, unlike most buffered networks, where a packet is the smallest independently-routed unit of traffic.
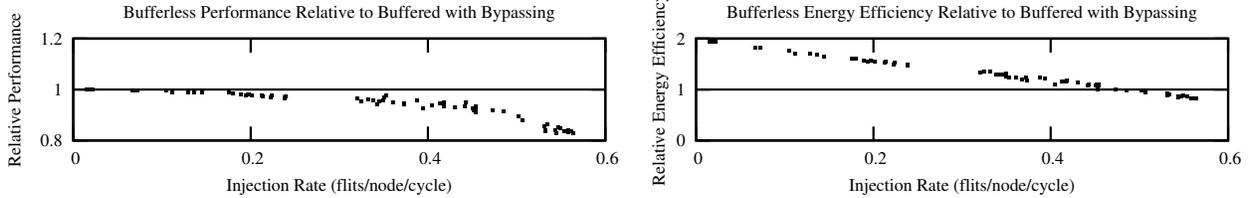
Figure 1. System performance and energy efficiency (performance per watt) of bufferless deflection routing, relative to conventional input-buffered routing (4 VCs, 4 flits/VC) that employs buffer bypassing, in a 4x4 2D mesh. Injection rate (X axis) for each workload is measured in the baseline buffered network.

buffered router, because only a relatively small fraction of flits are buffered (20% of all flits in our evaluations). As we show in our evaluations, MinBD provides higher energy efficiency while also providing high performance, compared to a comprehensive set of baseline router designs. **Our contributions** are:

- A new NoC router, *MinBD* (minimally-buffered deflection), that combines deflection routing with minimal buffering. The router performs deflection routing, but can choose to buffer up to one flit per cycle in a small side buffer, which significantly reduces deflection rate and enhances performance compared to a pure bufferless design while requiring smaller buffer space than a conventional input-buffered design.

- An evaluation of MinBD against aggressive NoC router baselines: a two-cycle virtual channel buffered router [9] with empty buffer bypassing [39], [27] at three buffer capacities (with a sensitivity analysis over many more configurations), a state-of-the-art bufferless deflection router, CHIPPER [12], and a hybrid bufferless-buffered design, AFC [21], with SPEC CPU2006 [34] multiprogrammed workloads on 16- and 64-node CMP systems. From our results, we conclude that MinBD has the best energy efficiency over all of these prior design points, while achieving competitive performance and critical path with the input-buffered router (the best-performing baseline) and competitive area and power with the pure-bufferless router (the smallest and most power-efficient baseline).

II. BACKGROUND

In this section, we give background on NoC-based cache-coherent CMPs, and on bufferless deflection routing, which we build upon. We assume the reader is familiar with the basic operation of conventional input-buffered routers; Dally and Towles [9] provide a good reference on these routers.

NoCs in cache-coherent CMPs: On-chip networks form the backbone of memory systems in many recently-proposed and prototyped large-scale CMPs (chip multiprocessors) [35], [19], [40]. Most such systems are cache-coherent shared memory multiprocessors. Packet-switched interconnect has served as the substrate for large cache-coherent systems for some time (e.g., for large multiprocessor systems such as SGI Origin [25]), and the principles are the same in a chip multiprocessor: each core, slice of a shared cache, or memory controller is part of one "node" in the network, and network nodes exchange packets that request and respond with data in order to fulfill memory accesses. For example, on a miss, a core's private cache might send a request packet to a shared L2 cache slice, and the shared cache might respond with a larger packet containing the requested cache block on an L2 hit, or might send another packet to a memory controller on an L2 miss. CMP NoCs are typically used to implement such a protocol between the cores, caches and memory controllers.

Bufferless Deflection Routers: Bufferless deflection routing was first proposed by Baran [2]. It has found renewed interest in NoC design because on-chip wires (hence, network links) are relatively cheap, in contrast to buffers, which consume significant die area and leakage power [28], [21], [6], [4], [5].

The fundamental unit of routing in a bufferless network is the *flit*, a packet fragment transferred by one link in one cycle. Flits are routed independently; thus, they must be reassembled after they are received. The basic operation of a bufferless deflection router is simple. In each cycle, flits arriving from neighbor routers enter the router pipeline. Because the router contains no buffers, flits are stored only in pipeline registers, and must leave the router at the end of the pipeline. Thus, the router must assign every input flit to some output port. When two flits request the same output port according to their ordinary routing function, the router deflects one of them to another port (this is always possible, as long as the router has as many outputs as inputs). This output will likely take the deflected flit further from its destination, and the flit will later have to work its way back. Thus, in such a design, care must be taken to avoid *livelock*, where a flit never arrives. Prior work uses priority schemes to ensure that every flit is eventually delivered. An early implementation of bufferless deflection routing for a NoC was BLESS [28], and CHIPPER [12] later provided a more efficient hardware implementation of the deflection routing and packet reassembly. CHIPPER is a baseline in our evaluations.

III. MOTIVATION

Previous NoC designs based on bufferless deflection routing [12], [28] were motivated largely by the observation that many NoCs in CMPs are overprovisioned for the common-case network load. In this case, a bufferless network can attain nearly the same application performance while consuming less power, which yields higher energy efficiency. We now examine the buffered-bufferless comparison in more detail. Fig. 1 shows (i) relative application performance (weighted speedup: see §V), and (ii) relative energy efficiency (performance per watt), when using a bufferless network, compared to a conventional buffered network. Both plots show these effects as a function of network load (average injection rate). Here we show a virtual channel buffered network (4 VCs, 4 flits/VC) (with buffer bypassing) and the CHIPPER bufferless deflection network [12] in a 4x4-mesh CMP (details on methodology in §V).

For low-to-medium network load, a bufferless network has performance close to a conventional buffered network, because the deflection rate is low: thus, most flits take productive network hops on every cycle, just as in the buffered network. In addition, the bufferless router has significantly reduced power (hence improved energy efficiency), because the buffers in a conventional router consume significant power. However, as network load increases, the deflection rate in a buffer-less deflection network also rises, because flits contend with
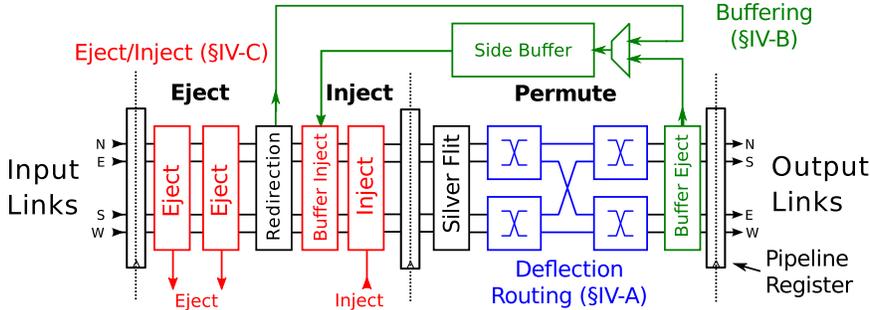
Figure 2. MinBD router pipeline.

each other more frequently. With a higher deflection rate, the dynamic power of a bufferless deflection network rises more quickly with load than dynamic power in an equivalent buffered network, because each deflection incurs some extra work. Hence, bufferless deflection networks lose their energy-efficiency advantage at high load. Just as important, the high deflection rate causes each flit to take a longer path to its destination, and this increased latency reduces the network throughput and system performance.

Overall, neither design obtains both good performance and good energy efficiency at all loads. If the system usually experiences low-to-medium network load, then the bufferless design provides adequate performance with low power (hence high energy efficiency). But, if we use a conventional buffered design to obtain high performance, then energy efficiency is poor in the low-load case, and even buffer bypassing does not remove this overhead because buffers consume static power regardless of use. Finally, simply switching between these two extremes at a per-router granularity, as previously proposed [21], does not address the fundamental inefficiencies in the bufferless routing mode, but rather, uses input buffers for *all* incoming flits at a router when load is too high for the bufferless mode (hence retains the energy-inefficiency of buffered operation at high load). We now introduce our *minimally-buffered deflection router* which combines bufferless and buffered routing in a new way to reduce this overhead.

IV. MINBD: MINIMALLY-BUFFERED DEFLECTION ROUTER

The *MinBD* (minimally-buffered deflection) router is a new router design that combines bufferless deflection routing with a small buffer, which we call the "side buffer." We start by outlining the **key principles** we follow to reduce deflection-caused inefficiency by using buffering:

1) When a flit would be deflected by a router, it is often better to buffer the flit and arbitrate again in a later cycle. Some buffering can avoid many deflections.
2) However, buffering every flit leads to unnecessary power overhead and buffer requirements, because many flits will be routed productively on the first try. The router should buffer a flit only if necessary.
3) Finally, when a flit arrives at its destination, it should be removed from the network (ejected) quickly, so that it does not continue to contend with other flits.

Basic High-Level Operation: The MinBD router does not use input buffers, unlike conventional buffered routers. Instead, a flit that arrives at the router proceeds directly to the routing and arbitration logic. This logic performs deflection routing, so that when two flits contend for an output port, one of

the flits is sent to another output instead. However, unlike a bufferless deflection router, the MinBD router can also *buffer* up to one flit per cycle in a single FIFO-queue side buffer. The router examines all flits at the output of the deflection routing logic, and if any are deflected, one of the deflected flits is removed from the router pipeline and buffered (as long as the buffer is not full). From the side buffer, flits are re-injected into the network by the router, in the same way that new traffic is injected. Thus, some flits that would have been deflected in a bufferless deflection router are removed from the network temporarily into this side buffer, and given a second chance to arbitrate for a productive router output when re-injected. This reduces the network's deflection rate (hence improves performance and energy efficiency) while buffering only a fraction of traffic.

We will describe the operation of the MinBD router in stages. First, §IV-A describes the deflection routing logic that computes an initial routing decision for the flits that arrive in every cycle. Then, §IV-B describes how the router chooses to buffer some (but not all) flits in the side buffer. §IV-C describes how buffered flits and newly-generated flits are injected into the network, and how a flit that arrives at its destination is ejected. Finally, §IV-D discusses correctness issues, and describes how MinBD ensures that all flits are eventually delivered.

A. Deflection Routing

The MinBD router pipeline is shown in Fig. 2. Flits travel through the pipeline from the inputs (on the left) to outputs (on the right). We first discuss the deflection routing logic, located in the Permute stage on the right. This logic implements deflection routing: it sends each input flit to its preferred output when possible, deflecting to another output otherwise.

MinBD uses the deflection logic organization first proposed in CHIPPER [12]. The *permutation network* in the Permute stage consists of two-input blocks arranged into two stages of two blocks each. This arrangement can send a flit on any input to any output. (Note that it cannot perform all possible permutations of inputs to outputs, but as we will see, it is sufficient for correct operation that at least one flit obtains its preferred output.) In each two-input block, arbitration logic determines which flit has a higher priority, and sends that flit in the direction of its preferred output. The other flit at the two-input block, if any, must take the block's other output. By combining two stages of this arbitration and routing, deflection arises as a distributed decision: a flit might be deflected in the first stage, or the second stage. Restricting the arbitration and routing to two-flit subproblems reduces complexity and allows for a shorter critical path, as demonstrated in [12].

In order to ensure correct operation, the router must arbitrate

between flits so that every flit is eventually delivered, despite deflections. We adapt a modified version of the Golden Packet priority scheme [12], which solves this *livelock-freedom problem*. This priority scheme is summarized in Ruleset 1. The basic idea of the Golden Packet priority scheme is that at any given time, at most one packet in the system is *golden*. The flits of this golden packet, called "golden flits," are prioritized above all other flits in the system (and contention between golden flits is resolved by the flit sequence number). While prioritized, golden flits are never deflected by non-golden flits. The packet is prioritized for a period long enough to guarantee its delivery. Finally, this "golden" status is assigned to one globally-unique packet ID (e.g., source node address concatenated with a request ID), and this assignment rotates through all possible packet IDs such that any packet that is "stuck" will eventually become golden. In this way, all packets will eventually be delivered, and the network is livelock-free. (See [12] for the precise way in which the Golden Packet is determined; we use the same rotation schedule.)

However, although Golden Packet arbitration provides correctness, a *performance* issue occurs with this priority scheme. Consider that most flits are not golden: the elevated priority status provides worst-case correctness, but does not impact common-case performance (prior work reported over 99% of flits are delivered without becoming golden [12]). However, when no flits are golden and ties are broken randomly, the arbitration decisions in the two permutation network stages are not coordinated. Hence, a flit might win arbitration in the first stage, and cause another flit to be deflected, but then lose arbitration in the second stage, and also be deflected. Thus, unnecessary deflections occur when the two permutation network stages are uncoordinated.

In order to resolve this performance issue, we observe that it is enough to ensure that in every router, at least one flit is prioritized above the others in every cycle. In this way, at least one flit will certainly not be deflected. To ensure this when no golden flits are present, we add a "silver" priority level, which wins arbitration over common-case flits but loses to the golden flits. One silver flit is designated randomly among the set of flits that enter a router at every cycle (this designation is local to the router, and not propagated to other routers). This modification helps to reduce deflection rate. Prioritizing a silver flit at every router does not impact correctness, because it does not deflect a golden flit if one is present, but it ensures that at least one flit will consistently win arbitration at both stages. Hence, deflection rate is reduced, improving performance.

B. Using a Small Buffer to Reduce Deflections

The key problem addressed by MinBD is *deflection inefficiency at high load*: in other words, when the network is highly utilized, contention between flits occurs often, and many flits will be deflected. We observe that adding a small buffer to a deflection router can reduce deflection rate, because the router can choose to buffer rather than deflect a flit when its output port is taken by another flit. Then, at a later time when output ports may be available, the buffered flit can re-try arbitration.

Thus, to reduce deflection rate, MinBD adds a "side buffer" that buffers only some flits that otherwise would be deflected. This buffer is shown in Fig. 2 above the permutation network. In order to make use of this buffer, a "buffer eject" block is placed in the pipeline after the permutation network. At this point, the arbitration and routing logic has determined which

flits to deflect. The buffer eject block recognizes flits that have been deflected, and picks up to one such deflected flit per cycle. It removes a deflected flit from the router pipeline, and places this flit in the side buffer, as long as the side buffer is not full. (If the side buffer is full, no flits are removed from the pipeline into the buffer until space is created.) This flit is chosen randomly among deflected flits (except that a golden flit is never chosen: see §IV-D). In this way, some deflections are avoided. The flits placed in the buffer will later be re-injected into the pipeline, and will re-try arbitration at that time. This re-injection occurs in the same way that new traffic is injected into the network, which we discuss below.

C. Injection and Ejection

So far, we have considered the flow of flits from router input ports (i.e., arriving from neighbor routers) to router output ports (i.e., to other neighbor routers). A flit must enter and leave the network at some point. To allow traffic to enter (inject) and leave (eject), the MinBD router contains inject and eject blocks in its first pipeline stage (see Fig. 2). When a set of flits arrive on router inputs, these flits first pass through the ejection logic. This logic examines the destination of each flit, and if a flit is addressed to the local router, it is removed from the router pipeline and sent to the local network node.[3] If more than one locally-addressed flit is present, the ejector picks one, according to the same priority scheme used by routing arbitration.

However, ejecting a single flit per cycle can produce a bottleneck and cause unnecessary deflections for flits that could not be ejected. In the workloads we evaluate, at least one flit is eligible to eject 42.8% of the time. Of those cycles, 20.4% of the time, at least two flits are eligible to eject. Hence, in ~8.5% of all cycles, a locally-addressed flit would be deflected rather than ejected if only one flit could be ejected per cycle. To avoid this significant deflection-rate penalty, we double the ejection bandwidth. To implement this, a MinBD router contains two ejector blocks. Each of these blocks is identical, and can eject up to one flit per cycle. Duplicating the ejection logic allows two flits to leave the network per cycle at every node.[4]

After locally-addressed flits are removed from the pipeline, new flits are allowed to enter. There are two injector blocks in the router pipeline shown in Fig. 2: (i) re-injection of flits from the side buffer, and (ii) injection of new flits from the local node. (The "Redirection" block prior to the injector blocks will be discussed in the next section.) Each block operates in the same way. A flit can be injected into the router pipeline whenever one of the four inputs does not have a flit present in a given cycle, i.e., whenever there is an "empty slot" in the network. Each injection block pulls up to one flit per cycle from an injection queue (the side buffer, or the local node's injection queue), and places a new flit in the pipeline when a slot is available. Flits from the side buffer are re-injected before new traffic is injected into the network. However, note that there is no guarantee that a free slot will be available for an injection in any given cycle. We now address this starvation problem for side buffer re-injection.

---

[3]Note that flits are reassembled into packets after ejection. To implement this reassembly, we use the Retransmit-Once scheme, as used by CHIPPER, which uses MSHRs (Miss-Status Handling Registers [24], or existing buffers in the cache system) to reassemble packets in place. See [12] for details.

[4]For fairness, because dual ejection widens the datapath from the router to the local node (core or cache), we also implement dual ejection in the baseline bufferless deflection network and input-buffered network to which we compare.

CPU cores	Out-of-order, 3-wide issue and retire (1 memory op/cycle), 16 MSHRs [24]
L1 caches	64 KB, 4-way associative, 32-byte blocks
L2 (shared) cache	Distributed across nodes; perfect (always hits) to penalize our design conservatively & isolate network performance from memory effects
Shared cache mapping	Consecutive cache blocks striped across L2 cache slices
Cache coherence scheme	Directory-based, perfect directory (SGI Origin protocol [25])
Data packet sizes	1-flit request packets, 4-flit data packets
Network Links	1-cycle latency (separate pipeline stage), 2.5mm, 128 bits wide
Baseline bufferless router	CHIPPER [12], 2-cycle router latency; 64-cycle Golden Epoch; Retransmit-Once [12]
Baseline buffered router	(m VCs, n flits/VC): $(8, 8)$, $(4, 4)$, $(4, 1)$. 2-cycle latency, buffer bypassing [39], [27]. Additional configurations evaluated in Fig. 4.
AFC (Adaptive Flow Control)	As described in [21]: 4 VCs/channel, 4 flits/VC. 2-cycle latency (buffered & bufferless modes). Implements buffer bypassing as well.
MinBD	2-cycle latency (§IV); 4-flit side buffer (single FIFO); $C_{threshold} = 2$; 64-cycle Golden Epoch; Retransmit-Once [12]

Table I. Simulated baseline system parameters.

D. Ensuring Side Buffered Flits Make Progress

When a flit enters the side buffer, it leaves the router pipeline, and must later be re-injected. As we described above, flit re-injection must wait for an empty slot on an input link. It is possible that such a slot will not appear for a long time. In this case, the flits in the side buffer are delayed unfairly while other flits make forward progress.

To avoid this situation, we implement *buffer redirection*. The key idea of buffer redirection is that when this side buffer starvation problem is detected, one flit from a randomly-chosen router input is *forced* to enter the side buffer. Simultaneously, the flit at the head of the side buffer is allowed to inject into the slot created by the forced flit buffering. In other words, one router input is "redirected" into the FIFO buffer for one cycle, in order to allow the buffer to make forward progress. This redirection is enabled for one cycle whenever the side buffer injection is starved (i.e., has a flit to inject, but no free slot allows the injection) for more than some threshold $C_{threshold}$ cycles (in our evaluations, $C_{threshold} = 2$). Finally, note that if a golden flit is present, it is never redirected to the buffer, because this would break the delivery guarantee.

E. Livelock and Deadlock-free Operation

MinBD provides livelock-free delivery of flits using Golden Packet and buffer redirection. If no flit is ever buffered, then Golden Packet [12] ensures livelock freedom (the "silver flit" priority never deflects any golden flit, hence does not break the guarantee). Now, we argue that adding side buffers does not cause livelock. First, the buffering logic never places a golden flit in the side buffer. However, a flit could enter a buffer and then become golden while waiting. Redirection ensures correctness in this case: it provides an upper bound on residence time in a buffer (because the flit at the head of the buffer will leave after a certain threshold time in the worst case). If a flit in a buffer becomes golden, it only needs to remain golden long enough to leave the buffer in the worst case, then progress to its destination. We choose the threshold parameter ($C_{threshold}$) and golden epoch length so that this is always possible. More details can be found in our extended technical report [13].

MinBD achieves deadlock-free operation by using Retransmit-Once [12], which ensures that every node always consumes flits delivered to it by dropping flits when no reassembly/request buffer is available. This avoids packet-reassembly deadlock (as described in [12]), as well as protocol level deadlock, because message-class dependencies [16] no longer exist.

V. EVALUATION METHODOLOGY

To obtain application-level performance as well as network performance results, we use an in-house CMP simulator. This simulator consumes instruction traces of x86 applications, and faithfully models the CPU cores and the cache hierarchy, with a directory-based cache coherence protocol (based on the SGI Origin protocol [8]) running on the modeled NoC.

The CPU cores model stalls, and interact with the caches and network in a closed-loop way. The modeled network routers are cycle-accurate, and are assumed to run in a common clock domain. The instruction traces are recorded with a Pin-tool [26], sampling representative portions of each benchmark as determined by PinPoints [30]. We find that simulating 25M cycles gives stable results with these traces. Detailed system parameters are shown in Table I.

Note that we make use of a *perfect shared cache* to stress the network, as was done in the CHIPPER [12] and BLESS [28] bufferless router evaluations. In this model, every request generated by an L1 cache miss goes to a shared cache slice, and the request is always assumed to hit and return data. This potentially increases network load relative to a real system, where off-chip memory bandwidth can also be a bottleneck. However, note that this methodology is *conservative*: because MinBD degrades performance relative to the buffered baseline, the performance degradation that we report is an upper bound on what would occur when other bottlenecks are considered. We choose to perform our evaluations this way in order to study the true capacity of the evaluated networks (if network load is always low because system bottlenecks such as memory latency are modeled, then the results do not give many insights about router design). Note that the cache hierarchy details (L1 and L2 access latencies, and MSHRs) are still realistic. We remove only the off-chip memory latency/bandwidth bottleneck.

Baseline Routers: We compare MinBD to a conventional input-buffered virtual channel router [9] with buffer bypassing [39], [27], a bufferless deflection router (CHIPPER [12]), and a hybrid bufferless-buffered router (AFC [21]). In particular, we sweep buffer size for input-buffered routers. We describe a router with m virtual channels (VCs) per input and n flits of buffer capacity per VC as an (m, n)-buffered router. We compare to a $(8, 8)$, $(4, 4)$, and $(4, 1)$-buffered routers in our main results. The $(8, 8)$ point represents a very large (overprovisioned) baseline, while $(4, 4)$ is a more reasonable general-purpose configuration. The $(4, 1)$ point represents the *minimum* buffer size for deadlock-free operation (two message classes [16], times two to avoid routing deadlock [20]). Furthermore, though 1-flit VC buffers would reduce throughput because they do not cover the credit round-trip latency, we optimistically assume zero-latency credit traversal in our simulations (which benefits the baseline design, hence is conservative for our claims). Finally, we simulate smaller (non-deadlock-free) designs with 1 and 2 VCs per input for our power-performance tradeoff evaluation in §VI-B (Fig. 4), solely for completeness (we were able to avoid deadlock at moderate loads and with finite simulation length for these runs).

Application Workloads: We use SPEC CPU2006 [34] benchmarks (26 applications[5]) to construct 75 multiprogrammed

---

[5] 410.bwaves, 416.gamess and 434.zeusmp were excluded because we were not able to collect representative traces from these applications.
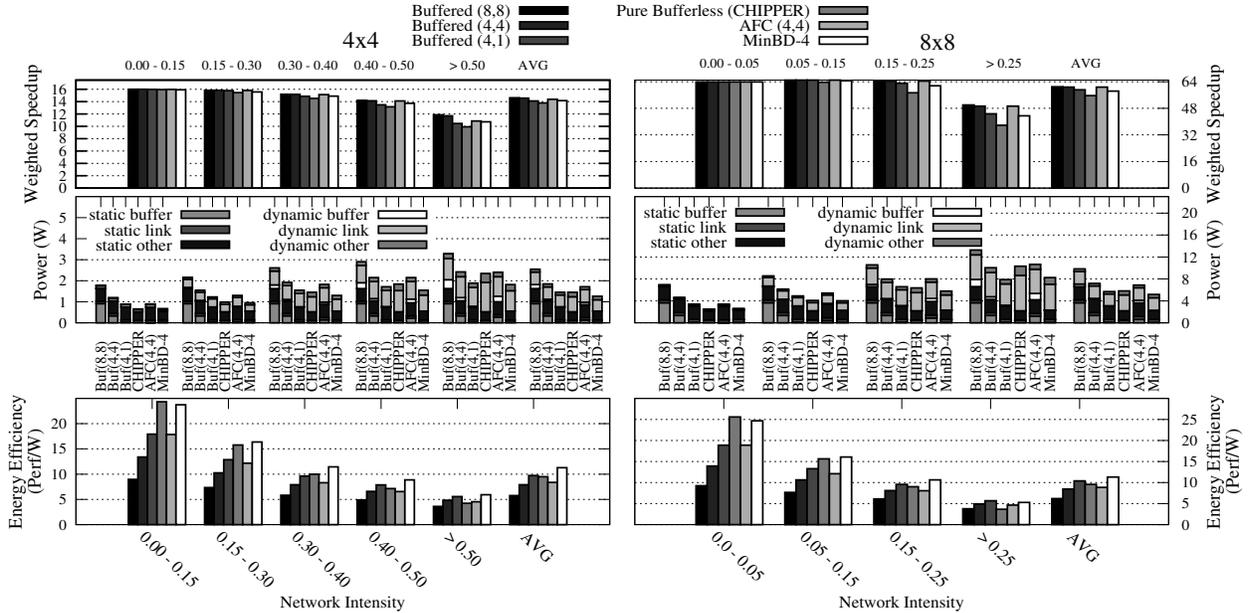
5

Figure 3. Performance (weighted speedup), network power, and energy efficiency (performance per watt) in 16 (4x4) and 64 (8x8)-node CMP systems.

workloads (which consist of many single-threaded benchmark instances that run independently). Each workload consists of 16 or 64 randomly-selected applications which are randomly mapped onto the mesh. Workloads are categorized by average network injection rate in the baseline $(4, 4)$-buffered system (measured in flits/cycle/node). For 4x4 workloads, these injection rate categories are $(0, 0.15)$, $(0.15, 0.3)$, $(0.3, 0.4)$, $(0.4, 0.5)$, and > 0.5; for 8x8 workloads, $(0, 0.05)$, $(0.05, 0.15)$, $(0.15, 0.25)$, and > 0.25 (an 8x8 mesh saturates at a lower load than a 4x4 mesh, due to limited bisection bandwidth). Each category contains 15 workloads.

Synthetic-Traffic Workloads: To show robustness under various traffic patterns, we evaluate 4x4 and 8x8 networks with uniform-random, bit-complement, and transpose synthetic traffic [9] in addition to application workloads. For each pattern, network injection rate is swept from zero to network saturation.

Performance Metrics: To measure system performance, we use the well-known *Weighted Speedup* metric [33]: $WS = \sum_{i=1}^{N} \frac{IPC_i^{shared}}{IPC_i^{alone}}$. All IPC_i^{alone} values are measured on the baseline bufferless network. Weighted speedup correlates to system throughput [11] and is thus a good general metric for multiprogrammed workloads.

Power Modeling: We use a modified and improved version of ORION 2.0 [38] (configured for 65nm), as developed by Grot et al. [15], as well as Verilog models synthesized with a commercial 65nm design library. We use Verilog models of router control logic, and add datapath area and power using ORION models for crossbars, buffers, links, and pipeline registers. CHIPPER and MinBD do not use a conventional crossbar, instead decoupling the crossbar into a permutation network, which we model using muxes. We rely on ORION's modeling for the baseline input-buffered router's control logic (e.g., arbitration). This hybrid synthesis / ORION approach models each portion of the router in a way that captures its key limitations. The control logic is logic- rather than wiring-dominated, and its arbitration paths determine the critical path; hence, standard-cell synthesis will model area, power and timing of MinBD and CHIPPER control logic with reasonable accuracy. The router datapath is wiring-dominated and relies

on heavily-optimized components with custom layouts such as large crossbars and wide muxes, pipeline registers and memories. ORION explicitly models such router components.

We report energy efficiency as performance-per-watt, computed as weighted speedup divided by average network power.

VI. EVALUATION

In this section, we evaluate MinBD against a bufferless deflection router [12] and an input-buffered router with buffer bypassing [39], [27], as well as a hybrid of these two, AFC [21], and demonstrate that by using a combination of deflection routing and buffering, MinBD achieves performance competitive with the conventional input-buffered router (and higher than the bufferless deflection router), with a smaller buffering requirement, and better energy efficiency than all prior designs.

A. Application Performance

Fig. 3 (top pane) shows application performance as weighted speedup for 4x4 (16-node) and 8x8 (64-node) CMP systems. The plots show average results for each workload category, as described in §V, as well as overall average results. Each bar group shows the performance of three input-buffered routers: 8 VCs with 8 flits/VC, 4 VCs with 4 flits/VC, and 4 VCs with 1 flit/VC. Next is CHIPPER, the bufferless deflection router, followed by AFC [21], a coarse-grained hybrid router that switches between a bufferless and a buffered mode. MinBD is shown last in each bar group. We make several observations:

1. MinBD improves performance relative to the bufferless deflection router by 2.7% (4.9%) in the 4x4 (8x8) network over all workloads, and 8.1% (15.2%) in the highest-intensity category. Its performance is within 2.7% (3.9%) of the $(4, 4)$ input-buffered router, which is a reasonably-provisioned baseline, and within 3.1% (4.2%) of the $(8, 8)$ input-buffered router, which has large, power-hungry buffers. Hence, adding a side buffer allows a deflection router to obtain significant performance improvement, and the router becomes more competitive with a conventional buffered design.

2. Relative to the 4-VC, 1-flit/VC input-buffered router (third bar), which is the smallest deadlock-free (i.e., correct) design,

MinBD performs nearly the same despite having less buffer space (4 flits in MinBD vs. 16 flits in $(4, 1)$-buffered). Hence, buffering only a portion of traffic (i.e., flits that would have been deflected) makes more efficient use of buffer space.

3. AFC, the hybrid bufferless/buffered router which switches modes at the router granularity, performs essentially the same as the 4-VC, 4-flit/VC input-buffered router, because it is able to use its input buffers when load increases. However, as we will see, this performance comes at an efficiency cost relative to our hybrid design.

B. Network Power and Energy Efficiency

Network Power: Fig. 3 (middle pane) shows average total network power, split by component and type (static/dynamic), for 4x4 and 8x8 networks across the same workloads. Note that static power is shown in the bottom portion of each bar, and dynamic power in the top portion. Each is split into buffer power, link power, and other power (which is dominated by datapath components, e.g., the crossbar and pipeline registers). We make several observations:

1. Buffer power is a large part of total network power in the input-buffered routers that have reasonable buffer sizes, i.e., $(4, 4)$ and $(8, 8)$ (VCs, flits/VC), even with empty-buffer bypassing, largely because static buffer power (bottom bar segment) is significant. Removing large input buffers reduces static power in MinBD as well as the purely-bufferless baseline, CHIPPER.[6] Because of this reduction, MinBD's total network power never exceeds that of the input-buffered baselines, except in the highest-load category in an 8x8-mesh (by 4.7%).

2. Dynamic power is larger in the baseline deflection-based router, CHIPPER, than in input-buffered routers: CHIPPER has 31.8% (41.1%) higher dynamic power than the $(4, 4)$-buffered router in the 4x4 (8x8) networks in the highest-load category. This is because bufferless deflection-based routing requires more network hops, especially at high load. However, in a 4x4 network, MinBD consumes less dynamic power (by 8.0%) than the $(4, 4)$-buffered baseline in the highest-load category because reduced deflection rate (by 58%) makes this problem relatively less significant, and allows savings in buffer dynamic energy and a simplified datapath to come out. In an 8x8 network, MinBD's dynamic power is only 3.2% higher.

3. MinBD and CHIPPER, which use a permutation network-based datapath rather than a full 5x5 crossbar, reduce datapath static power (which dominates the "static other" category) by 31.0%: the decoupled permutation network arrangement has less area, in exchange for partial permutability (which causes some deflections). Input-buffered routers and AFC require a full crossbar because they cannot deflect flits when performing buffered routing (partial permutability in a non-deflecting router would significantly complicate switch arbitration, because each output arbiter's choice would be limited by which other flits are traversing the router).

4. AFC, the coarse-grained hybrid, has nearly the same network power as the $(4, 4)$ buffered router at high load: 0.6% (5.7%) less in 4x4 (8x8). This is because its buffers are enabled most of the time. At low load, when it can power-gate its buffers frequently, its network power reduces. However, AFC's

network power is still higher than the pure bufferless router (CHIPPER) or MinBD because (i) it still spends some time in its buffered mode, and (ii) its datapath power is higher, as described above. On average, AFC still consumes 36.8% (18.1%) more network power than CHIPPER, and 33.5% (33.0%) more than MinBD, in the lowest-load category.

Energy efficiency: Fig. 3 (bottom pane) shows energy efficiency. We make two key observations:

1. MinBD has the highest energy efficiency of any evaluated design: on average in 4x4 (8x8) networks, 42.6% (33.8%) better than the reasonably-provisioned $(4, 4)$ input-buffered design. MinBD has 15.9% (8.7%) better energy-efficiency than the most energy-efficient prior design, the $(4, 1)$-buffered router.

2. At the highest network load, MinBD becomes less energy-efficient compared to at lower load, and its efficiency degrades at a higher rate than the input-buffered routers with large buffers (because of deflections). However, its per-category energy-efficiency is still better than all baseline designs, with two exceptions. In the highest-load category (near saturation) in an 8x8-mesh, MinBD has nearly the same efficiency as the $(4, 1)$-buffered router (but, note that MinBD is much more efficient than this baseline router at lower loads). In the lowest-load category in a 4x4 mesh, the purely-bufferless router CHIPPER is slightly more energy-efficient (but, note that CHIPPER's performance and efficiency degrade quickly at high loads).
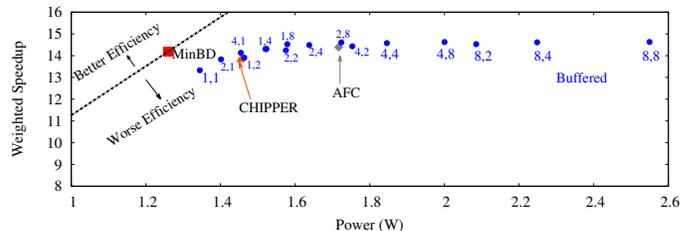

Figure 4. Power (X) vs. application performance (Y) in 4x4 networks. The line represents all points with equivalent performance-per-watt to MinBD.

We conclude that, by achieving competitive performance with the buffered baseline, and making more efficient use of a much smaller buffer capacity (hence reducing buffer power and total network power), MinBD provides better energy efficiency than prior designs. To summarize this result, we show a 2D plot of power and application performance in Fig. 4 for 4x4 networks, and a wider range of buffered router designs, as well as MinBD and CHIPPER. (Recall from §V that several of the baseline input-buffered designs are not deadlock free (too few VCs) or have a buffer depth that does not cover credit round-trip latency, but we evaluate them anyway for completeness.) In this plot, with power on the X axis and application performance on the Y axis, a line through the origin represents a fixed performance-per-watt ratio (the slope of the line). This equal-efficiency line is shown for MinBD. Points above the line have better efficiency than MinBD, and points below have worse efficiency. As shown, MinBD presents the best energy efficiency among all evaluated routers. The trend in an 8x8 network (not shown for space) is similar (see technical report [13]).

C. Performance Breakdown

To understand the observed performance gain in more detail, we now break down performance by each component of MinBD. Fig. 5 shows performance (for 4x4 networks) averaged across all workloads for eight *deflection systems*, which constitute all possible combinations of MinBD's mechanisms

---

[6]Note that network power in the buffered designs takes buffer bypassing into account, which reduces these baselines' dynamic buffer power. The $(4, 4)$-buffered router bypasses 73.7% (83.4%) of flits in 4x4 (8x8) networks. Without buffer bypassing, this router has 7.1% (6.8%) higher network power, and 6.6% (6.4%) worse energy-efficiency.
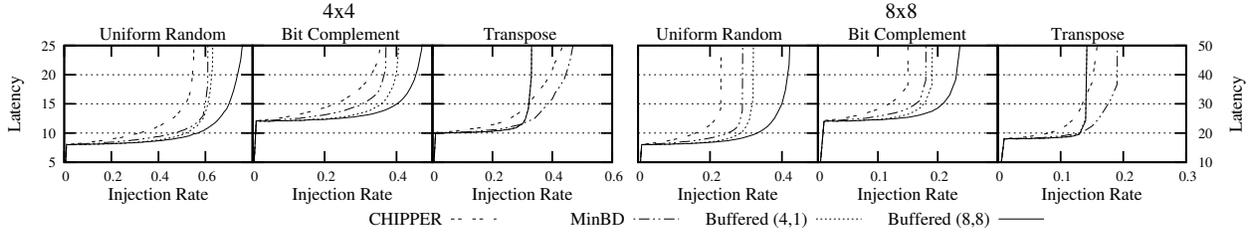
Figure 6. Synthetic traffic evaluations for MinBD, CHIPPER and input-buffered routers (with small and large input buffers), in 4x4 and 8x8 meshes. (For the transpose pattern, the curves for both input-buffered routers overlap.)

added to the baseline (CHIPPER) router: dual-width ejection (**D**), silver-flit prioritization (**S**), and the side buffer (**B**), shown with the same three input-buffered configurations as before. The eighth bar (**D+S+B**), which represents all three mechanisms added to the baseline deflection network, represents MinBD. Table II shows deflection rate for the same set of systems.
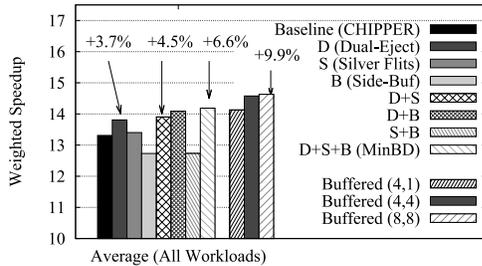

Figure 5. Breakdown of performance gains for each mechanism in MinBD.

Table II. Average deflection rates for deflection-based routers.

Baseline (CHIPPER)	D	S	B	D+S	D+B	S+B	D+S+B (MinBD)
0.28	0.22	0.27	0.17	0.22	0.11	0.16	0.10

We draw three main conclusions:

1. All mechanisms individually contribute to performance and reduce deflection rate. Dual ejection (**D**) increases performance by 3.7% over baseline CHIPPER.[7] Adding silver-flit prioritization (**D+S**) increases performance by 0.7% over dual-ejection alone. Adding a side buffer to the other two mechanisms (**D+S+B**) increases performance by 2.0% on average.

2. Adding a side buffer by itself (**B**) is not sufficient to attain the performance that MinBD achieves. In fact, when only the side buffer is added, performance drops by 4.3% relative to baseline CHIPPER. The degradation occurs primarily because the increased in-network throughput (enabled by a reduced deflection rate) in the network with side buffers places more ejection pressure on nodes, which exacerbates the ejection bottleneck that we observed in §IV-A. This performance reduction comes despite a reduced deflection rate: even though flits are not deflected as frequently, they must still be removed from the network efficiently for performance to increase.

3. Adding dual ejection to the side buffered system (**D+B**) to address the ejection bottleneck increases performance to 5.8% above baseline. Silver-flit prioritization in addition to this configuration point yields the MinBD router (eighth bar), which attains 6.6% performance above baseline (2.7% above dual-ejection alone) on average for all workload intensities. Overall, deflection rate reduces by 64% from baseline CHIPPER to MinBD (and 54% from CHIPPER with dual-ejection (**D**) to MinBD, as shown in our primary results).

[7]The main results presented in Fig. 3 use this data point (with dual ejection) in order to make a fair (same external router interface) comparison.

D. Synthetic Traffic Performance

We study the network-level performance of MinBD and baseline designs by applying synthetic traffic patterns: uniform random, bit-complement, and transpose [9]. Fig. 6 shows latency curves with injection rate swept from zero to saturation for the bufferless deflection router, MinBD, and the $(4,1)$ and $(8,8)$ input-buffered router (other input-buffered routers are omitted for clarity; these are the smallest and largest evaluated input-buffered routers, to show lower and upper bounds). Under uniform random traffic (which most closely resembles our multiprogrammed application workloads with striped data mapping), MinBD performs better than the bufferless baseline, with a higher saturation point. MinBD has a lower network saturation point than the input-buffered network with large input buffers, but very similar saturation point to the small $(4,1)$ input-buffered router, as our earlier results indicated. We conclude that with only 4 flits of buffering per router, MinBD closes nearly half of the gap in network saturation throughput between the bufferless router (CHIPPER) and the largest input-buffered router (with 256 flits of total buffer space), and performs similarly to a smaller input-buffered router with 16 flits of total buffer space.

In addition, non-uniform traffic patterns demonstrate the robustness and adaptivity of deflection routing: in particular, the transpose traffic pattern demonstrates a lower saturation point in the input-buffered router than either deflection-based router (MinBD or CHIPPER). This advantage occurs because deflection routing is *adaptive* (a flit's path can change based on network conditions). Deflections spread traffic away from hotspots and balance load in unbalanced traffic patterns. While adaptive routing is also possible in input-buffered routers, it is more complex because it requires the router to track network load or congestion (locally or globally) and make decisions accordingly. In contrast, deflection routing provides adaptivity by virtue of its ordinary operation.

E. Sensitivity to Side Buffer Size

As side buffer size is varied from 1 to 64 flits, mean weighted speedup (application performance) changes only 0.2% on average across all workloads (0.9% in the highest-intensity category) in 4x4 networks. We conclude that the *presence* of the buffer (to buffer at least one deflected flit) is more important than its size, because the average utilization of the buffer is low. In a 4x4 MinBD network with 64-flit side buffers at saturation (61% injection rate, uniform random), the side buffer is empty 48% of the time on average; 73% of the time, it contains 4 or fewer flits; 93% of the time, 16 or fewer. These measurements suggest that a very small side buffer captures most of the benefit. Furthermore, total network power increases by 19% (average across all 4x4 workloads) when a 1-flit buffer per router is increased to a 64-flit buffer per router. Hence, a larger buffer wastes power without significant performance benefit.

8

Router Design	CHIPPER	MinBD	Buffered $(8, 8)$	Buffered $(4, 4)$	Buffered $(4, 1)$
Normalized Area	1.00	1.03	2.06	1.69	1.60
Normalized Critical Path Length	1.00	1.07	0.99	0.99	0.99

Table III. Normalized router area and critical path for bufferless and buffered baselines, compared to MinBD.

We avoid a 1-flit side buffer because of the way the router is pipelined: such a single-flit buffer would either require for a flit to be able to enter, then leave, the buffer in the same cycle (thus eliminating the independence of the two router pipeline stages), or else could sustain a throughput of one flit only every two cycles. (For this sensitivity study, we optimistically assumed the former option for the 1-flit case.) The 4-flit buffer we use avoids this pipelining issue, while increasing network power by only 4% on average over the 1-flit buffer.

F. Hardware Cost: Router Area and Critical Path

We present normalized router area and critical path length in Table III. Both metrics are normalized to the bufferless deflection router, CHIPPER, because it has the smallest area of all routers. MinBD adds only 3% area overhead with its small buffer. In both CHIPPER and MinBD, the datapath dominates the area. In contrast, the large-buffered baseline has 2.06x area, and the reasonably-provisioned buffered baseline has 1.69x area. Even the smallest deadlock-free input-buffered baseline has 60% greater area than the bufferless design (55% greater than MinBD). In addition to reduced buffering, the reduction seen in CHIPPER and MinBD is partly due to the simplified datapath in place of the 5x5 crossbar (as also discussed in §VI-B). Overall, MinBD reduces area relative to a conventional input-buffered router both by significantly reducing the required buffer size, and by using a more area-efficient datapath.

Table III also shows the normalized critical path length of each router design, which could potentially determine the network operating frequency. MinBD increases critical path by 7% over the bufferless deflection router, which in turn has a critical path 1% longer than an input-buffered router. In all cases, the critical path is through the flit arbitration logic (the permutation network in MinBD and CHIPPER, or the arbiters in the input-buffered router). MinBD increases critical path relative to CHIPPER by adding logic in the deflection-routing stage to pick a flit to buffer, if any. The buffer re-injection and redirection logic in the first pipeline stage (eject/inject) does not impact the critical path because the permutation network pipeline stage has a longer critical path.

VII. RELATED WORK

To our knowledge, MinBD is the first NoC router design that combines deflection routing with a small side buffer that reduces deflection rate. Other routers combine deflection routing with buffers, but do not achieve the efficiency of MinBD because they either continue to use input buffers for all flits (Chaos router) or switch all buffers on and off at a coarse granularity with a per-router mode switch (AFC), in contrast to MinBD's fine-grained decision to buffer or deflect each flit.

Buffered NoCs that also use deflection: Several routers that primarily operate using buffers and flow control also use deflection routing as a secondary mechanism under high load. The Chaos Router [23] deflects packets when a packet queue becomes full to probabilistically avoid livelock. However, all packets that pass through the router are buffered; in contrast, MinBD performs deflection routing first, and only buffers some flits that would have been deflected. This key aspect of our design reduces buffering requirements and buffer power. The Rotary Router [1] allows flits to leave the router's inner ring on a non-productive output port after circulating the ring enough times, in order to ensure forward progress. In this case, again, deflection is used as an escape mechanism to ensure probabilistic correctness, rather than as the primary routing algorithm, and all packets must pass through the router's buffers.

Other bufferless designs: Several prior works propose bufferless router designs [12], [28], [37], [14], [17]. We have already extensively compared to CHIPPER [12], from which we borrow the deflection routing logic design. BLESS [28], another bufferless deflection network, uses a more complex deflection routing algorithm. Later works showed BLESS to be difficult to implement in hardware [12], [17], [27], thus we do not compare to it in this work. Other bufferless networks drop rather than deflect flits upon contention [14], [17]. Some earlier large multiprocessor interconnects, such as those in HEP [31] and Connection Machine [18], also used deflection routing. The HEP router combined some buffer space with deflection routing [32]. However, these routers' details are not well-known, and their operating conditions (large off-chip networks) are significantly different than those of modern NoCs.

Improving high-load performance in bufferless networks: Some work has proposed *congestion control* to improve performance at high network load in bufferless deflection networks [29], [7]. Both works used source throttling: when network-intensive applications cause high network load which degrades performance for other applications, these intensive applications are prevented from injecting some of the time. By reducing network load, source throttling reduces deflection rate and improves overall performance and fairness. These congestion control techniques and others (e.g., [36]) are orthogonal to MinBD, and could improve MinBD's performance further.

Hybrid buffered-bufferless NoCs: AFC [21] combines a bufferless deflection router based on BLESS [28] with input buffers, and switches between bufferless deflection routing and conventional input-buffered routing based on network load at each router. While AFC has the performance of buffered routing in the highest-load case, with better energy efficiency in the low-load case (by power-gating buffers when not needed), it misses opportunity to improve efficiency because it switches buffers on at a coarse granularity. When an AFC router experiences high load, it performs a mode switch which takes several cycles in order to turn on its buffers. Then, it pays the buffering energy penalty for every flit, whether or not it would have been deflected. It also requires buffers as large as the baseline input-buffered router design in order to achieve equivalent high-load performance. As a result, its network power is nearly as high as a conventional input-buffered router at high load, and it requires fine-grained power gating to achieve lower power at reduced network load. In addition, an AFC router has a larger area than a conventional buffered router, because it must include both buffers and buffered-routing control logic as well as deflection-routing control logic. In contrast, MinBD does not need to include large buffers and the associated buffered-mode control logic, instead using only a smaller side buffer. MinBD also removes the dependence on efficient buffer power-gating that AFC requires for energy-efficient operation at low loads. We quantitatively compared MinBD to AFC in §IV and demonstrated better energy efficiency at all network loads.

Reducing cost of buffered routers: Empty buffer bypassing [39], [27] reduces buffered router power by allowing flits to bypass input buffers when empty. However, as our evaluations (which faithfully model the power reductions due to buffer bypassing) show, this scheme reduces power less than our new router design: bypassing is only effective when buffers are empty, which happens more rarely as load increases. Furthermore, buffers continue to consume static power, even when unused. Though both MinBD and empty-buffer-bypassed buffered routers avoid buffering significant traffic, MinBD further reduces router power by using much smaller buffers.

Kim [22] proposed a low-cost buffered router design in which a packet uses a buffer only when turning, not when traveling straight along one dimension. Unlike our design, this prior work does not make use of deflection, but uses deterministic X-Y routing. Hence, it is not adaptive to different traffic patterns. Furthermore, its performance depends significantly on the size of the turn-buffers. By using deflection, MinBD is less dependent on buffer size to attain high performance, as we argued in §VI-E. In addition, [22] implements a token-based injection starvation avoidance scheme which requires additional communication between routers, whereas MinBD requires only per-router control to ensure side buffer injection.

VIII. CONCLUSION

In this paper, we present MinBD, a new, minimally-buffered deflection router design. MinBD combines deflection routing with a small buffer, such that some network traffic that would have been deflected is placed in the buffer instead. By using the buffer for only a fraction of network traffic, MinBD makes more efficient use of a given buffer size than a conventional input-buffered router. Its average network power is also greatly reduced: relative to an input-buffered router, buffer power is much lower, because buffers are smaller. Relative to a bufferless deflection router, dynamic power is lower, because deflection rate is reduced with the small buffer. Our evaluations show that MinBD performs competitively with the input-buffered router design, and has the best energy efficiency of all evaluated networks. We conclude that a router design which augments bufferless deflection routing with a small buffer to reduce deflection rate is a compelling design point for energy-efficient, high-performance on-chip interconnect.

REFERENCES

[1] P. Abad et al. Rotary router: an efficient architecture for cmp interconnection networks. *ISCA-34*, 2007.

[2] P. Baran. On distributed communications networks. *IEEE Trans. on Comm.*, 1964.

[3] S. Borkar. NoCs: What's the point? NSF Workshop on Emerging Tech. for Interconnects (WETI), Feb. 2012.

[4] S. Borkar. Thousand core chips: a technology perspective. *DAC-44*, 2007.

[5] S. Borkar. Future of interconnect fabric: a contrarian view. *SLIP '10*, 2010.

[6] P. Bose. The power of communication: Trends, challenges (and accounting issues). NSF WETI, Feb. 2012.

[7] K. Chang et al. Adaptive cluster throttling: Improving high-load performance in bufferless on-chip networks. SAFARI technical report TR-2011-006: http://safari.ece.cmu.edu/tr.html, Sep 2011.

[8] D. E. Culler et al. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufmann, 1999.

[9] W. Dally and B. Towles. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, 2004.

[10] W. J. Dally and B. Towles. Route packets, not wires: On-chip interconnection networks. *DAC-38*, 2001.

[11] S. Eyerman and L. Eeckhout. System-level performance metrics for multiprogram workloads. *IEEE Micro*, 28:42–53, May 2008.

[12] C. Fallin, C. Craik, and O. Mutlu. CHIPPER: A low-complexity bufferless deflection router. *HPCA-17*, 2011.

[13] C. Fallin et al. MinBD: Minimally-buffered deflection routing for energy-efficient interconnect. SAFARI technical report TR-2011-008: http://safari.ece.cmu.edu/tr.html, Sep 2011.

[14] C. Gómez et al. Reducing packet dropping in a bufferless noc. *Euro-Par-14*, 2008.

[15] B. Grot et al. Kilo-NOC: a heterogeneous network-on-chip architecture for scalability and service guarantees. *ISCA-38*, 2011.

[16] A. Hansson et al. Avoiding message-dependent deadlock in network-based systems-on-chip. *VLSI Design*, 2007.

[17] M. Hayenga, N. E. Jerger, and M. Lipasti. Scarab: A single cycle adaptive routing and bufferless network. *MICRO-42*, 2009.

[18] W. Hillis. *The Connection Machine*. MIT Press, 1989.

[19] Y. Hoskote et al. A 5-GHz mesh interconnect for a teraflops processor. *IEEE Micro*, 2007.

[20] W. Dally and C. Seitz. Deadlock-free message routing in multiprocessor interconnection networks. *IEEE Trans. Comp.*, 1987.

[21] S. A. R. Jafri et al. Adaptive flow control for robust performance and energy. *MICRO-43*, 2010.

[22] J. Kim. Low-cost router microarchitecture for on-chip networks. *MICRO-42*, 2009.

[23] S. Konstantinidou and L. Snyder. Chaos router: architecture and performance. *ISCA-18*, 1991.

[24] D. Kroft. Lockup-free instruction fetch/prefetch cache organization. *ISCA-8*, 1981.

[25] J. Laudon and D. Lenoski. The SGI Origin: a ccNUMA highly scalable server. *ISCA-24*, 1997.

[26] C.-K. Luk et al. Pin: building customized program analysis tools with dynamic instrumentation. *PLDI*, 2005.

[27] G. Michelogiannakis et al. Evaluating bufferless flow-control for on-chip networks. *NOCS*, 2010.

[28] T. Moscibroda and O. Mutlu. A case for bufferless routing in on-chip networks. *ISCA-36*, 2009.

[29] G. Nychis et al. Next generation on-chip networks: What kind of congestion control do we need? *Hotnets-IX*, 2010.

[30] H. Patil et al. Pinpointing representative portions of large Intel Itanium programs with dynamic instrumentation. *MICRO-37*, 2004.

[31] B. Smith. Architecture and applications of the HEP multiprocessor computer system. *SPIE*, 1981.

[32] B. Smith. Personal communication, 2008.

[33] A. Snavely and D. M. Tullsen. Symbiotic jobscheduling for a simultaneous multithreaded processor. *ASPLOS-9*, 2000.

[34] Standard Performance Evaluation Corporation. SPEC CPU2006. http://www.spec.org/cpu2006.

[35] M. Taylor, J. Kim, J. Miller, and D. Wentzlaff. The raw microprocessor: A computational fabric for software circuits and general-purpose programs. *IEEE Micro*, Mar 2002.

[36] M. Thottethodi, A. Lebeck, and S. Mukherjee. Self-tuned congestion control for multiprocessor networks. *HPCA-7*, 2001.

[37] S. Tota et al. Implementation analysis of NoC: a MPSoC trace-driven approach. *GLSVLSI-16*, 2006.

[38] H. Wang et al. Orion: a power-performance simulator for interconnection networks. *MICRO-35*, 2002.

[39] H. Wang, L. Peh, and S. Malik. Power-driven design of router microarchitectures in on-chip networks. *MICRO-36*, 2003.

[40] D. Wentzlaff et al. On-chip interconnection architecture of the tile processor. *IEEE Micro*, 27(5):15–31, 2007.