DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling

<u>Chen Sun</u>, Chia-Hsin Owen Chen, George Kurian, Lan Wei, Jason Miller, Anant Agarwal, Li-Shiuan Peh, Vladimir Stojanovic



NoC Cost Evaluation is Critical



• Many recent works utilize photonics





DAP

CAP

Photonics on-chip [Vantrease '08, Kurian '10]

Photonics to DRAM [Beamer '10, Udipi '11]

From processor

Mem

Stack

To processor

Fiber

Coupler

DIMM

Ring' modulators

• Many recent works utilize photonics





Photonics on-chip [Vantrease '08, Kurian '10]

Photonics to DRAM [Beamer '10, Udipi '11]

Tradeoffs of photonics not well explored

• Many recent works utilize photonics





Photonics on-chip [Vantrease '08, Kurian '10]

Photonics to DRAM [Beamer '10, Udipi '11]

- Tradeoffs of photonics not well explored
- At risk of being too optimistic

DAP

Bank

CAP

• Many recent works utilize photonics





Photonics to DRAM [Beamer '10, Udipi '11]

To processor

Mem

Stac

Fiber

Couple

Ring' modulators

- Tradeoffs of photonics not well explored
- At risk of being too optimistic
- Device/circuit designers need feedback











All these costs need to be visible to the network architect! 9/21/2012









• Original plan for DSENT, but...

- Original plan for DSENT, but...
- Photonics is dependent on electronics
 - Modulator drivers, Receivers
 - Serialize/Deserialize from core to link
 - Thermal ring resonance tuning



- Original plan for DSENT, but...
- Photonics is dependent on electronics
 - Modulator drivers, Receivers
 - Serialize/Deserialize from core to link
 - Thermal ring resonance tuning
- Need to compare electronics <u>fairly</u> with photonics...

X Scaling factors no longer valid for advanced processes

X Scaling factors no longer valid for advanced processes

X Very difficult to add technology or extend existing models

X Scaling factors no longer valid for advanced processes

- X Very difficult to add technology or extend existing models
- X Incomplete architectural models and timing for the router

- X Scaling factors no longer valid for advanced processes
- X Very difficult to add technology or extend existing models
- X Incomplete architectural models and timing for the router
- X All links are optimized for min-delay

- X Scaling factors no longer valid for advanced processes
- X Very difficult to add technology or extend existing models
- X Incomplete architectural models and timing for the router
- X All links are optimized for min-delay
- X Very low accuracies for modern technologies
 - 3X power overestimate for 65 nm, 400 MHz [Jeong, Kahng, et al. 2010]
 - 7X power, 2X area overestimate for 45 nm, 1 GHz
 - 5X+ power overestimate for links
 - Skewed breakdowns

- X Scaling factors no longer valid for advanced processes
- X Very difficult to add technology or extend existing models
- X Incomplete architectural models and timing for the router
- X All links are optimized for min-delay
- X Very low accuracies for modern technologies
 - 3X power overestimate for 65 nm, 400 MHz [Jeong, Kahng, et al. 2010]
 - 7X power, 2X area overestimate for 45 nm, 1 GHz
 - 5X+ power overestimate for links
 - Skewed breakdowns

A 10-year-old model that worked well, but insufficient now

Design Space Exploration of Networks Tool

Overview

- Overview
- Methodology
 - Improvements to electrical modeling frameworks
 - Incorporate photonics models

- Overview
- Methodology
 - Improvements to electrical modeling frameworks
 - Incorporate photonics models
- Example cross-hierarchical network evaluation

- Overview
- Methodology
 - Improvements to electrical modeling frameworks
 - Incorporate photonics models
- Example cross-hierarchical network evaluation
- Conclusion

Structure of DSENT

- Written in C++ (Object-Oriented)
- Fast Evaluations, few seconds
- ASIC-driven approach
- Made flexible, extensible

Two Ways to Use DSENT

• Stand-alone for design space exploration

Two Ways to Use DSENT

- Stand-alone for design space exploration
 - Takes network parameters, queries, technology, give back area, power



Two Ways to Use DSENT

- Stand-alone for design space exploration
 - Takes network parameters, queries, technology, give back area, power


Two Ways to Use DSENT

- Use with architectural simulator for app-driven power traces
 - Uses event counts [Kurian, IPDPS 2012]



DSENT

Design Space Exploration of Networks Tool

- Overview
- Methodology
 - Improvements to electrical modeling frameworks
 - Incorporate photonics models
- Example cross-hierarchical network evaluation
- Conclusion

ASIC-like modeling flow, generates primitives/standard cells



ASIC-like modeling flow, generates primitives/standard cells

Keep relevant tech parameters, simplify technology entry



ASIC-like modeling flow, generates primitives/standard cells

Keep relevant tech parameters, simplify technology entry

Delay model, timing-constrained cell sizing, electrical links



ASIC-like modeling flow, generates primitives/standard cells

Keep relevant tech parameters, simplify technology entry

Delay model, timing-constrained cell sizing, electrical links



- ASIC-like flow, standard cell based
- Keep relevant tech parameters, simplify technology entry
- Delay model, timing-constrained cell sizing, electrical links
- Able to model more generic digital, beyond just routers
- Methodology targeted for 45 nm and below

- ASIC-like flow, standard cell based
- Keep relevant tech parameters, simplify technology entry
- Delay model, timing-constrained cell sizing, electrical links
- Able to model more generic digital, beyond just routers
 - Methodology targeted for 45 nm and below

0

Power/Area estimates accurate to ~20% of SPICE simulation

	Model	Reference Point	DSENT
Router (6x6)	Buffer (mW)	SPICE – 6.93	7.55 (+9%)
	Xbar (mW)	SPICE – 2.14	2.06 (+4%)
	Control (mW)	SPICE – 0.75	0.83 (+11%)
	Clock (mW)	SPICE – 0.74	0.63 (-15%)
	Total (mW)	SPICE – 10.6	11.2 (+6%)
	Area (mm ²)	Encounter – 0.070	0.062 (-11%)

45 nm SOI

•

٠

•

- 6 Input ports, 6 output ports
- 64-bit flit width
- 8 VCs/Port, 16 Buffer FIFO
- 1 GHz clock
- 0.16 flit injection rate



- Four different sources of power consumption
 - Modulator, receivers
 - Laser power
 - Thermal tuning
 - Serialize, deserialize backends



- Modulator becomes more expensive with:
 - High data-rate
 - Higher modulation depth (extinction ratio)

9/21/2012

Lower insertion loss



- Receiver becomes more expensive with:
 - High data-rate

- Receiver sensitivity degrades with:
 - High data-rate
 - Lower modulation depth
 - Higher bit error rate requirement



- Laser power requirement gets worse with:
 - Higher receiver sensitivity requirement
 - Higher channel losses, e.g. higher modulator insertion loss



• Ring resonator devices are sensitive to process, temperature, active tuning is needed



- Ring resonator devices are sensitive to process, temperature, active tuning is needed
- Not necessarily a fixed cost per ring!
 - [Georgas CICC 2011, Nitta HPCA 2011]



- Ring resonator devices are sensitive to process, temperature, active tuning is needed
- Not necessarily a fixed cost per ring!
 - [Georgas CICC 2011, Nitta HPCA 2011]

DSENT models schemes for tuning, impact of process sigmas



- Ring resonator devices are sensitive to process, temperature, active tuning is needed
- Not necessarily a fixed cost per ring!
 - [Georgas CICC 2011, Nitta HPCA 2011]

DSENT models schemes for tuning, impact of process sigmas

Serializer/Deserializers are taken care of by electrical framework

DSENT

Design Space Exploration of Networks Tool

- Overview
- Methodology
 - Improvements to electrical modeling frameworks
 - Incorporate photonics models
- Example cross-hierarchical network evaluation
- Conclusion

Example Study

• 256-core clos network, energy per bit as metric – Pclos, EClos normalized to same throughput/latency



- 128-bit Flit Width
- 16 ingress, middle, egress routers, k, n, r = 16, 16, 16
- 2 GHz
- 1 dB/cm waveguide loss

Compare at

- 45nm (present)
- 11nm (future)

[Joshi, NOCS 2009]

Two Types of Power

- Data-dependent vs. non-data-dependent power
- Optical components (laser, thermal tuning) are non-data-dependent

Data-Dependent	Non-Data-Dependent
Router data-path/control	Leakage
Electrical links	Un-gated clocks
Gated clocks	Laser
Receiver/Modulator	Thermal tuning, ring heating
SerDes	

Effect of Utilization



Effect of Utilization



Effect of Utilization



Energy Breakdown at Max Network Throughput (33 Tb/s)





60



Energy Breakdown at Low Network Throughput (4.5 Tb/s)



"Wow non-data-dependent laser really hurts, can I make it better?"

Energy Breakdown at Low Network Throughput (4.5 Tb/s)



"Wow non-data-dependent laser really hurts, can I make it better?"

Optimistic device guy: "No problem, I go make my devices better!"

Evaluate the effect of waveguide losses

"How much better does he need to do in order to beat the competing 11nm electrical?"



9/21/2012



9/21/2012



Evaluate the effect of waveguide losses



9/21/2012

• Story doesn't end here...

- Story doesn't end here...
 - Thermal tuning strategies

- Story doesn't end here...
 - Thermal tuning strategies
 - Data-rates, change number of optical devices

- Story doesn't end here...
 - Thermal tuning strategies
 - Data-rates, change number of optical devices
 - Modulator, laser balance
 - Modulator is DD, laser is NDD
Tech Parameter Study

- Story doesn't end here...
 - Thermal tuning strategies
 - Data-rates, change number of optical devices
 - Modulator, laser balance

- Modulator is DD, laser is NDD

These are examples of DSENT models

• Design decisions in NoCs require evaluation



- Design decisions in NoCs require evaluation
- We created DSENT to bridge photonics and electronics
 - Generalized methodology for digital components
 - Moves beyond fixed number evaluations for photonics
 - Includes power/area models for several networks



- Design decisions in NoCs require evaluation
- We created DSENT to bridge photonics and electronics
 - Generalized methodology for digital components
 - Moves beyond fixed number evaluations for photonics
 - Includes power/area models for several networks
- We showed how DSENT can be used to capture the tradeoffs for an example photonic clos network
 - Utilization-dependent energy plots
 - Data-dependent and non-data-dependent power
 - Investigate network sensitivity to optical parameters



- Design decisions in NoCs require evaluation
- We created DSENT to bridge photonics and electronics
 - Generalized methodology for digital components
 - Moves beyond fixed number evaluations for photonics
 - Includes power/area models for several networks
- We showed how DSENT can be used to capture the tradeoffs for an example photonic clos network
 - Utilization-dependent energy plots
 - Data-dependent and non-data-dependent power
 - Investigate network sensitivity to optical parameters
- Continuing and future work
 - Ease user model specification to aid microarchitecture studies
 - Automatically form estimates for local interconnect



Thank You

- We would like to acknowledge
 - Integrated Photonics teams at MIT and University of Colorado, Boulder for models
 - Prof. Dmitri Antoniadas's group for their sub-45nm transistor models
- Support
 - DARPA, NSF, FCRP IFC, SMART LEES, Trusted Foundry, Intel, APIC, MIT CICS, NSERC

For more info, visit

https://sites.google.com/site/mitdsent/

(we will make it downloadable following the conference)



Backups

Evaluation Configuration

Network Configuration	Values
Number of tiles	256
Chip area (divided equally amongst tiles)	$400 \mathrm{mm^2}$
Packet length	80 Bytes
Flit width	128 bits
Core frequency	2 GHz
Clos configuration (m, n, r)	16, 16, 16
Link latency	2 cycles
Link throughput	128 bits/core-cycle
Router Configuration	Values
Number pipelines stages	3
Number virtual channels (VC)	4
Number buffers per VC	4

Evaluation Parameters

Technology Parameters	Default Values		
Process technology	11 nm TG		
Optical link data-rate	2 Gb/s		
Laser efficiency	0.25		
Coupler loss	$2 \mathrm{dB}$		
Waveguide loss	1 dB/cm		
Ring drop loss	1 dB		
Ring through loss	0.01 dB		
Modulator loss (optimized)	0.01 - 10.0 dB		
Modulator extinction (optimized)	0.01 - 10.0 dB		
Photodetector Capacitance	5 fF		
Link bit error rate	10^{-15}		
Ring tuning model	Bit-Reshuffled [12, 41]		
Ring heating efficiency	100 K/mW		

Orion Specifics

- Missing decoder and mux for register-type buffer
- Flops based on cross-coupled NOR gates
 - Uses old Cacti decoder sizing
- Missing pipeline flops energy on the data-path
 - Though clock power of those is added
- Clock H-tree optimized by data link model
 - Optimal delay H-tree

DSENT Modeling Methodology



Technology Characterization



Optical Models

- Models for major optical components
 - Waveguide, ring, coupler, modulator, photodetecter ...
- Models for peripheral circuitry
 - Modulator driver, receiver, SerDes, thermal tuning



Timing Optimization

- A greedy algorithm to select the standard cell sizes
 - Make circuit meet the timing constraint



Expected Transitions

 A simplified expected transition probability model



Power Breakdown (Half)

Energy Break-Down at Half Network Throughput (16 Tb/s)



 Photonics (P45, P11) are roughly even with electronics

Network Case Study



Fig. 7: A comparison of three thermal-tuning strategies discussed in Section V-C. Link data-rate is used as a degree of freedom to balance tuning power with other sources of power consumption. Since the throughput of each link is 128 bits/core-cycle at a 2GHz core clock, a data-rate of 2, 4, 8, 16, 32 Gb/s per wavelength (λ) implies 128, 64, 32, 16, 8 wavelengths per link, respectively. All energy breakdowns are shown for half of saturation throughput (16.5 Tb/s.

Photonic Technology Scaling

• Waveguide loss



(a) Sensitivity to waveguide loss. Energy per bit vs throughput (left) and energy per bit breakdown at 16 Tb/s throughput (right)

• Ring heating efficiency



(b) Sensitivity to heating efficiency. Energy per bit vs throughput (left) and energy per bit breakdown at 16 Tb/s throughput (right)

Tool Validation (45nm SOI)

	Model	Reference Point	DSENT	Orion2.0 +	Orion2.0 Mod*
Rin	g Modulator Driver (fJ/b)	50 (11 Gb/s)	60.87	N/A	N/A
Rec	eiver (fJ/b)	52 (3.5 Gb/s 45nm)	43.02	N/A	N/A
Router (6x6)	Buffer (mW)	SPICE – 6.93	7.55	34.4	3.57
	Xbar (mW)	SPICE – 2.14	2.06	14.5	1.26
	Control (mW)	SPICE – 0.75	0.83	1.39	0.31
	Clock (mW)	SPICE – 0.74	0.63	28.8	0.36
	Total (mW)	SPICE – 10.6	11.2	91.3	5.56
	Area (mm²)	Encounter – 0.070	0.062	0.129	0.067

Default Orion 2.0 technology parameters for 45nm
*Correctly specified 45nm tech params



- 6 Input ports, 6 output ports
- 64-bit flit width
- 8 VCs/Port, 16 Buffer FIFO
- 1 GHz clock
- 0.16 flit injection rate

DSENT Framework



 Required technology input mostly limited to what is attainable through ITRS projections and other roadmaps
 9/21/2012

DSENT Framework



DSENT Framework



9/21/2012

[Georgas, CICC 2051]

Misc

Error in Cacti 6.5

Validation Targets	Published Pleakage	CACTI [1] result	CACTI [1] Error
Xeon 16MB 65nm L3	6.6 W	27.2 W	412%
Penryn 6MB 45nm L2	1.7 W	9.65 W	568%
Intel 32nm 128 Kb subarray	5 mW	36.5 mW	630%

[S. Li, ICCAD 2011]