

V. RELATED WORK

Peh *et al.* proposed a speculative technique for VC routers allowing a packet to simultaneously arbitrate for both VCA and SA giving a higher priority for non-speculative packets to win SA; therefore reducing zero-load latency in which the probability of failed speculation is small [13]. This low latency, however, comes with the high complexity of SA circuitry and also wastes more power each time the speculation fails. A packet must stall if even it wins SA but fails VCA, and then has to redo both arbitration at next cycle. Reversely, RoShaQ is non-speculative architecture. An incoming packet in RoShaQ only stalls if it fails both OPA and SQA; therefore it has high chance to advance either to be written to a shared queue (if it wins SQA) or be sent to output port (if it wins OPA) instead of stalling at an input port, and also reducing re-arbitration times.

Increasing crossbar input ports, that allows to directly connect to all virtual-channels of an input port instead of muxing them, improves much network throughput for VC routers. Using a large-radix crossbar is feasible and low-cost than adding more buffers as the results reported by DeMicheli *et al.* [19]. Recently, Passas *et al.* designed a 128×128 crossbar allowing to connect 128 tiles while occupying only 6% of their area [20]. This fact encourages us to build RoShaQ that has two crossbars while sharing cost-expensive buffer queues. The additional costs of crossbars are compensated by the simplicity of allocators and reducing the number of routing computation circuits that make our router better VC routers in many-fold: throughput, latency and packet energy.

IBM Colony router has a shared central buffer which is built from a time-multiplexed multi-bank SRAM array with wide word-width in order that it can be simultaneously written/read multiple flits (defined as a chunk) by input/output ports [21]. As a result, the central buffer is high cost and not identical with input queue design. RoShaQ has all buffer queues (both input and share queues) to be the same structure that allows to reuse the existing generic simple queues reducing practical design and test costs.

Latif *et al.* implemented a router with input ports sharing all queues [10] that is similar to the architecture illustrated in Fig. 6(a). Its implementation on FPGA shows more power and area-efficient than typical input VC routers. However, no router performance was reported and compared to VC routers. A similar approach is proposed by Tran *et al.* [15]; however, due to the high complexity of its allocators and also inter-router round-trip request/grant signaling, its performance is actually poorer than a typical router.

Ramanujam *et al.* recently proposed a router architecture with shared-queues named DSB which emulates an output-buffered router [11]. This router is similar to one illustrated in Fig. 6(b) that has higher zero-load latency than a VC router. This is because a packet has to travel through both two crossbars and be buffered in both input and shared queues at each router. Besides that, the timestamp-based flow control of DSB router design is highly complicated and hence consumes much larger area and power than a typical VC router (that are 35% and 58%, respectively). RoShaQ allows input packets to bypass shared-queues hence achieves lower zero-load latency compared to VC routers. RoShaQ also achieves much higher saturation throughput than VC routers, with only small area and power overheads while consuming less average energy per packet.

VI. CONCLUSION

We have presented RoShaQ, a new router architecture which allows to share multiple buffer queues for improving network throughput. Input packets also can bypass shared queues to achieve low latency in the case that the network load is low. Compared to a typical VC

router, it is 21% lower zero-load latency and 14% higher saturation throughput with only 4% higher power and 16% larger area. It is also 2% higher throughput than a full-crossbar VC router with 3% less than power and area. While targeting the same average packet latency of 60 cycles, RoShaQ has 7% and 5% less energy dissipated per packet than typical VC and full-crossbar VC routers, respectively, while having the same buffer space. Its low latency, high throughput and low energy are achieved without the need of pipeline speculation.

ACKNOWLEDGMENTS

This work was supported by a VEF Fellowship, SRC GRC Grant 1971 and CSR Grant 1659, ST Microelectronics, UC Micro, NSF Grants 0903549, 1018972 and CAREER Award 0546907, and Intel.

REFERENCES

- [1] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*, Morgan Kaufmann, San Francisco, USA, 2007.
- [2] S. Borkar, "Thousand core chips: a technology perspective," in *Design Automation Conference (DAC)*, June 2007, pp. 746–749.
- [3] C. H. V. Berkel, "Multi-core for mobile phones.," in *Design, Automation and Test in Europe (DATE)*, 2009, pp. 1260–1265.
- [4] W. J. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," in *Design Automation Conference (DAC)*, 2001, pp. 684–689.
- [5] W. J. Dally, "Virtual-channel flow control," *IEEE TPDS*, vol. 3, pp. 194–205, Mar. 1992.
- [6] Y. Hoskote *et al.*, "A 5-GHz mesh interconnect for a teraflops processor," *IEEE Micro*, vol. 27, no. 5, pp. 51–61, Sept. 2007.
- [7] T. Moscibroda and O. Mutlu, "A case for bufferless routing in on-chip networks," in *Int. Symp. on Computer Architecture (ISCA)*, 2009, pp. 196–207.
- [8] M. Hayenga *et al.*, "SCARAB: A single cycle adaptive routing and bufferless network," in *IEEE/ACM Int. Symp. on Microarchitecture (MICRO)*, 2009, pp. 244–254.
- [9] G. Michelogiannakis *et al.*, "Evaluating bufferless flow control for on-chip networks," in *ACM/IEEE Int. Symp. on Networks-on-Chip (NOCS)*, 2010, pp. 9–16.
- [10] K. Latif *et al.*, "Power and area efficient design of network-on-chip router through utilization of idle buffers," in *IEEE Int Conf. and Workshops on Engineering of Computer Based Systems (ECBS)*, 2010, pp. 131–138.
- [11] R. S. Ramanujam *et al.*, "Design of a high-throughput distributed shared-buffer NoC router," in *ACM/IEEE Intl. Symp. on Networks-on-Chip (NOCS)*, 2010, pp. 69–78.
- [12] M. Galles, "Spider: a high-speed network interconnect," *IEEE Micro*, vol. 17, no. 1, pp. 34–39, 1997.
- [13] L. Peh and W. J. Dally, "A delay model and speculative architecture for pipelined routers," in *Intl. Symp. on High-Performance Computer Architecture (HPCA)*, Jan. 2001, pp. 255–266.
- [14] R. Mullins, A. West, and S. Moore, "Low-latency virtual-channel routers for on-chip networks," in *Intl. Symposium on Computer Architecture (ISCA)*, Mar. 2004, p. 188.
- [15] A. T. Tran and B. M. Baas, "DLABS: A dual-lane buffer-sharing router architecture for networks on chip," in *IEEE Workshop on Signal Processing Systems (SiPS)*, Oct. 2010, pp. 327–332.
- [16] W. J. Dally and B. Towles, *Principles and Practices of Interconnection Networks*, Morgan Kaufmann, San Francisco, USA, 2004.
- [17] N. Enright-Jerger and L. Peh, *On-Chip Networks*, Morgan-Claypool, 2009.
- [18] A. Banerjee *et al.*, "An energy and performance exploration of network-on-chip architectures," *IEEE TVLSI*, vol. 17, no. 3, pp. 319–329, 2009.
- [19] G. De Micheli *et al.*, "Networks on chips: from research to products," in *ACM/IEEE Design Automation Conference (DAC)*, 2010, pp. 300–305.
- [20] G. Passas *et al.*, "A $128 \times 128 \times 24$ gb/s crossbar interconnecting 128 tiles in a single hop and occupying 6% of their area," in *ACM/IEEE Int. Symp. on Networks-on-Chip (NOCS)*, 2010, pp. 87–95.
- [21] C. B. Stunkel *et al.*, "A new switch chip for IBM RS/6000 SP systems," in *ACM/IEEE Conference on Supercomputing (CS)*, 1999.