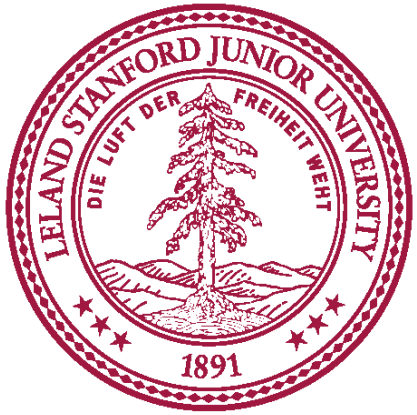

Evaluating Bufferless Flow Control for On-Chip Networks

George Michelogiannakis, Daniel Sanchez,

William J. Dally, Christos Kozyrakis

Stanford University

In a nutshell

- Many researchers report high buffer costs.
- Motivates bufferless networks.
- We compare bufferless networks with VC networks.
- We perform simple optimizations on both sides and a thorough analysis.
- We show that bufferless networks:
  - Consume only marginally less energy than buffered networks at very low loads.
  - Have higher latency and provide less throughput per unit power.
  - Are more complex.

Outline

- Methodology.
  - Evaluation infrastructure.
- Background.
- Optimizing routing in BLESS.
- Router microarchitecture.
- Network evaluation.
- Discussion.
- Conclusion.

Methodology

- Cycle-accurate network simulator.

- Balfour and Dally [ICS '06] power and area models.
  - Based on first-order principles.
  - We validate our models against HSPICE.
- 32nm ITRS high performance models, as a worst case for leakage power.
  - Also, a 45nm low-power commercial library.
- 2D 8x8 mesh.

Outline

- Methodology.
- Background.
  - A quick overview.
- Optimizing routing in BLESS.
- Router microarchitecture.
- Network evaluation.
- Discussion.
- Conclusion.

Bufferless flow control

- Flits can't wait in routers.
- Contention is handled by:
  - Dropping and retransmitting from the source.
  - Deflecting to a free output.

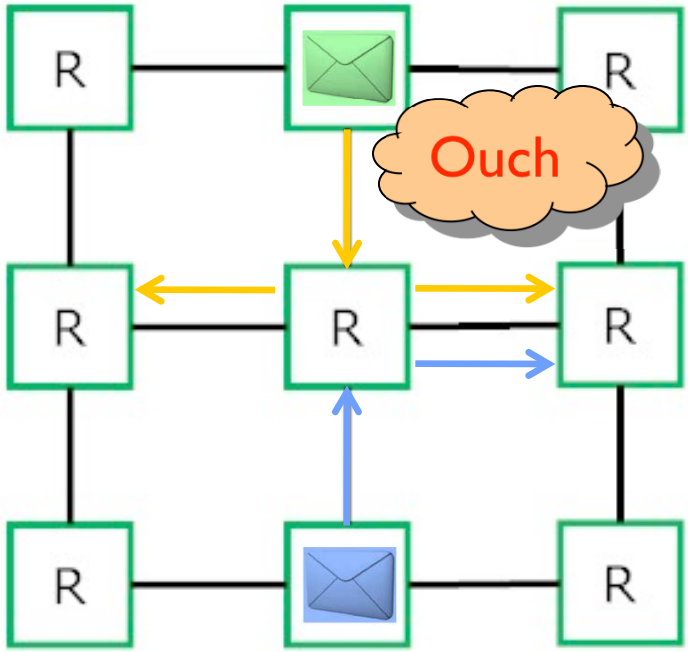

BLESS deflection network [ISCA '09]

- Flits bid for a single output using dimension-ordered routing (DOR).
- Body flits may get deflected.
  - They must contain destination information.
  - They may arrive out of order.
- Oldest flits are prioritized to avoid livelocks.
- We compare virtual channel (VC) networks against BLESS.

Outline

- Methodology.
- Background.
- Optimizing routing in BLESS.
  - Dimension-order revisited.
- Router microarchitecture.
  - Implications in router design.
- Network evaluation.
- Discussion.
- Conclusion.

Optimizing routing in BLESS

- Deadlocks impossible in bufferless networks, thus DOR unnecessary.
- Multidimensional routing (MDR) requests all productive outputs.
- 5% lower latency, equal throughput compared to DOR.

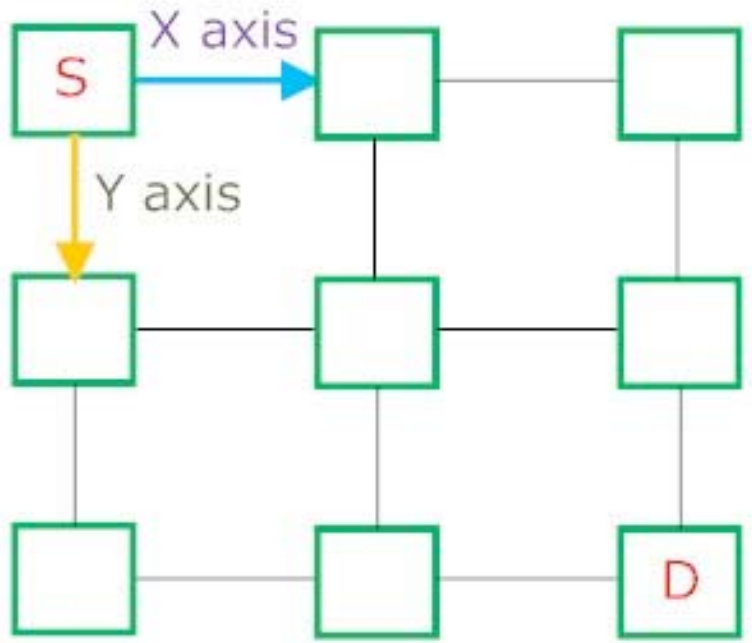

Allocator complexity

- Deflection networks require a complete matching.
  - Critical path through each output arbiter.

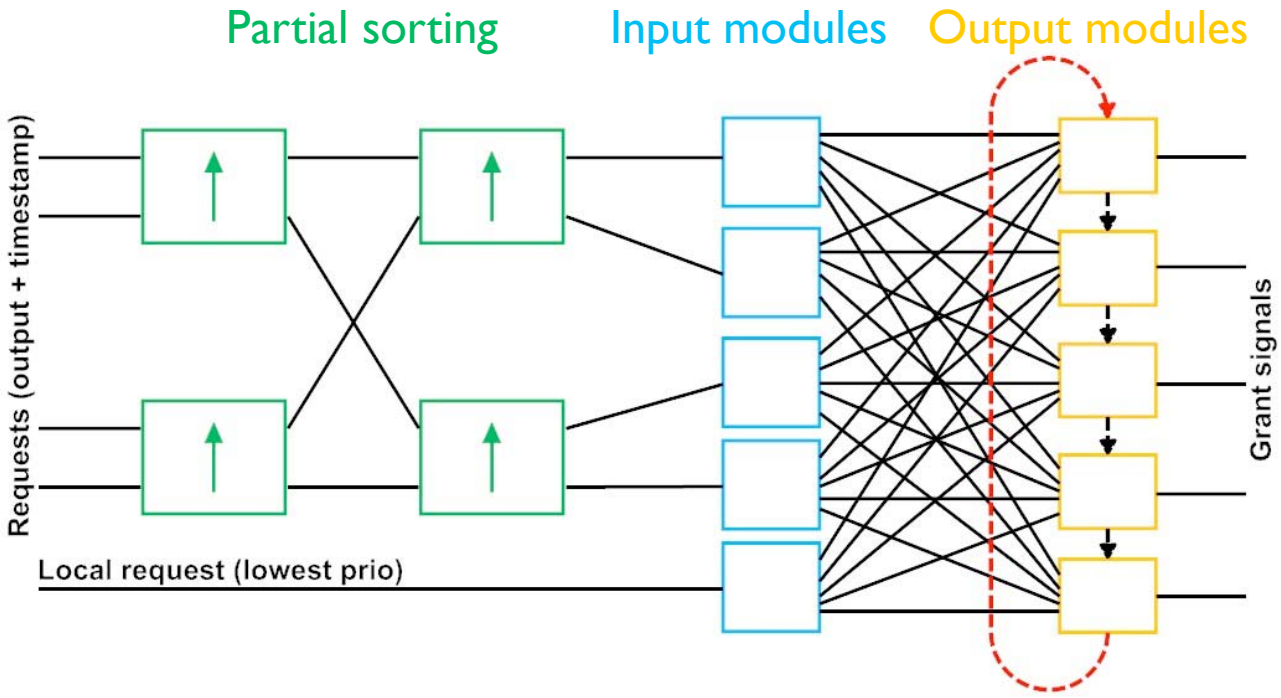

- BLESS allocator increases cycle time by 81% compared to input-first, round-robin switch allocator.

Buffer cost

- We assume efficient custom SRAMs.
- We use empty buffer bypassing.

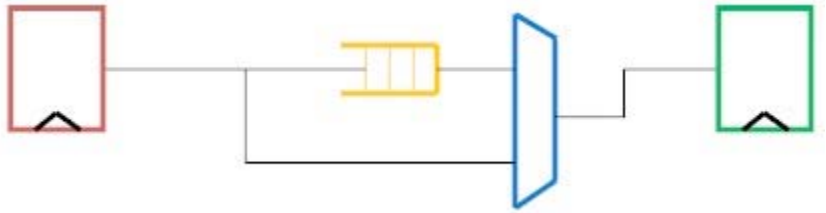

- Thus, at very low loads the extra power is only buffer leakage.
  - 1.5% of the overall network power.

Outline

- Methodology.
- Background.
- Optimizing routing in BLESS.
- Router microarchitecture.
- Network evaluation.
  - Let's talk numbers.
- Discussion.
- Conclusion.

Power versus injection rate

- BLESS: less power for flit injection rates lower than 7%.
- Higher than that, activity factor from deflections costs more.

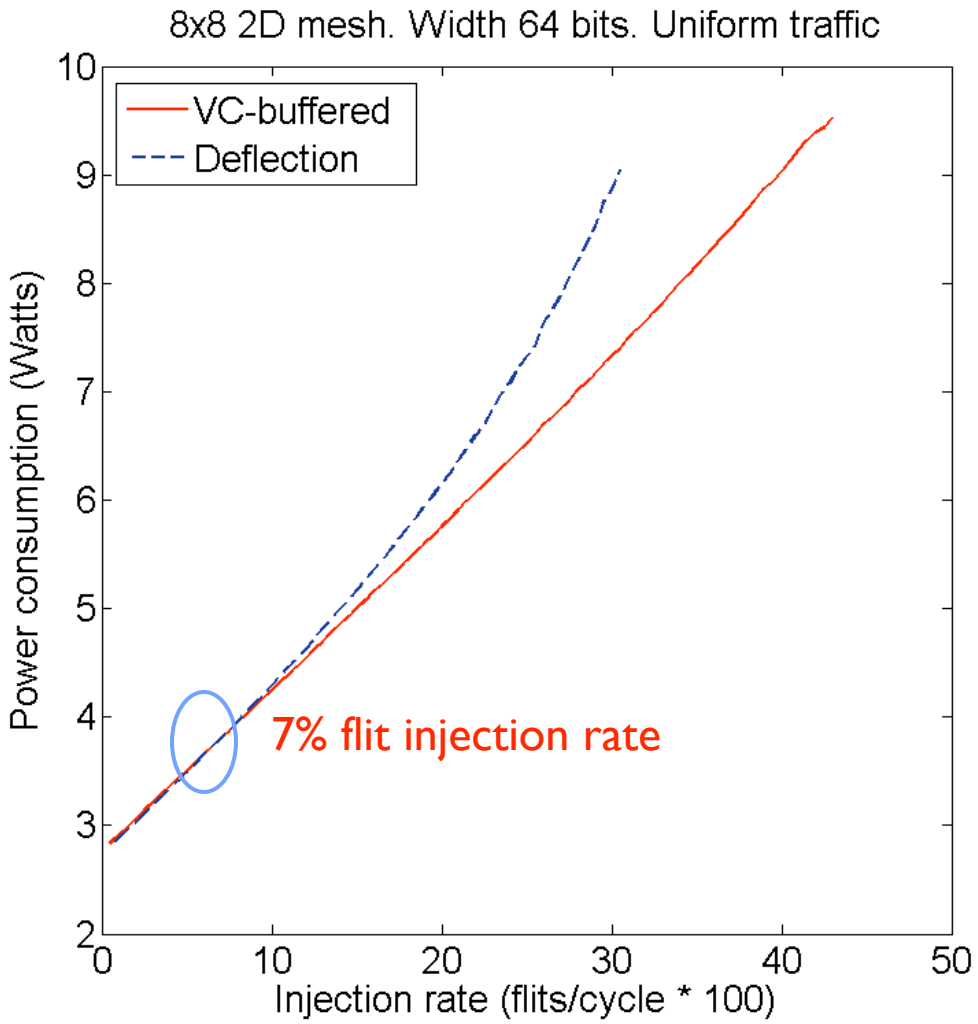

Throughput efficiency

Swept datapath width.

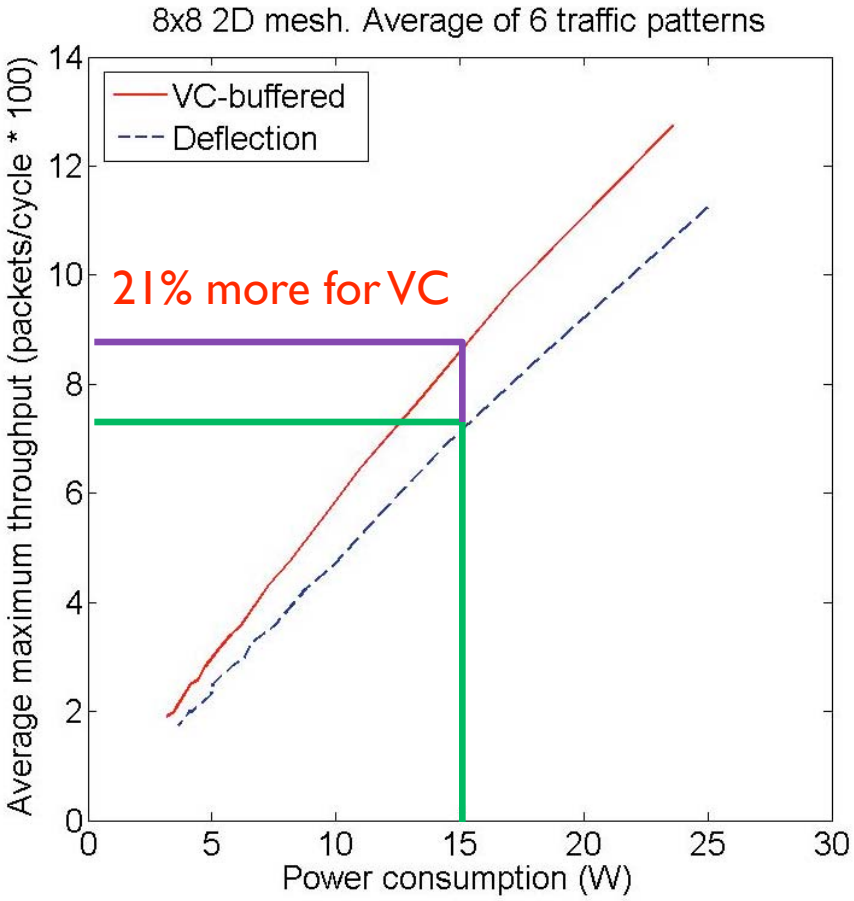

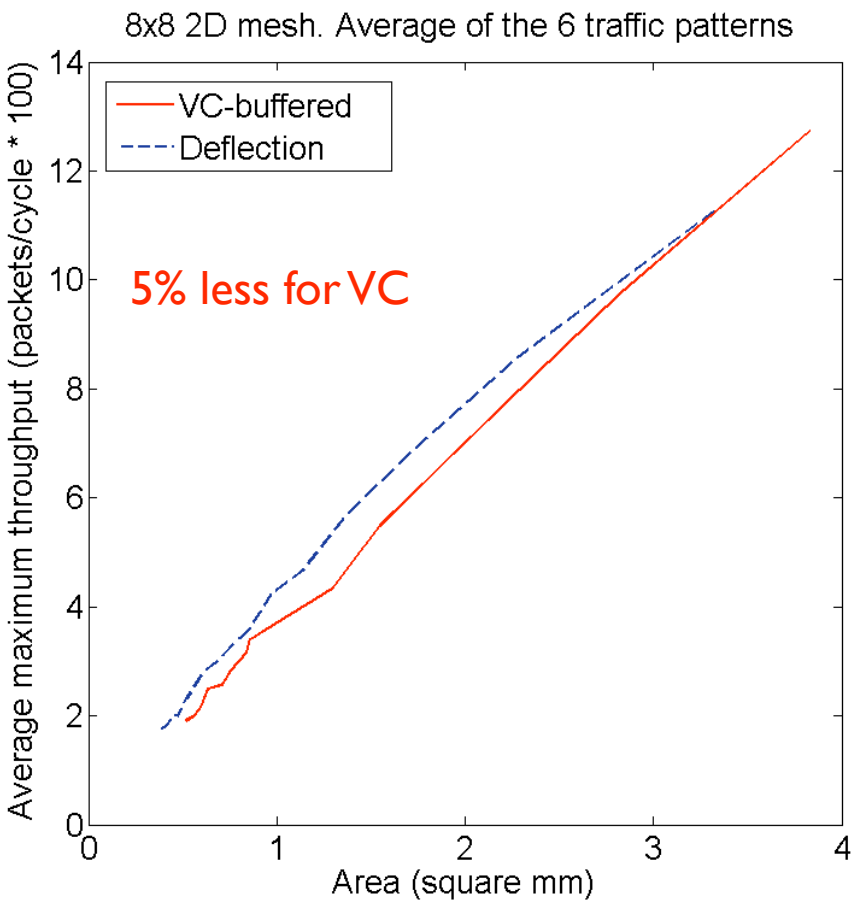

Latency distribution

- Blocking or deflection latency:

	Avg.	Max.	Std.
VC	0.75	13	1.18
Deflect.	4.87	108	8.09

- One deflection costs 6 cycles (2 hops)

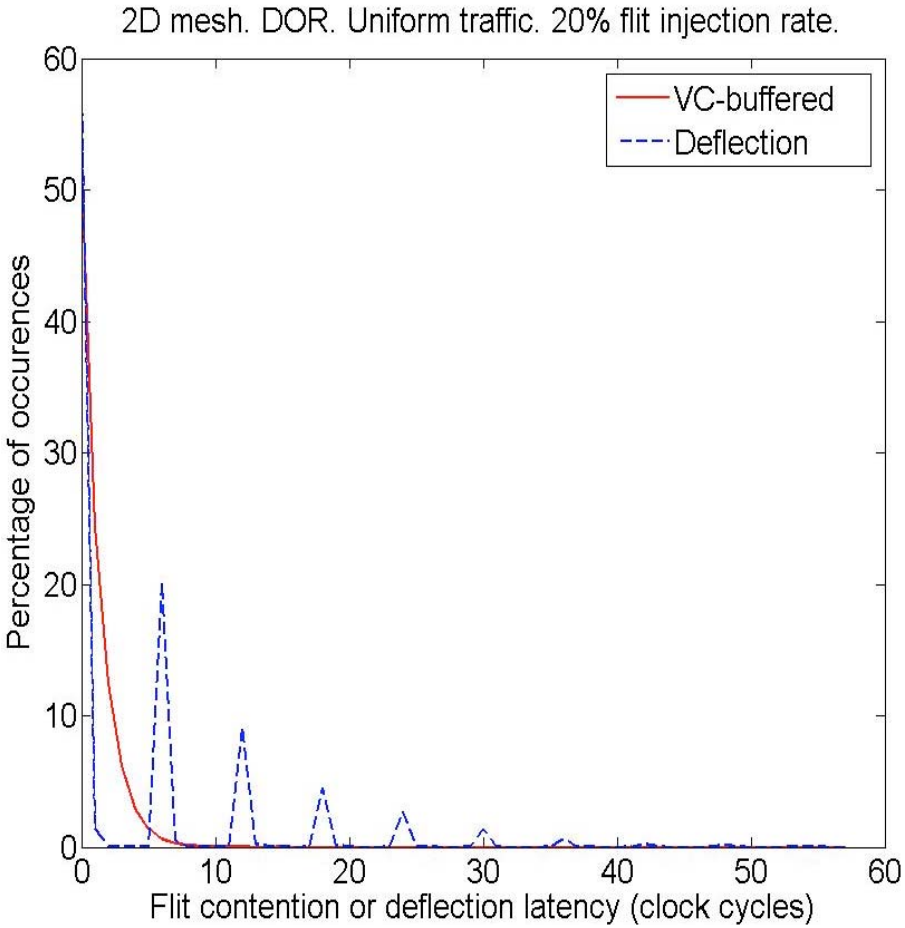

Power breakdown

BLESS: 4.6% activity factor increase.

20% flit injection rate

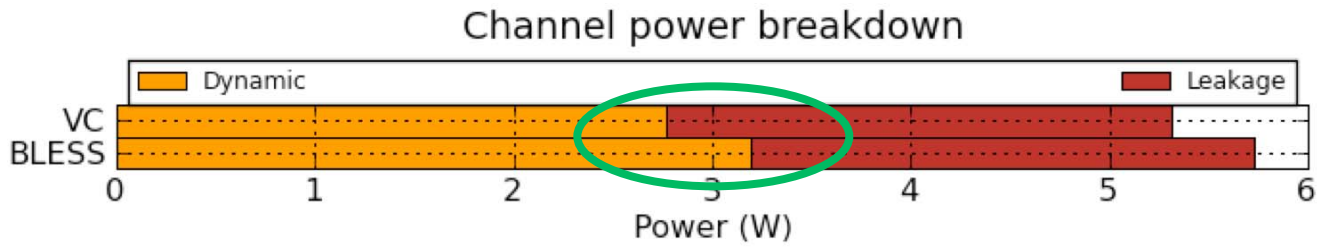

Buffer power: 2% compared to channel power. 7% without bypassing.

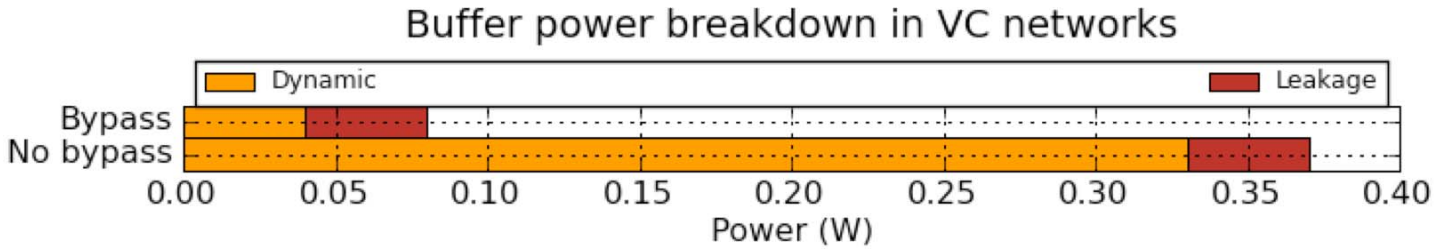

❑ Underlying cause:

- Reading & writing a buffer: 6.2pJ.
- One deflection: 42pJ. 6.7x the above.

Outline

- Methodology.
- Background.
- Optimizing routing in BLESS.
- Router microarchitecture.
- Network evaluation.
- Discussion.
  - Many parameters in such networks.
- Conclusion.

Discussion

- Topics covered in the paper in detail but not in this presentation:

- Low-swing channels: Favor deflection.
  - Never more than 1.5% less than VC power.
  - VC:16% more throughput per unit power.
  - VC becomes more area efficient.

- Endpoint complexity: Need complexity, such as backpressure if ejection buffers are full, or very large ejection buffers.

Discussion

- Points briefly mentioned in our study:

- Dropping networks: Same fundamental hop-buffering energy tradeoff.
  - Average hop count in dropping networks is affected more from topology and routing.
- Self-throttling sources: Hide network performance inefficiencies.
  - But CPU execution time really matters.
- Sub-networks, network size, more traffic classes: No clear trend.

Conclusion

- We compare VC and deflection networks. We show:
- Deflection network consumes marginally (1.5%) less energy at very low loads.
- VC network:
  - 12% lower average latency. Smaller std. dev.
  - 21% more throughput per unit power.
- Deflection network are more complex.
  - E.g. endpoint complexity & age-based allocation.
- Unless buffer cost unusually high, bufferless networks less efficient & more complex.
  - Designers should focus on optimizing buffers.

That's all folks

QUESTIONS?

Area breakdown

- Buffers 30% of the network area.