NoRD: Node-Router Decoupling for Effective Power-gating of On-Chip Routers

Lizhong Chen   Timothy M. Pinkston
Ming Hsieh Department of Electrical Engineering
University of Southern California
Los Angeles, CA
{lizhongc, tpink}@usc.edu

Abstract

While power-gating is a promising technique to mitigate the increasing static power of a chip, a fundamental requirement is for the idle periods to be sufficiently long enough to compensate for the power-gating and performance overhead. On-chip routers are potentially good targets for power optimizations, but few works have explored effective ways of power-gating them due to the intrinsic dependence between the node and router – any packet (sent, received or forwarded) must wakeup the router before transferring, thus breaking the potentially long idle period into fragmented intervals. Simulation shows that directly applying conventional power-gating techniques would cause frequent state-transitions and significant energy and performance overhead. In this paper, we propose NoRD (Node-Router Decoupling), a novel power-aware on-chip network approach that provides for power-gating bypass to decouple the node's ability for transferring packets from the powered-on/off status of the associated router, thereby maximizing the length of router idle periods. Full system evaluation using PARSEC benchmarks shows that the proposed approach can substantially reduce the number of state-transitions, completely hide wakeup latency from the critical path of packet transport and eliminate node-network disconnection problems. Compared to an optimized conventional power-gating technique applied to on-chip routers, NoRD can further reduce the router static energy by 29.9% and improve the average packet latency by 26.3%, with only 3% additional area overhead.

1. Introduction

In recent years, power has become a critical design constraint, driving the microarchitecture design toward the paradigm of chip multiprocessors (CMPs). As a key component in CMPs, the network-on-chip (NoC) is the backbone for supporting communications among multiple cores. It is thus very important for NoCs to work efficiently and effectively to achieve both high performance and low power. However, recent studies show that NoCs can draw a substantial percentage of a chip's power, by up to 10%~36% [8, 9, 28]. In particular, the static power of routers has become a significant contributor of power consumption, consisting of more than 35% and 43% of the total NoC power at 45nm and 32nm processes, respectively (more details in Section 2). Unfortunately, the impact of static power will only get worse with continued scaling of transistor feature size and chip operating voltage.

Power-gating is a useful circuit-level technique applicable to power-aware architectures to mitigate the increasing static power, especially for circuit blocks that exhibit enough idleness [10, 19]. While power-gating in general is a promising technique, applying it directly to on-chip routers has been elusive as doing so requires several fundamental problems to be addressed in order to maximize energy-savings and minimize performance penalties. First, intermittent packet arrivals may cause a large number of idle periods to fall below the breakeven time needed to compensate for power-gating overhead, reducing the opportunity to apply power-gating techniques usefully. Second, packets encountering gated-off routers suffer additional transport latency to wait for routers to wake up and are likely to experience successive wakeup latencies on the critical path if routed over multiple hops. Third, as a gated-off router essentially disconnects the associated node from the rest of the network, the power-gating opportunity is upper bounded by the local node's traffic and none of the local resources (e.g., cache and directory) can be accessed by other nodes, unless connectivity is somehow supported another way. Without solving these fundamental problems, the effectiveness of applying power-gating to on-chip routers is severely limited.

The above problems are all caused by node-router dependence – whether a node can send, receive or forward a packet depends directly on the on/off status of the associated router of that node. In this paper, we propose *NoRD (Node-Router Decoupling)*, a novel approach that provides separate power-gating bypass to decouple the node's ability for transferring packets from the status of the router. This approach avoids unnecessary router wakeups and, more importantly, the associated performance penalty and energy overhead. NoRD effectively increases the length of idle periods, removes wakeup latency from the critical path, and eliminates power-gating disconnection problems. The main contributions of this paper are the following:

- Fundamental and critical problems of applying conventional power-gating techniques to on-chip network routers are identified;
- The concept of Node-Router Decoupling and a power-gating bypass technique to implement NoRD are proposed, which provides a unified and effective solution to the aforementioned problems;
- Full system simulations show a significant improvement in the use of power-gating with NoRD as compared to directly applying power-gating with conventional techniques.

The rest of the paper is organized as follows. Section 2 provides more background on the static power of routers and the power-gating technique. Section 3 highlights the problems of power-gating on-chip routers and motivates the need for a better approach. Section 4 explains the details of the proposed NoRD design. Section 5 discusses our evaluation methodology, and Section 6 presents simulation results. Finally, related work is summarized in Section 7, and Section 8 concludes the paper.
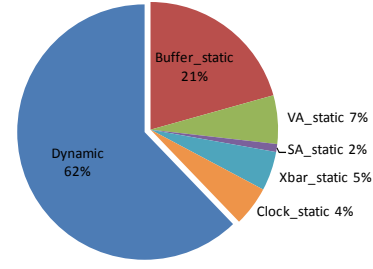
2. Background

2.1 Static Power of On-chip Routers

The static power of CMOS circuitry has been increasing substantially in recent years due to the continued scaling of transistor feature size and chip operating voltage. As a major
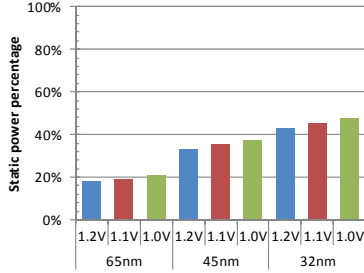
(a) Static power percentage   (b) Router power decomposition

Figure 1: Static power vs. dynamic power of on-chip routers.

component of multicore chips, on-chip networks consume around 10%~36% of a chip's power, as shown in recent industrial and research chips [8, 9, 28]. A considerable amount of NoC power comes from static consumption. To study the significance of NoC static power and the impact of technology scaling, Figure 1(a) plots the percentage of static power of on-chip routers at 3GHz for various manufacturing generations and operating voltages. Results are obtained from the Orion 2.0 [13] power model. To reflect realistic workloads, Orion is fed with statistics from full system simulation – Simics [21] plus GEMS [22] – running multi-threaded PARSEC 2.0 benchmarks [2] (more details of the simulation infrastructure are described in Section 5). As shown in the figure, the percentage of static power consumption increases as the feature size and operating voltage decrease, from 17.9% at 65nm and 1.2V, to 35.4% at 45nm and 1.1V, to 47.7% at 32nm and 1.0V. This trend clearly illustrates that the static power of on-chip routers has become a significant part of the overall router power consumption, and only worsens for process technologies beyond 45nm. Figure 1(b) further breaks down the total power consumption of on-chip routers at 45nm with 1.0V into dynamic and various static components. As can be seen, buffers consume 55% of the static power (21% of the total power) while other router components consume 45% of the static power (17% of the total power). This indicates that the static power consumption in router components other than buffers is significant and that appropriate techniques need to be adopted to reduce all contributors to static power.

2.2 Power-gating Techniques

One of the most effective techniques to mitigate the static power of a circuit block is power-gating as it cuts off the power supply of that block, which is the source of leakage currents in both subthreshold conduction and reverse-biased diodes. It is implemented by inserting appropriately sized header or footer transistor(s) with high threshold voltage (non-leaky "sleep switch") between Vdd and the block, or the block and GND, as illustrated in Figure 2(a). By asserting or de-asserting the sleep signal, the supply voltage to the power-gated block can be turned on and off.

Figure 2(b) depicts the key intervals of power-gating. At time t_0, the sleep signal is asserted and distributed to the sleep transistor with certain overhead energy. At t_1, this signal arrives at the sleep transistor and turns it off, so the virtual Vdd starts to drop. Correspondingly, the leakage current also decreases and the cumulative energy savings start to increase. From this moment, the block stays in the power-gated off state until t_2 when the sleep signal is de-asserted and distributed again, initiating the wakeup process. From t_2 to t_3, another energy overhead is incurred in distributing the sleep signal and waking up the gated-off block. The cumulative energy savings stop increasing at t_3 when the

virtual Vdd restores to full Vdd and the wakeup process concludes. Consequently, an important parameter in power-gating is the "breakeven time" (BET), which is defined to be the minimum number of consecutive cycles that a gated block needs to remain in idle state before being awoken to offset power-gating energy overhead [19, 20]. Prior research using analytical modeling and simulation [10, 23] estimate the BET value to be around 10 cycles for functional units and on-chip routers under current technology parameters.

2.3 Use of Power-gating

Although power-gating can reduce power, it can also reduce system performance. This is because a powered-off block cannot perform the assumed functions temporarily, and waking up the block takes an additional wakeup delay, thus potentially stalling system progress. Therefore, effective use of power-gating should achieve two objectives in a balanced way:

(1) Maximize net energy savings, which means to maximize the idleness of unneeded functional blocks in order to increase the cumulative energy savings while reducing the associated energy overhead as much as possible;

(2) Minimize performance penalty, which means to partially or completely reduce/hide the wakeup latency of needed functional blocks, so that execution can continue with minimal delay.

While power-gating has been used successfully in cores and execution units [10, 19, 20], only recently has research started to investigate its application to on-chip network routers [23, 24, 25]. However, as discussed shortly in the next section, due to the node-router dependence in on-chip networks, the conventional way of power-gating routers is ineffective in achieving the energy and performance objectives. Several fundamental and critical problems must be addressed to mitigate costly frequent state-transitions and performance overhead that comes with applying the conventional technique.

3. Motivation

3.1 Conventional Power-gating of On-chip Routers

The on-chip network is responsible for connecting the various components within a CMP, where each node may consists of a processor core, caches, and an associated router. *Node-router dependence* means that the ability for a node to send, receive or forward a packet depends directly on the on/off status of the associated router. For example, a node can inject a packet into the network only when the associated router is in the powered-on state. Conversely, routers become idle when the associated nodes have no packet to send, receive or forward. Our full system simulation results show that on-chip routers can be idle 30%~70% of the time
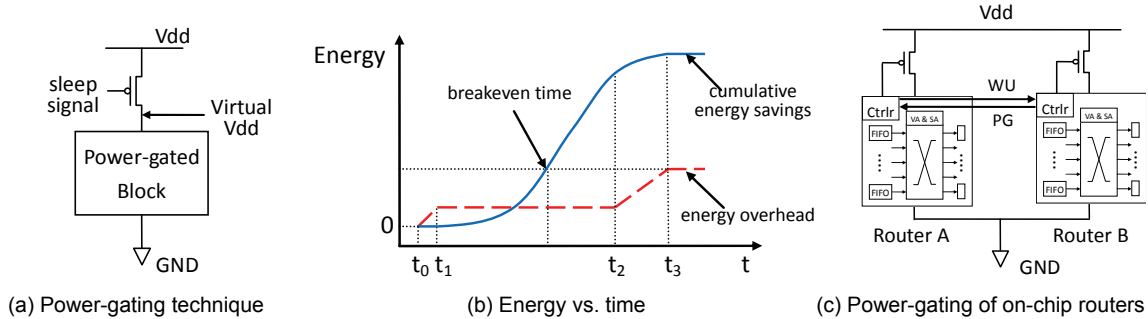
Figure 2: Power-gating technique and its application to on-chip routers.

(a) Power-gating technique   (b) Energy vs. time   (c) Power-gating of on-chip routers

(with *x264* having the lowest of 30.4% and *blackscholes* having the highest of 71.2%), depending on the physical location of the routers in the NoC and the load intensity of the applications. Therefore, power-gating techniques can be applied to on-chip routers to take advantage of their idleness.

When the internal datapaths of a router are empty (i.e., input ports, output latches, and the crossbar), the router microarchitecture can be power-gated off to save static power after notification of all its neighbors. Figure 2(c) shows an example of power-gating router *B* and handshaking with one of its upstream routers, router *A*. A canonical wormhole router [3] is assumed, which consists of routing computation (RC), VC allocation (VA), switch allocation (SA) and switch traversal (ST), with another stage of link traversal and buffered writing (LT). A small non-power-gated controller is added in the router to monitor the emptiness of the datapath and the wakeup signals from neighbors. When the datapath of router *B* is detected as empty and the *WU* (wakeup) signals are clear, the controller in router *B* asserts a sleep signal to put router *B* into gated-off state and asserts a *PG* (power-gate) signal to notify router *A*. Upon detecting the asserted *PG* signal, router *A* tags the output port that leads to router *B* as being power-gated and hence becomes unavailable in the SA stage[1]. Later, after router *B* is power-gated, some packet in router *A* or another neighbor of router *B* may request an output port to router *B* in the SA stage, triggering the *WU* signal to be asserted which causes the controller in router *B* to de-assert its sleep signal. The packet will then be stalled in the SA stage while waiting for router *B* to wake up and de-assert the *PG* signal. According to previous studies [23, 25], the wakeup latency for on-chip routers under typical technology parameters is a few nanoseconds, or around 10~20 cycles depending on the frequency. In what follows, we use the term *conventional power-gating* of routers to refer to the above mechanism of applying conventional power-gating to on-chip routers.

3.2 Intensified BET Limitation

A major obstacle to achieving effective power-gating of on-chip routers is the intensified limitation caused by breakeven time (BET). It has been observed that, when applying power-gating to functional units, the BET limitation may cause large energy penalty for some applications where functional units do not exhibit long enough idle periods [19]. Unfortunately, when applying conventional power-gating to on-chip routers, the BET limitation
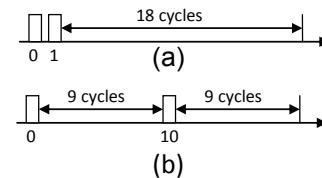

Figure 3: Intermittent packet arrival.

becomes much more prevalent due to intermittent packet arrivals seen by the routers. Figure 3 illustrates the problem even in the case where the NoC has substantial idleness, as given by a low average arrival rate of 0.1 flits/cycle (i.e., 10% traffic load). In (a), with two successive single-flit packets arriving in the first two cycles, the router has up to 18 idle cycles for useful power-gating; whereas in (b), discrete packet arrivals cut down idle periods to below the BET, leading to an energy *penalty* as opposed to savings if power-gated. Our evaluation on PARSEC benchmarks shows that the number of idle periods having a length less than or equal to the BET constitutes more than 61% of the total number of idle periods. Thus, on the one hand, routers on average exhibit very good idleness that could benefit from applying power-gating, but on the other hand, a large percentage of these idle periods are too short to meet the BET requirement as any sending, receiving or forwarding operation of a node would generate packets for the associated router to process, thus severely limiting the effectiveness of conventional power-gating of routers.

One direct way to address this problem is to reduce the BET through better circuit-level design or advanced manufacturing processes, which unavoidably have physical limitations (e.g., transistor sizing of the inverter-chain has limited ability in mitigating the energy overhead of sleep-signal distribution). Another possibility is to apply conventional power-gating to smaller individual components within each router, such as per input port or per virtual channel [24, 25]. This method, however, can only mitigate the impact of the BET problem as individual components have only slightly longer idle period, and even if the BET condition is satisfied, many power-gated cycles are wasted to offset the energy overhead. Moreover, this requires prohibitive hardware implementation overhead. For example, there are 35 power domains in a single router in [25] to implement this method of power-gating in addition to the complex coordination needed among different components, which incurs significant energy and area overhead with considerable design effort. Thus, a much more effective way of removing the dependence between the node and router is needed, so as to combat the BET limitation from the source by reducing the number of wakeups while maintaining the ability to transport packets in the NoC.

---

[1] To ensure the receiving of packets that are already in ST and LT stages, either router *B* needs to wait two more cycles before deciding to enter gated-off state, or *WU* should be generated early enough.

3.3 Cumulative Wakeup Latency in Multi-hop Networks

Just as the BET limitation of energy-savings is magnified in power-gated on-chip routers, the wakeup latency problem is also exacerbated in NoC environments, which affects performance negatively. Due to the node-router dependence, conventional power-gating of routers requires routers to be in on-state to forward packets, which makes the wakeup latency exposed directly to the critical path of packet transport to downstream routers. A packet routed in a multi-hop NoC can experience wakeup latency multiple times as routers at many hops along the path could be gated-off. To make things worse, power-gating works best when load rates are low, but in those situations more routers are in the gated-off state, making packets more likely to encounter multiple wakeups. One approach is to use early wakeup signal generation (e.g., generate the wakeup signal as soon as the output port is computed). However, this has limited ability to hide router wakeup latency, e.g., 3 cycles maximum out of the 10~20 cycles of wakeup latency for a 4-stage pipeline. Look-ahead wakeup is also possible [23, 25], in which the candidate router monitors all the wakeup signals two hops away so that it can hide at most 6 cycles of wakeup latency. This still limited technique requires monitoring hardware that is very complex and expensive to implement as every router essentially has to monitor every input port in up to 12 routers within a 2-hop distance, assuming a 2-D mesh topology. A much better approach would be to effectively *remove* the wakeup latency from the critical path by providing bypass of powered-off routers, as proposed in Section 4.

3.4 Disconnection Problem

The third major and most obvious problem in applying conventional power-gating to on-chip routers is the network disconnection problem. This problem is caused also by the node-router dependence, as whenever a router is power-gated off, the associated node is disconnected from the rest of the network. The disconnection problem impacts system in two ways. First, the local node cannot send/receive packets to/from the network if the associated router is powered-off, which limits the opportunity of power-gating to only those cases when the core and cache associated with the node are completely idle. Second, remote nodes cannot access any resource on the local node either, particularly the cache line and coherence directory. For a typical shared last level cache (LLC) configuration, this essentially decreases the effective cache size. For example, if half of the routers are power-gated off, the accessible LLC size available to the remaining nodes is reduced by 50%. Especially worth noting is that a private LLC does not help much due to the maintaining of cache coherence protocols. For instance, a dirty line in the private LLC of the local node is the unique last copy of the data in the entire system. Any other request to this line from remote nodes must wakeup the local router to access the data and resume correct execution, even if the local core is idle. Therefore, a more effective way to circumvent powered-off routers and maintain the connectivity of on-chip resources using some alternative path is needed.

4. Proposed Scheme: NoRD

In this section, we propose *NoRD (Node-Router Decoupling)*, a novel approach that removes the intrinsic dependence between nodes and routers, solving all the aforementioned problems unaddressed by conventional power-gating of on-chip routers.

4.1 The Basic Idea

The proposed approach is based on the simple idea of breaking node-router dependence via wakeup-avoidance decoupling bypass paths. Recall that in conventional power-gating of routers, due to the node-router dependence, any incoming packet from either a local node or other nodes would first have to wake up the gated-off router before further packet transport could occur. This wakeup incurs energy overhead and performance penalty on each occurrence. By providing decoupling bypass for each router, the ability to transport packets in the network is decoupled from the on/off status of the routers. This solves all three problems of conventional power-gating of routers. First, packets (sent, received or forwarded) have the option to go through bypass paths instead of powering-on the routers to continue progress, thus avoiding unnecessary wakeups and the associated energy overhead which causes BET in the first place. Second, bypass allows packets to be transferred *while* the router is being awoken, which removes the wakeup latency completely from the critical path of packet transport. Third, when the associated router is powered-off, the local node can still be connected with the rest of the network through the decoupling bypass paths, thus eliminating the disconnection problem.

While NoRD conceptually is a simple yet attractive solution, implementing decoupling bypass that provides chip-wide connectivity even when many or all routers are gated-off and transition between the gated-on/off state is not straightforward. In the proposed design, we add internal bypass paths in each router that can forward packets directly from a selected input port to the network interface (NI) and then forward the packets from the NI back to a selected output port. The input/output port pairs from all routers form – in the worst case – a unidirectional ring across the chip, so that all the NIs are always connected. The resulting bypass paths, together with all remaining paths provided by the normal deadlock-free routing algorithm, allow packets to be transported without deadlock in NoCs comprised of any combination of powered-on and powered-off routers. In the rest of this section, we present the detailed design of NoRD, addressing the construction of bypass paths, the implementation of NI forwarding, the transition and interface between routers in bypass mode and normal mode, the avoidance of deadlock and other network abnormalities under the presence of both on and off routers, and asymmetric wakeup threshold to further increase the efficiency of NoRD.

4.2 Decoupling Bypass

Without loss of generality, we start by describing the microarchitecture of bypass using a 4x4 2D mesh as an example. Decoupling bypass is achieved through two-level coordination. At the chip level, an input port (referred to as a Bypass Inport) and an output port (referred to as a Bypass Outport) from each router are chosen in a way such that, collectively across the network, they form a unidirectional ring (referred to as Bypass Ring) connecting all nodes, as shown in Figure 4(a). At individual router level, two datapaths are added as follows. In order to inject packets from the local node (e.g., processor core), a datapath is added from the NI input to the Bypass Outport (the bottom bold line in Figure 4(b)). In order to receive packets destined to the local node from the network, a second datapath is added from the Bypass Inport to the NI outport to eject packets from the router (the top bold line in Figure 4(b)). The bypass paths consisting of minimal hardware described here are not power-gated.
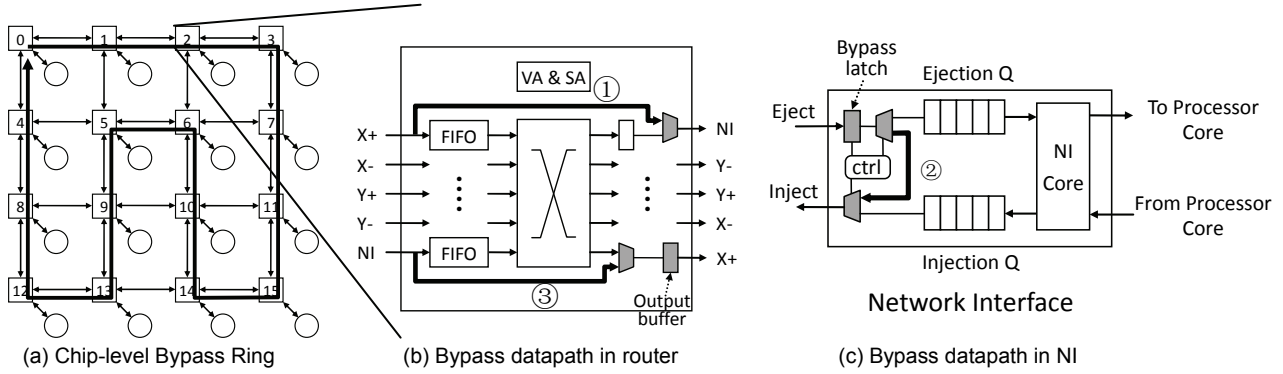
Figure 4: Decoupling bypass (shaded components in (b) and (c) are not power-gated).

(a) Chip-level Bypass Ring   (b) Bypass datapath in router   (c) Bypass datapath in NI

To forward packets through a powered-off router, a bypass path from the router's Bypass Inport to its Bypass Outport is established through the node's NI. Flits are ejected from the powered-off router to the NI and injected back into the same router along the path of the Bypass Ring, as shown in Figure 4(c). In a typical NoC with wormhole switching, the NI is responsible for accepting data from the node and encapsulating it into packets and flits (NI core), allocating a virtual channel and checking flow control credits in the NI input port of the associated router, and injecting the formatted flits into the network. Receiving data from the network to the node has a similar but reversed process. Now, to implement router bypassing through the NI of the node, we add a latch and a demultiplexer ahead of the ejection queue, insert a multiplexer after the NI's injection queue, and create a path between the input and output ports of the NI according to Figure 4(c). With this forwarding path, a flit can now be forwarded from the gated-off router's Bypass Inport to its Bypass Outport in three stages, as annotated in Figure 4(b) and (c): ① at the end of link traversal, instead of being written into the router's input buffer as done when the router is powered-on, the flit is written directly into the NI bypass latch through the bypass datapath; ② based on the packet's destination header bits, the NI either sinks this flit in the local node or forwards the flit by allocating a VC (and checking its credits); ③ the flit is re-injected into the power-gated router's Bypass Output through the bypass datapath. The bypass datapath is enabled only when the router is in the power-gated off state.

The above two-level coordination essentially decouples nodes from the on/off status of routers, as now a node can send, receive and forward packets through the decoupling bypass even if the associated router is in the gated-off state. Moreover, it ensures the connectivity of all nodes. Packets can route through a combination of Bypass Ring paths to circumvent gated-off routers and normal paths of gated-on routers to minimize hop count. Even in the extreme case of all routers being gated-off, packets can still traverse along the Bypass Ring to reach any destination.

Owing to the decoupling bypass that provides network connectivity in all cases, deadlock-free adaptive routing based on Duato's Protocol [4] is easily supported. Escape resources are comprised of the unidirectional ring formed by the (Bypass Inport, Bypass Outport) pairs in both gated-on and gated-off router state, where two VCs can be used to break cyclic dependence. Additional VCs can be used as adaptive resources for adaptive routing over the NoC.

The deadlock- and livelock-free routing of NoRD is as follows. Every router has adaptive VCs and escape VCs (powered-off routers have no VCs but still have the corresponding adap-

tive/escape latches for bypassing). At normal routers, packets on adaptive VCs use minimal adaptive routing to choose the next hop, but packets on escape VCs are confined to choose the Bypass Outport (i.e., move along the bypass ring) and confined to escape VCs until destination. For packets on adaptive VCs, misrouting occurs only when all of the downstream routers on the minimal path are powered-off AND the Bypass Outport forces a detour (note that the Bypass Outport could, in fact, also be on the minimal path). In that case, packets must choose the Bypass Outport to traverse to next router (could be either normal or off) misrouted by one hop. However, packets are still allowed to remain on adaptive VCs for normal routers or the corresponding adaptive latches for bypassed routers (i.e., the entire set of adaptive resources) if the total misrouted hops are below a threshold; otherwise packets are forced to enter escape VCs (or the corresponding escape latches for bypassed router) and route along the unidirectional ring without returning to adaptive resources until the destination is reached. At the next router, if packets are still on adaptive VCs, they will repeat the above process (i.e., use minimal adaptive routing if available on the bypass ring or mesh, or enter escape resources on the Bypass Ring if needed) until reaching the destination. No U-turns are allowed at any hop. The above routing for NoRD follows Duato's Protocol for deadlock-free adaptive routing as the escape VCs on the Bypass Ring have no cycles in the extended channel dependence graph and the adaptive channels allow for fully adaptive routing. As detoured packets have a cap on the number of misroutes allowed before being forced to enter escape VCs with a bounded hop count, NoRD avoids both deadlock and livelock. Also, any additional hops from detours are partially offset by gains in completely hiding router wakeup latency as compared to conventional power-gating and reduced per hop latency of the bypass path. Finally, starvation for NI resources by the local node is easily avoided by granting priority over bypass traffic to the local node if not served for a predetermined number of consecutive cycles. However, this should happen rarely as the router is assumed to be power-gated off only when the load is low and contention is minimal.

4.3 Transition between Gated-on and Gated-off States

To transition between gated-on and gated-off states and to interface with neighboring routers for correct flow control, several handshaking signals are needed as illustrated in Figure 5. In this example, we focus on the state-transition of router B, and the bypass of router B is from router A through the NI of router B to router D.
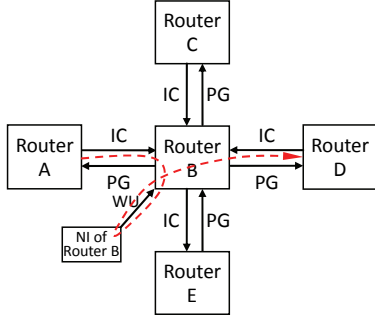
Figure 5: Handshaking in NoRD.
PG: power-gate, WU: wakeup, IC: incoming

Figure 6: Impact of powering-on routers.

To transition from gated-on to gated-off state, similar to the conventional power-gating mechanism described in Section 3.1, if router *B* is empty and both *IC* and *WU* are clear (these two signals will be explained shortly), it asserts the *PG* signals, enables bypass and goes into gated-off state by asserting the sleep signal (not shown). Upon detecting the asserted *PG* signal, routers *C*, *D* and *E* tag the output port that leads to router *B* as power-gated (and becomes unavailable in the SA stage) and stop tracking credits, while router *A*, which is the Bypass Ring upstream router, sets the credit of each VC in that output port to 1 as router *B* now has only one output buffer available as shown in Figure 4(b). To ensure the receiving of packets that are already in the ST and LT stages of the neighboring routers, an *IC* (incoming) signal is generated at the beginning of SA if there is a flit in the SA stage and propagates to router *B*. In this way, the *IC* signal is always two cycles ahead of flits to notify router *B* that a flit is incoming and router *B* should not enter into gated-off state. Finally, for any flits that are in the VA and SA stages of routers *C*, *D* and *E*, they will restart the pipeline from RC using the new output port availability information as they are still in the input channel. Note that these flits must be head flits; otherwise if the head flits have left router *C*/*D*/*E* to *B* but body/tail flits have not yet arrived at router *B*, then the virtual channel is not de-allocated and router *B* is not considered as empty.

To transition router *B* from gated-off state to gated-on state, the *WU* signal first needs to be generated according to a wakeup metric. Ideally, the wakeup metric should de-assert *WU* when the load is low, and assert the signal when it is above a threshold when the load becomes high. A naïve way is to use the number of flits transmitted by the gated-off router in a fixed period of time, but this may not necessarily generate a wakeup signal when the load is high as flits could be stalled due to network congestion. Another traditional metric is to use router buffer utilization [27], which also is not suitable as input buffers are not used in the gated-off state. As all traffic to gated-off routers are forwarded through the NI and allocated a VC there to (re)inject into the network, we use as a threshold parameter the number of VC requests at the local NI over a period of time (10 cycles) for the wakeup metric. This metric works for both low and high load as the number of VC requests goes up even if the flits are stalled, and it remains valid in the extreme case when all the routers are gated-off, as the wakeup signal is generated locally.

With the number of VC requests used as threshold wakeup metric, the operation of turning on a gated-off router is straight-forward. When the *WU* signal is asserted, router *B* starts to wake up while the bypass is still functioning. When wakeup finishes, router *B* de-asserts the *PG* signal. Upon detecting the de-asserted *PG* signal, routers *C*, *D* and *E* reset the credits to full while router
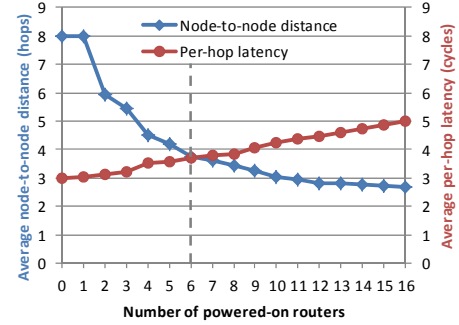
A adds back (full-1) credits. Once the flit in the NI bypass datapath is written into the input buffer of router *B*, the bypass of router *B* is disabled to complete the state-transition.

4.4 Asymmetric Wakeup Thresholds

While previous subsections describe the necessary operations to keep NoRD functional, the efficiency of NoRD can be increased using *asymmetric wakeup thresholds*. For certain topologies and constructions of the Bypass Ring, some routers may have greater impact on performance than others based on their location in the NoC. For example, powering on Routers 4 and 5 in Figure 4(a) has larger performance benefits than powering on Routers 0 and 1, as the former provide a shortcut to route packets that would otherwise be detoured through 9->13->12->8. Therefore, taking the placement of bypass paths and routers into account, additional performance gains can be obtained.

To differentiate between routers in NoRD, asymmetric wakeup thresholds can be used. For example, NoC routers can fall broadly under two classes – *performance-centric* and *power-centric* – based on their importance, where a low wakeup threshold is assigned to the *performance-centric* class and a high wakeup threshold is assigned to the *power-centric* class. The intuition behind this is to wake up early a few performance-critical routers while waking up late the rest (majority) of the routers. In this way, not only performance improves due to the added shortcuts in routing paths, but also more static power can be saved by allowing non-performance-critical routers to stay in the gated-off state for a longer time. As a threshold metric is needed for wakeup anyway, no additional hardware is required.

To select the set of routers that are more critical to performance, we wrote a short off-line program based on the Floyd-Warshall all-pair shortest path algorithm [7]. Figure 6 plots the best node-to-node average distance and per-hop latency that can be achieved with a given number of powered-on routers for the 2-D mesh example in Figure 4(a). As expected, with more routers turned on, the average hop distance between nodes in NoRD decreases rapidly due to the added flexibility in routing paths. Meanwhile, more packets are routed through the normal pipeline of powered-on routers instead of the simpler and shorter bypass pipeline, thus gradually increasing the per-hop latency. Figure 6 also shows that, by turning on six routers, the average hop distance can be greatly reduced with moderate increase in the per-hop latency, indicating a viable trade-off point. The corresponding router set that achieves this data point consists of Routers 4, 5, 6, 7, 13 and 14 in Figure 4(a). In this example, these routers are designated as the *performance-centric* routers, and the remaining routers are classified as the *power-centric* routers. Other classifica-

tions are still possible and an optimal classification could be determined dynamically with comprehensive consideration of topology, traffic patterns, bypass placement, and routing algorithm. For instance, the routing algorithm may adaptively steer packets to a few *performance-centric* routers and the rest of the routers can be designated as *power-centric* routers. While further work can be conducted to investigate the design complexity of finding the optimal classification and the trade-off in doing so, this falls outside the scope of this paper. Here, we intend only to show that asymmetric wakeup threshold, even with a simple dual-mode classification, can provide additional benefits in both performance and energy to complement the proposed decoupling bypass mechanism.

4.5 Impact of NoRD on Energy and Performance

We mentioned before that there are two primary objectives when using power-gating techniques. Here, we analyze the impact of NoRD on achieving these two objectives to highlight the benefits of the proposed decoupling approach.

Impact on Net Energy Savings

NoRD maximizes the opportunity for saving energy by allowing fragmented idle periods that are even shorter than the BET to be exploited, which is not possible in conventional power-gating of routers. Moreover, by steering short packet spikes to bypass paths without waking up the routers, the energy overhead in distributing the sleep signal and powering-on the router is also largely avoided. Therefore, NoRD is able to increase the cumulative energy savings while reducing the power-gating energy overhead.

Impact on Performance

NoRD minimizes the performance penalty of power-gating techniques from the following aspects: (1) the use of decoupling bypass reduces the number of state-transitions and, hence, avoids the wakeup latency when routers do not need to be turned on; (2) when router wakeup is unavoidable, decoupling bypass provides temporary paths for packets while the router is being awoken, thus hiding wakeup latency; (3) a few performance-centric routers with low thresholds can be awoken earlier to guard performance. With these features, NoRD can greatly reduce the performance penalty of conventional power-gating of routers, as the following analysis shows.

5. Evaluation Methodology

5.1 Simulator Configuration

The proposed NoRD scheme is evaluated quantitatively under full-system simulation using Simics [21], with GEMS [22] and Garnet [1] for detailed timing of the memory system and on-chip network. Orion 2.0 [13] is integrated in Garnet for NoC power and area estimation using technology parameters from an industrial standard 45nm CMOS process and 1.1V operating voltage. The saved static power is modeled after [10] and the overhead is modeled after [10, 13]. A wakeup latency of 12 cycles is used assuming a 4ns wakeup delay and 3GHz frequency, and 3 cycles can be hidden when the early wakeup technique [23] is applied. We modify the simulators to model all the key additional hardware for power-gating and bypass, including the extra power consumption in the NI buffering and forwarding logic. The additional dynamic (static) power of the NI in NoRD is lumped into router dynamic (static) power to provide fair comparison across different schemes. Step 2 in Figure 4(c) that checks VC availability in the

Table 1: Key parameters used in simulation.

Core model	Sun UltraSPARC III+, 3GHz
Private I/D L1\$	32KB, 2-way, LRU, 1-cycle latency
Shared L2 per bank	256KB, 16-way, LRU, 6-cycle latency
Cache block size	64Bytes
Coherence protocol	MOESI
Network topology	4x4 and 8x8 mesh
Router	4-stage, 3GHz
Virtual channel	4 per protocol class
Input buffer	5-flit depth
Link bandwidth	128 bits/cycle
Memory controllers	4, located one at each corner
Memory latency	128 cycles

NI is assumed to take one cycle, as this step essentially reuses the original function in the NI which is modeled as one cycle in Garnet. Wormhole switching with credit-based flow control is assumed, although NoRD is agnostic to the switching and flow control mechanism used. Table 1 lists the key parameters used in the evaluations. Full system simulation uses a 16-node mesh, and synthetic traffic simulation uses both 16- and 64-node configurations to evaluate scalability.

We compare the following designs: (1) No_PG: baseline design with no power-gating; (2) Conv_PG: applying conventional power-gating to routers; (3) Conv_PG_OPT: conventional power-gating optimized with early wakeup (this optimized design not only improves performance by partially hiding wakeup latency, but also reduces power-gating overhead by avoiding powering-off all idle periods that are shorter than 4 cycles); (4) NoRD: our proposed approach based on node-router decoupling. In addition, all designs under evaluation are augmented with adaptive routing algorithms using Duato's Protocol [4]. The only difference is that (1)~(3) use adaptive routing in adaptive VCs and XY routing in escape VCs, whereas (4) uses adaptive routing and the ring-based escape mechanism described in Section 4.2.

5.2 Workloads

Multi-threaded PARSEC 2.0 benchmarks [2] are used for the majority of simulations, as the performance and power consumption of realistic workloads are of primary concern. Each core is warmed up for sufficiently long time (with a minimum of 10 million cycles) and then run until completion. We also perform simulations with synthetic traffic (uniform random and bit-complement [3]) to provide insight on the behavior of different designs across a wide range of load rates and parameter values. In those cases, packets are uniformly assigned two lengths. Short packets are single-flit while long packets have 5 flits. For synthetic traffic, the simulator is warmed up for 10,000 cycles and then the statistics are collected over another 100,000 cycles.

6. Results and Analysis

6.1 Wakeup Thresholds

To simulate NoRD, the appropriate wakeup thresholds must first be found. This is done empirically. All routers are forced into sleep mode without waking up − concentrating traffic on the Bypass Ring − and the number of VC requests (averaged over all routers) is recorded while varying the load rate. It can be seen from Figure 7 that the maximum achievable throughput of the Bypass
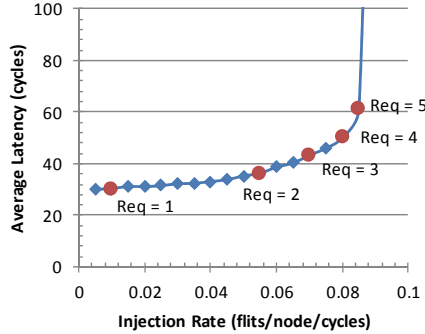
Figure 7: Determining wakeup threshold.

Figure 8: Static energy comparison (normalized to No_PG).

Ring is low (i.e., 14% of the throughput when all routers are turned on), indicating that some routers need to be awoken when network traffic increases, as measured by VC requests.

The objective of choosing the wakeup thresholds is to maximize the static power savings opportunity while not significantly increasing packet latency. In this sense, the dual-threshold technique in asymmetric wakeup thresholding provides more flexibility in achieving a good trade-off. In the current implementation of NoRD, the performance-centric routers are assigned a threshold of 1 as they are critical to performance and need to be awoken early. The remaining power-centric routers can use a higher threshold to enable more power-savings. Considering that a threshold value of 4 VC requests can lead to nearly 60% increase in packet latency, the power-centric routers are assigned a threshold of 3 to avoid large performance penalty. Although the thresholds here are determined empirically, they work very well across all benchmarks.

6.2 Impact on Static Energy

Figure 8 presents the results of static energy of different designs normalized to No_PG. It can be seen that, Conv_PG reduces the static energy slightly more than Conv_PG_OPT by 4.2% on average (51.2% vs. 47.0%). This is because Conv_PG does power-gating as long as the routers are empty whereas Conv_PG_OPT power-gates routers only if the idle periods are longer than 3 cycles as indicated by the early wakeup signal. As shown later, early wakeup pays off for Conv_PG_OPT in terms of performance. The lowest static power is achieved in the proposed NoRD approach for all benchmarks, with an average reduction of 62.9% compared with No_PG. When comparing relatively, NoRD provides savings relative to Conv_PG and Conv_PG_OPT of 23.9% and 29.9% on average, respectively. This improvement mainly comes from the increased opportunity in utilizing short idle periods and the reduced number of wakeups through decoupling bypass.

6.3 Reducing Power-gating Overhead

To provide more insight of the effectiveness of NoRD in reducing power-gating overhead, Figure 9(a) compares the energy overhead caused by router wakeup for conventional power-gating designs and the bypass design, normalized to Conv_PG (No_PG is not shown in the figure as it does not have any wakeups). As can be seen, the power-gating overhead in NoRD is considerably reduced by 80.7% and 74.0% compared with Conv_PG and Conv_PG_OPT, respectively. Figure 9(b) shows the reduction in the total number of wakeups in different designs normalized to
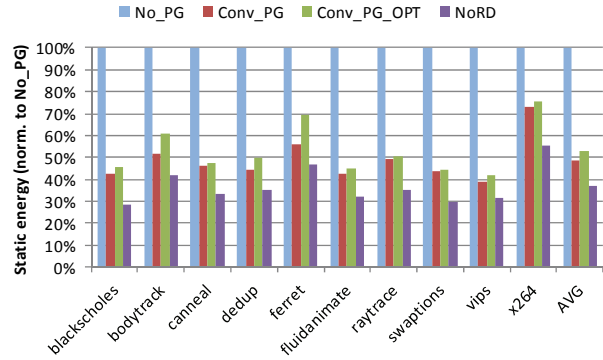
Conv_PG. NoRD decreases the number of wakeups by 81.0% and 73.3% over Conv_PG and Conv_PT_OPT, respectively, which explains the above substantial reduction of power-gating overhead and demonstrates the usefulness of the decoupling approach.

6.4 Impact on Dynamic Energy

Due to the detour of some packets in bypassing powered-off routers, the dynamic energy of NoRD may increase. Figure 10 plots the breakdown of NoC energy across the benchmarks, so that the relative impact of each NoC energy component can be examined. For the NoC dynamic energy (routers plus links), NoRD incurs an overhead of 10.2% on average, which constitutes 4.0% of the total NoC energy consumption. However, the static energy and wakeup overhead savings offered by NoRD constitutes 24.7% of the total NoC energy. Compared to No_PG, Conv_PG and Conv_PG_OPT, this renders NoRD a net savings of NoC energy of 9.1% and 9.4% and 20.6%, respectively. As on-chip networks consume a varying percentage of chip's overall energy (e.g., around 10%~36% as mentioned in Section 2), the impact of NoRD on overall chip energy depends on particular chip microarchitectures.

6.5 Impact on Performance

After presenting the energy statistics, we now compare the performance impact of different designs, which is another importance objective of power-gating techniques. Figure 11 shows the average packet latency, and Figure 12 compares the execution time of the four designs. No_PG does not have any performance penalty as there is no power-gating, and hence provides a lower bound on average packet latency and execution time. As can be seen, the aggressive power-gating scheme, Conv_PG, significantly degrades the average packet latency by 63.8% on average; whereas Conv_PG_OPT with early wakeup mitigates this degradation to 41.5% on average. These large penalties in conventional power-er-gating designs mainly come from the fact that once a router is power-gated off, any packet from either local traffic or in-network traffic suffers additional wakeup latency before being processed by the node. The comparison between Conv_PG_OPT and Conv_PG indicates that early wakeup does help a lot in reducing the performance penalty, but still cannot mask entirely the negative effects of wakeup latency. In contrast, NoRD decouples nodes from routers, effectively removing the wakeup latency from the critical path. The latency overhead in NoRD is caused by packet detours, which is partially offset by reduced per hop latency and avoidance of long wakeup latency as discussed before. As a result,
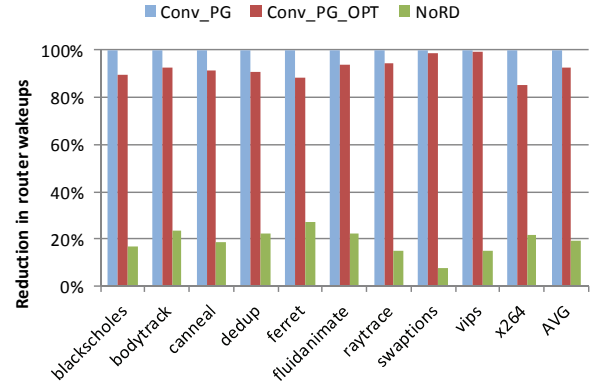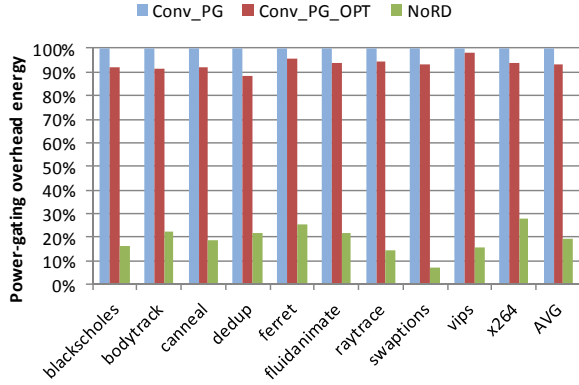
(a) Power-gating energy overhead   (b) Reduction in router of wakeups

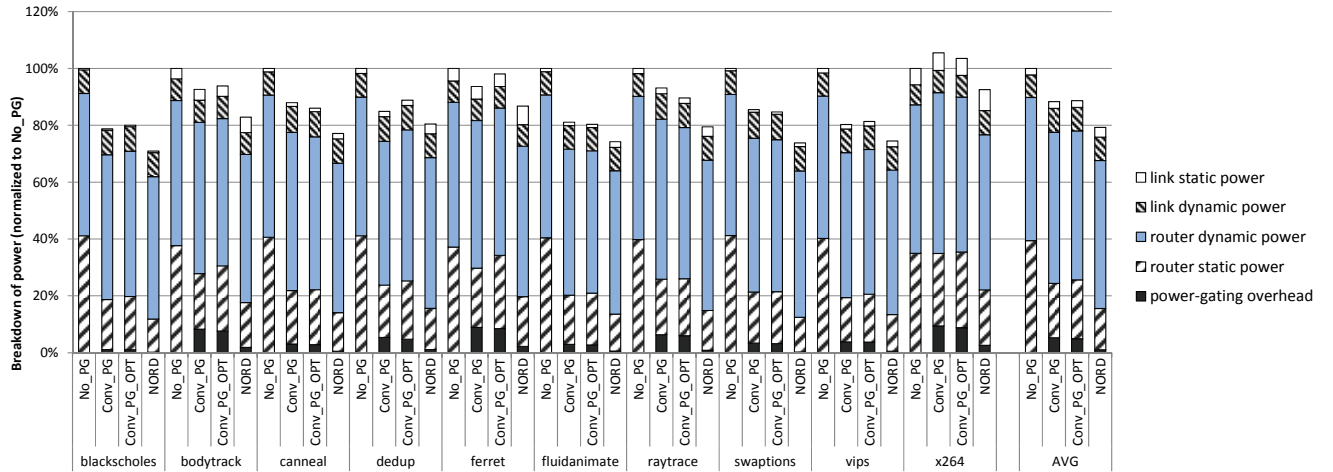**Figure 9: Reduction of power-gating overhead.**

Figure 10: Overall NoC energy breakdown.

the overall degradation of average packet latency in NoRD is only 15.2%, on average. The disparities in average packet latency among these designs result in different execution time, as shown in Figure 12. Although different benchmarks exhibit variations in the specific percentage of degradation due to their difference in network sensitivity, the trend is similar that NoRD has the smallest performance penalty compared to Conv_PG and Conv_PG_OPT. Overall, the Conv_PG, Conv_PG_OPT and NoRD increase the execution time by 11.7%, 8.1% and 3.9%, respectively, in order to achieve the energy saving described previously.

6.6 Effects on Hiding Wakeup Latency

So far, the effectiveness of NoRD has been demonstrated in real applications using full system simulations. In addition to the above primary results, we also perform simulations with synthetic uniform random traffic to highlight key characteristics of NoRD. Recall that cumulative wakeup latency is one of the big obstacles to power-gating routers, particularly in multi-hop networks. To illustrate that NoRD fundamentally solves this problem, Figure 13 shows the average packet latency of Conv_PG, Conv_PG_OPT and NoRD while varying the wakeup latency across a wide range. The load rate is set to the average load rate of PARSEC benchmarks. As can be seen, the latency of Conv_PG and

Conv_PG_OPT increases by nearly 1.5X and when the wakeup latency increases from 9 to 18 cycles; whereas the latency of NoRD remain similar for different wakeup latencies, which clearly demonstrates its ability to hide wakeup latency.

6.7 Behavior across Full Range of Network Loads

Next, we investigate the behavior of different designs across the entire network load range: from zero load to saturation loads. Figure 14 presents the performance and power results of a 16-node mesh under uniform random traffic, and Figure 15 presents for 64-node under uniform random and bit-complement traffic. Here, while the behavior of No_PG is very typical, interesting results are found for Conv_PG_OPT and NoRD. These are explained by separating the loads into three regions.

(1) Low to medium load region: When the load is very low, many routers are in the gated-off state for the majority of the time in both Conv_PG_OPT and NoRD. For Conv_PG_OPT, packets are likely to experience wakeup latency once or multiple times, so the average packet latency is high. For NoRD, packets use bypass more often, so average latency is increased due to detours. When load gradually increases, more routers are in the on-state, which tends to reduce the latency. This factor actually offsets the effect of increased load on average latency, leading to a net decrease in
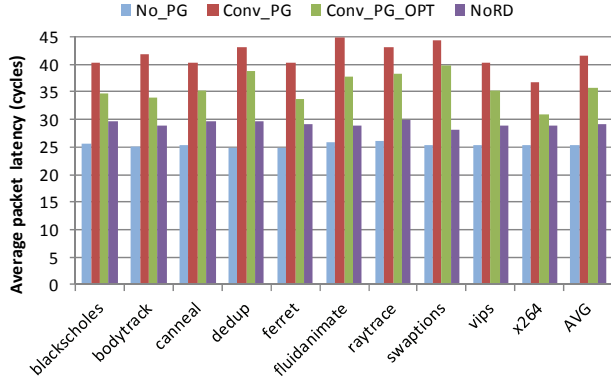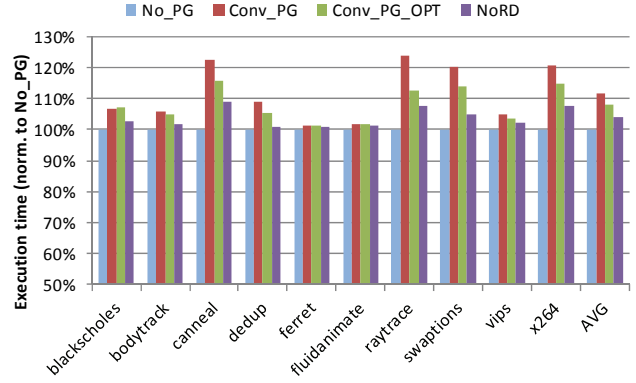
Figure 11: Average packet latency.

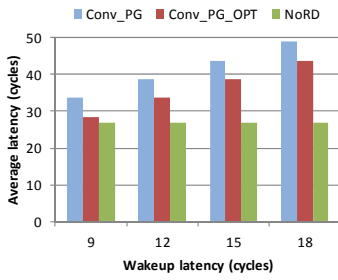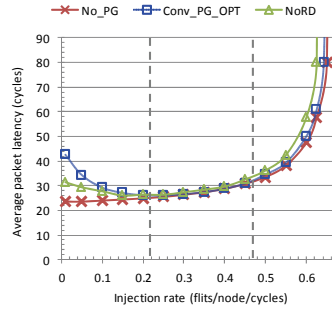**Figure 12: Execution time.**

Figure 13: Impact of wakeup latency.

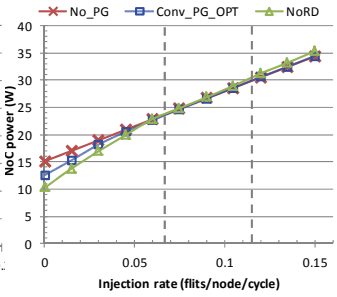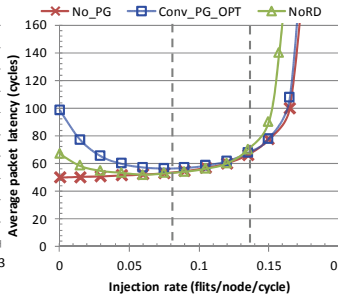**Figure 14: Packet latency and power of 16-node for different load ranges.**
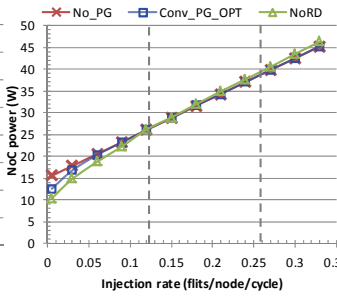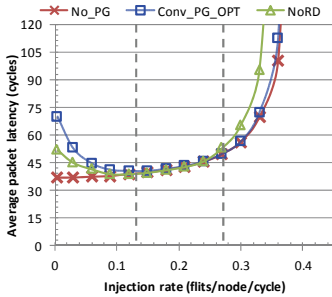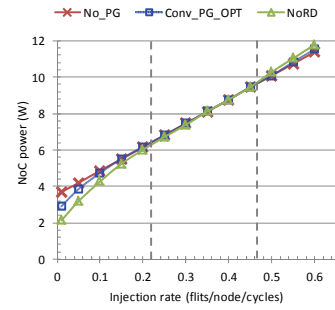

Figure 15: Packet latency and power for 64-node. Left two figures: uniform random; right two figures: bit-complement.

latency for Conv_PG_OPT and NoRD. As can be seen, NoRD achieves both lower average latency and lower power than Conv_PG_OPT. Note that, in this region, NoRD has increased benefits compared to Conv_PG_OPT for larger networks. This is because the cumulative wakeup latency problem in Conv_PG_OPT is more severe due to the increased NoC diameter in larger networks. A gated-off router at any hop of a packet's route adds extra wakeup latency, and every router has a high probability of being gated-off under low load. For instance, at 10% injection rate under uniform traffic, the latency for No_PG, Conv_PG_OPT and NoRD for a 4x4 mesh is 24, 34 and 29 cycles, respectively; whereas for an 8x8 mesh, it is 36, 52 and 44 cycles, respectively. This indicates that for a 64 node network, the latency of NoRD is lower than Conv_PG_OPT with an increased difference compared to the 16 node network. Curves for power for an 8x8 NoC are also similar in shape as for 4x4, indicating that the net energy-savings of NoRD that considers all energy contributors is still more favorable than conventional PG for larger networks.

(2) Medium to high load region: In this region, the three schemes have very similar latency and power characteristics. The relatively high load causes most of the routers to be turned on, making little difference between the designs with or without power-gating.

(3) Saturation region: In this region, as nearly all routers are in the on-state, both Conv_PG_OPT and NoRD are reduced to No_PG, except that they use different escape mechanisms. In this regard, as the escape ring in NoRD has less flexibility in routing packets as compared to escape XY routing of Conv_PG_OPT, NoRD saturates a little earlier. However, this is not an inherent limitation of node-router decoupling, as more efficient dead-lock-free routing algorithms such as in [5] can be used for the bypass ring to close the throughput difference.

Our full system simulations show that real application loads, in practice, typically stay within the low-to-medium region where NoRD has clear advantages over Conv_PG_OPT in both performance and power.

6.8 Discussion

Area Overhead

For any power-gating technique, there is hardware overhead for the sleep switch and the distribution of the sleep signal. While it greatly depends on the optimization level of circuit design, the area overhead of a well-designed power-gating block is usually between 4~10% [10, 12]. More of a concern for NoRD is the area overhead of the added bypass and related hardware. In evaluating this, the Orion 2.0 [13] on-chip network model is used with 45nm technology parameters. We modified the simulator to model all the additional key components of NoRD, including the added forwarding logic in the NI. Results show that NoRD has an area overhead of only 3.1% compared with Conv_PG_OPT.

Other Conventional Power-Gating Techniques

We have compared NoRD with conventional power-gating of routers optimized with early wakeup, which is one of the most effective optimizations so far. Another trade-off that can be done for conventional power-gating is to power-gate smaller individual components within a router, as mentioned in Section 3.2. As investigated in [25], this approach can reduce static energy by an additional 17.6% on top of conventional power-gating with early wakeup, but at the cost of 15.9% area overhead using a commercial standard cell library. In comparison, the proposed NoRD can reduce static energy by 29.9% with only 3.1% area overhead compared to Conv_PG_OPT, indicating NoRD is a much more cost-effective approach.

Bufferless Routing

Recently, bufferless routing has been proposed as a means of reducing router power consumption [6]. Although the bufferless approach may introduce livelock, deflection and packet reassembly issues, it can eliminate buffers and their associated power consumption. However, as shown in Figure 1(b), while buffers are the largest contributor of static power, other router components consume a considerable percentage (e.g., 45%) of total static power, which would remain even if a bufferless approach is used. In fact, bufferless routing is complementary to power-gating techniques in general, as both can be applied at the same time to reduce router power consumption. For example, flits in bufferless routing have the option to be deflected through the bypass paths in NoRD if needed.

Shorter router pipelines and aggressive NoRD design

In the baseline, a canonical router is used which takes 4 cycles for the pipeline plus 1 cycle for LT; whereas the bypass for gated-off routers in NoRD takes 2 cycles plus 1 cycle for LT. There are some techniques such as look-ahead routing [15] and speculative SA [26] that can potentially shorten the 4-cycle router pipeline to 2-cycle. However, NoRD is still competitive in that case for the following reasons. First, shortening the pipeline by two also reduces the number of cycles that can hide wakeup latency by two, making the total time (pipeline delay plus wakeup latency) to go through a gated-off router to remain the same. Second, these techniques come with overheads. Look-ahead routing requires contention information to be propagated one-hop ahead, while speculative SA may not always succeed, making 2 cycles a best-case scenario. Ironically, speculative SA is likely to succeed at low load, in which routers are also likely to be gated-off and the wakeup latency dominates the delay at those routers. Third, the bypass in NoRD can also be optimized to become more aggressive by directly connecting the Bypass Inport to the Bypass Outport. This has a similar rationale as for speculation in that the forwarding of flits optimistically assumes that there is no local flit to inject,

thereby bypassing the router in just one cycle. In case of conflict, additional cycles are needed, just like that in speculative SA. Therefore, when optimizations are used for both the baseline and NoRD, there are no clear advantages for the baseline, and NoRD remains competitive.

7. Related Work

Power-gating as a circuit-level technique has been proposed for some time and has been applied to cores and execution units in CMPs [10, 19, 20]. Only recently has it been investigated for on-chip network routers [23, 24, 25]. These works apply power-er-gating to routers, but are severely limited by the BET requirement, wakeup delay and disconnection problem. In contrast, as our approach breaks node-router dependence, it provides a unified solution to these problems and enables effective use of power-gating to on-chip routers.

Bypass has been used for various purposes in on-chip networks. In [17], default backup paths are proposed to allow fault-tolerance with graceful performance degradation. This scheme assumes all routers are notified each time a router becomes faulty and requires re-computing the routing table for all routers for each fault occurrence. Therefore, it is not suitable for run-time power-gating in which the status of routers may change more frequently. In comparison, each router in the proposed NoRD approach can be powered-on/off independently without notifying all other routers or re-computing any routing tables. A modular router architecture is proposed in [16] that can bypass some internal faults within a router. However, this design does not provide chip-wide connectivity and does not explore the application of power-gating techniques as proposed in this paper. Express VC [18] also makes use of bypass in that it virtually bypasses routers to improve both performance and dynamic power. However, it does not reduce router static power. Another bypass design is proposed in [11] for adaptive flow control between bufferless and buffered router modes. It is based on bufferless design and is subject to the associated constraints, such as flit-by-flit routing, livelock and packet reassembly issues. Moreover, it only targets the buffers in a router and applies power-gating techniques conventionally, whereas our approach is able to bypass the entire router and implement node-router decoupling.

Many prior works have investigated techniques to save dynamic and static power of links [14, 27, 29]. These techniques can readily be used together with NoRD to provide more energy-efficient NoC designs. These works and other general-purpose dynamic power-saving techniques (such as clock-gating) have different targets other than router static power and, therefore, are orthogonal and complementary to this work.

8. Conclusion

While power-gating is a promising technique to reduce static power, node-router dependence severely limits its effective use in on-chip routers due to the BET limitation, wakeup delay and disconnection problem. In this paper, a novel approach that provides separate power-gating bypass to decouple the node's ability for sending, receiving and forwarding packets from the on/off status of the associated router is proposed. The resulting design can significantly reduce the number of state transitions, increase the length of idle periods, completely hide the wakeup latency from the critical path and eliminate node-network disconnection problems. Full system simulations show that,

compared to an optimized conventional power-gating technique applied to on-chip routers, NoRD can further reduce the router static energy by 29.9% and improve average packet latency by 26.3%, with only 3% additional area overhead.

Acknowledgements

References

[1] N. Agarwal, T. Krishna, L.-S. Peh, and N. K. Jha, "GARNET: A detailed on-chip network model inside a full-system simulator," in *International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pp. 33-42, 2009.

[2] C. Bienia and K. Li, "Parsec 2.0: A new benchmark suite for chip-multiprocessors," in *Proceedings of the 5th Annual Workshop on Modeling, Benchmarking and Simulation*, 2009.

[3] W. Dally and B. Towles, *Principles and Practices of Interconnection Networks*: Morgan Kaufmann Publishers Inc., 2003.

[4] J. Duato, "A new theory of deadlock-free adaptive routing in wormhole networks," *IEEE Transactions on Parallel and Distributed Systems (TPDS),* vol. 4, pp. 1320-31, 1993.

[5] J. Duato and T. M. Pinkston, "A general theory for deadlock-free adaptive routing using a mixed set of resources," *IEEE Transactions on Parallel and Distributed Systems (TPDS),* vol. 12, pp. 1219-1235, 2001.

[6] C. Fallin, C. Craik, and O. Mutlu, "CHIPPER: A low-complexity bufferless deflection router," in *17th International Symposium on High Performance Computer Architecture (HPCA)*, pp. 144-55, 2011.

[7] R. W. Floyd, "Algorithm 97: shortest path," *Communications of the ACM,* vol. 5, p. 345, 1962.

[8] Y. Hoskote, S. Vangal, A. Singh, N. Borkar, and S. Borkar, "A 5-GHz mesh interconnect for a Teraflops processor," *IEEE Micro,* vol. 27, pp. 51-61, 2007.

[9] J. Howard, S. Dighe, S. R. Vangal, G. Ruhl, N. Borkar, S. Jain, V. Erraguntla, M. Konow, *et al.*, "A 48-Core IA-32 Processor in 45 nm CMOS Using On-Die Message-Passing and DVFS for Performance and Power Scaling," *IEEE Journal of Solid-State Circuits,* vol. 46, pp. 173-83, 2011.

[10] Z. Hu, A. Buyuktosunoglu, V. Srinivasan, V. Zyuban, H. Jacobson, and P. Bose, "Microarchitectural techniques for power gating of execution units," in *International Symposium on Lower Power Electronics and Design (ISLPED)*, pp. 32-37, 2004.

[11] S. A. R. Jafri, Y.-J. Hong, M. Thottethodi, and T. N. Vijaykumar, "Adaptive flow control for robust performance and energy," in *43rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 433-444, 2010.

[12] H. Jiang, M. Marek-Sadowska, and S. R. Nassif, "Benefits and costs of power-gating technique," in *International Conference on Computer Design (ICCD)*, pp. 559-566, 2005.

[13] A. Kahng, L. Bin, L.-S. Peh, and K. Samadi, "ORION 2.0: A fast and accurate NoC power and area model for early-stage design space exploration," in *Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pp. 423-428, 2009.

[14] E. J. Kim, K. H. Yum, G. M. Link, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, M. Yousif, and C. R. Das, "Energy Optimization Techniques in Cluster Interconnects," in *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 459-464, 2003.

[15] J. Kim, D. Park, T. Theocharides, N. Vijaykrishnan, and C. R. Das, "A low latency router supporting adaptivity for on-chip interconnects," in *42nd Design Automation Conference (DAC)*, pp. 559-564, 2005.

[16] J. Kim, C. Nicopoulos, D. Park, V. Narayanan, M. S. Yousif, and C. R. Das, "A gracefully degrading and energy-efficient modular router architecture for on-chip networks," in *33rd International Symposium on Computer Architecture (ISCA)*, pp. 4-15, 2006.

[17] M. Koibuchi, H. Matsutani, H. Amano, and T. M. Pinkston, "A lightweight fault-tolerant mechanism for network-on-chip," in *2nd ACM/IEEE International Symposium on Networks-on-Chip (NOCS)*, pp. 13-22, 2008.

[18] A. Kumar, L.-S. Peh, P. Kundu, and N. K. Jha, "Express virtual channels: Towards the ideal interconnection fabric," in *34th Annual International Symposium on Computer Architecture (ISCA)*, pp. 150-161, 2007.

[19] A. Lungu, P. Bose, A. Buyuktosunoglu, and D. J. Sorin, "Dynamic power gating with quality guarantees," in *International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 377-382, 2009.

[20] N. Madan, A. Buyuktosunoglu, P. Bose, and M. Annavaram, "A case for guarded power gating for multi-core processors," in *17th International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 291-300, 2011.

[21] P. S. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Hallberg, J. Hogberg, F. Larsson, A. Moestedt, *et al.*, "Simics: A full system simulation platform," *IEEE Computer,* vol. 35, pp. 50-58, 2002.

[22] M. K. Martin, D. J. Sorin, B. M. Beckmann, M. R. Marty, M. Xu, A. R. Alameldeen, K. E. Moore, M. D. Hill, *et al.*, "Multifacet's general execution-driven multiprocessor simulator toolset," *ACM SIGARCH Computer Architecture News,* vol. 33, pp. 92-99, 2005.

[23] H. Matsutani, M. Koibuchi, W. Daihan, and H. Amano, "Run-time power gating of on-chip routers using look-ahead routing," in *13th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 55-60, 2008.

[24] H. Matsutani, M. Koibuchi, D. Wang, and H. Amano, "Adding slow-silent virtual channels for low-power on-chip networks," in *2nd ACM/IEEE International Symposium on Networks-on-Chip (NOCS)*, pp. 23-32, 2008.

[25] H. Matsutani, M. Koibuchi, D. Ikebuchi, K. Usami, H. Nakamura, and H. Amano, "Ultra fine-grained run-time power gating of on-chip routers for CMPs," in *4th ACM/IEEE International Symposium on Networks on Chip (NOCS)*, pp. 61-68, 2010.

[26] L. S. Peh and W. J. Dally, "A delay model and speculative architecture for pipelined routers," in *7th International Symposium on High Performance Computer Architecture (HPCA)*, pp. 255-66, 2001.

[27] V. Soteriou and P. Li-Shiuan, "Design-space exploration of power-aware on/off interconnection networks," in *2nd International Conference on Computer Design (ICCD)*, pp. 510-17, 2004.

[28] M. B. Taylor, J. Kim, J. Miller, D. Wentzlaff, F. Ghodrat, B. Greenwald, H. Hoffman, P. Johnson, *et al.*, "The Raw microprocessor: a computational fabric for software circuits and general-purpose programs," *IEEE Micro,* vol. 22, pp. 25-35, 2002.

[29] B. Zafar, J. Draper, and T. M. Pinkston, "Cubic Ring networks: A polymorphic topology for network-on-chip," in *39th International Conference on Parallel Processing (ICPP)*, pp. 443-452, 2010.