Enabling Technology for On-Chip Networks

Prof. William J. Dally, Daniel U. Becker, James Chen, Nan Jiang, George Michelogiannakis
Concurrent VLSI Architecture Group, Stanford University
<http://nocs.stanford.edu/>

Overview – Future large-scale chip multiprocessors will rely on sophisticated on-chip interconnection networks to provide efficient core-to-core communication. These networks must provide high throughput and low latency while meeting stringent power and area constraints at the same time.

Our research is focusing on the development of enabling technology for such on-chip networks; in particular, we are investigating aspects ranging from the circuit level, where we design efficient implementations of basic building blocks including channels, buffers and crossbar switches, to high-level aspects like network topologies, flow-control techniques, routing algorithms and fairness considerations.

Age-based Priority – We are investigating techniques for improving fairness in the allocation of network resources. Our results indicate that allocation using age-based priorities has many beneficial effects on network performance. Age-based priority can be used to resolve network starvation effects that occur under adversarial traffic patterns (Figure 1). Furthermore, it provides a synergistic effect with some adaptive routing algorithms that leads to better load balancing across the network (Figure 2). Under high load, age-based priority also reduces the variance in network latency across packets, which allows for tighter bounds on maximum packet delay at the cost of slightly increased minimum delay (Figure 3).

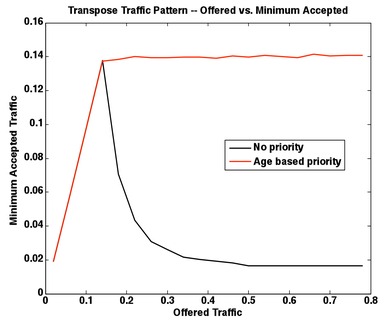

Figure 1. Resolving network starvation using age-based priorities.

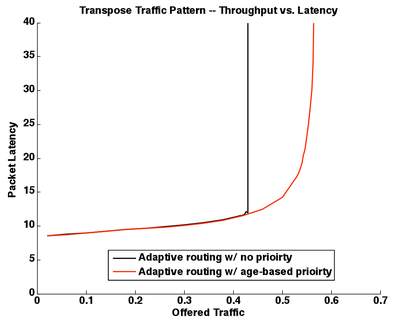

Figure 2. Improved saturation throughput with age-based priorities.

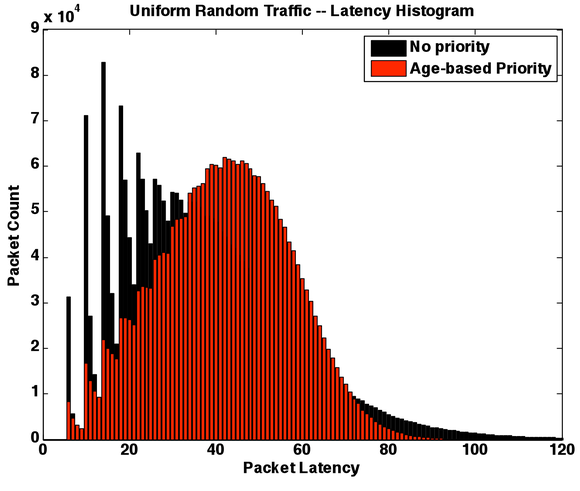

Figure 3. Shift in packet latency distribution under heavy network load with age-based priorities.

Router Microarchitecture – We have developed a highly parameterized open source RTL implementation of a virtual channel router, which we use to evaluate microarchitectural design tradeoffs; for example, Figure 4 shows the impact of different allocator implementations on network performance and router cost. Based on these evaluations, we have developed several improvements to VC and switch allocation and a simplified router pipeline. Furthermore, we are investigating schemes for reducing router cost by sharing part of the input buffers among multiple VCs (Figure 5).

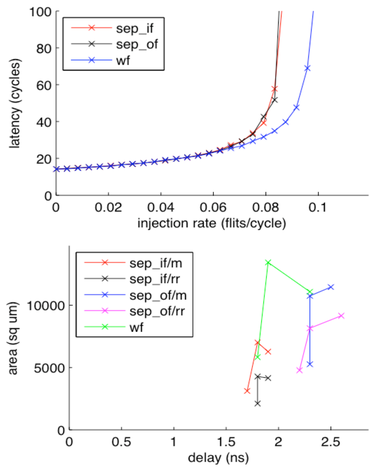

Figure 4. Network-level performance for mesh and FBFly.

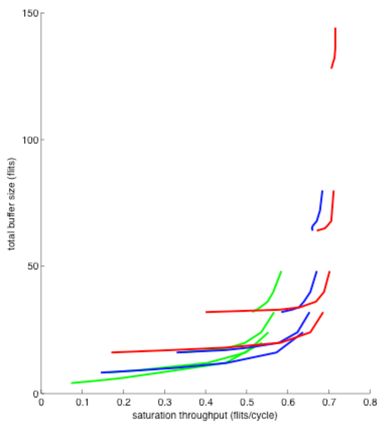

Figure 5. Buffer partitioning tradeoffs for 10x10 FBFly router.

Network Topologies – Using our parameterized router implementation, we have conducted a detailed study on the energy efficiency of common network topologies for a 64-node network. Our results are based on post place and route simulations in a commercial 45nm process. For a variety of traffic patterns the *flattened butterfly* (*FBFly)* topology provides the best energy efficiency (Figure 6).

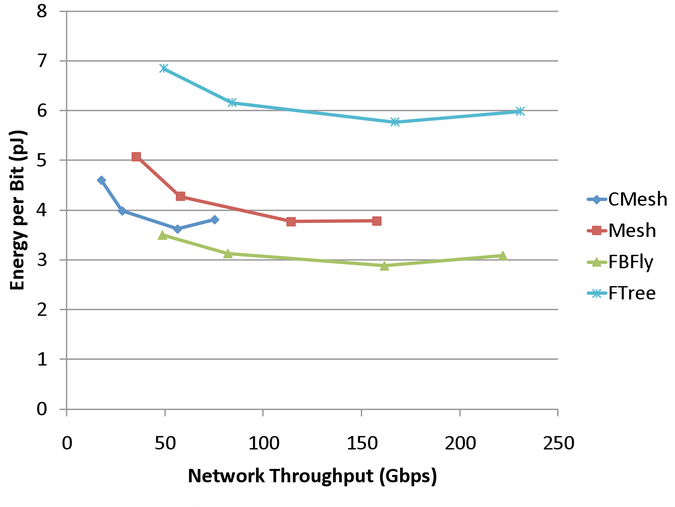

Figure 6. Energy per bit transported for different network topologies and channel sizes.

Channel Circuits – We are investigating circuit techniques that improve the efficiency of channel circuits, including *low-swing signaling*, *high-bandwidth transmission* and *elastic buffering*. We have developed low-swing repeaters that reduce transmission energy by up to 6x in the channels and 5x in the switches compared to conventional signaling (Figures 7-8).

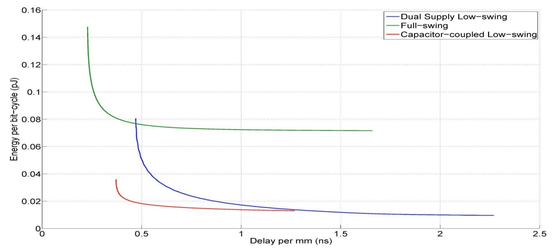

Figure 7. Channel energy vs. delay tradeoff.

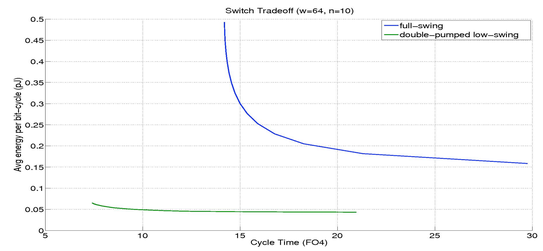

Figure 8. Switch energy vs. delay tradeoff.

Flow Control – *Elastic buffer* flow control reuses existing pipeline flip-flops in the channels to form distributed FIFOs (Figure 9), eliminating the need for input buffers at the routers and credit flow control. Compared to VC networks, EB networks can provide up to 45% shorter cycle time, 43% more throughput per unit power and 22% more throughput per unit area (Figures 10-11).

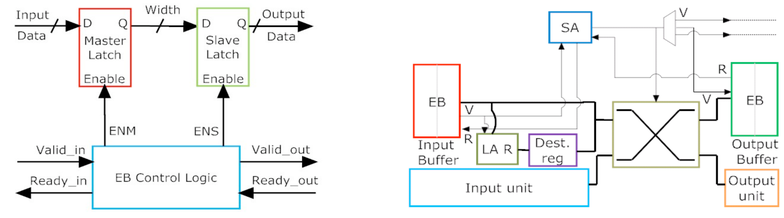

Figure 9. Elastic buffer and single-stage EB router block diagram.

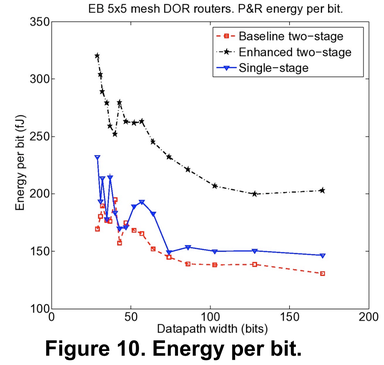

Figure 10. Energy per bit.

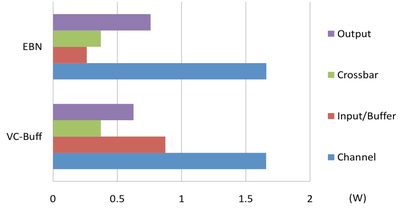

Figure 11. Power breakdown for low-swing mesh at 2% injection rate.