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

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

rle AT MIT

CSAIL

Massachusetts Institute of Technology

# NoC Cost Evaluation is Critical

**Performance**

- TFLOPS in a 100+ core chip
- Tb/s of on-chip traffic
- Tb/s of off-chip I/O

**Power**

- Energy-constrained chips
- Moving bits costs energy
- Wires are not scaling much

**Area**

- Network should not consume too much area

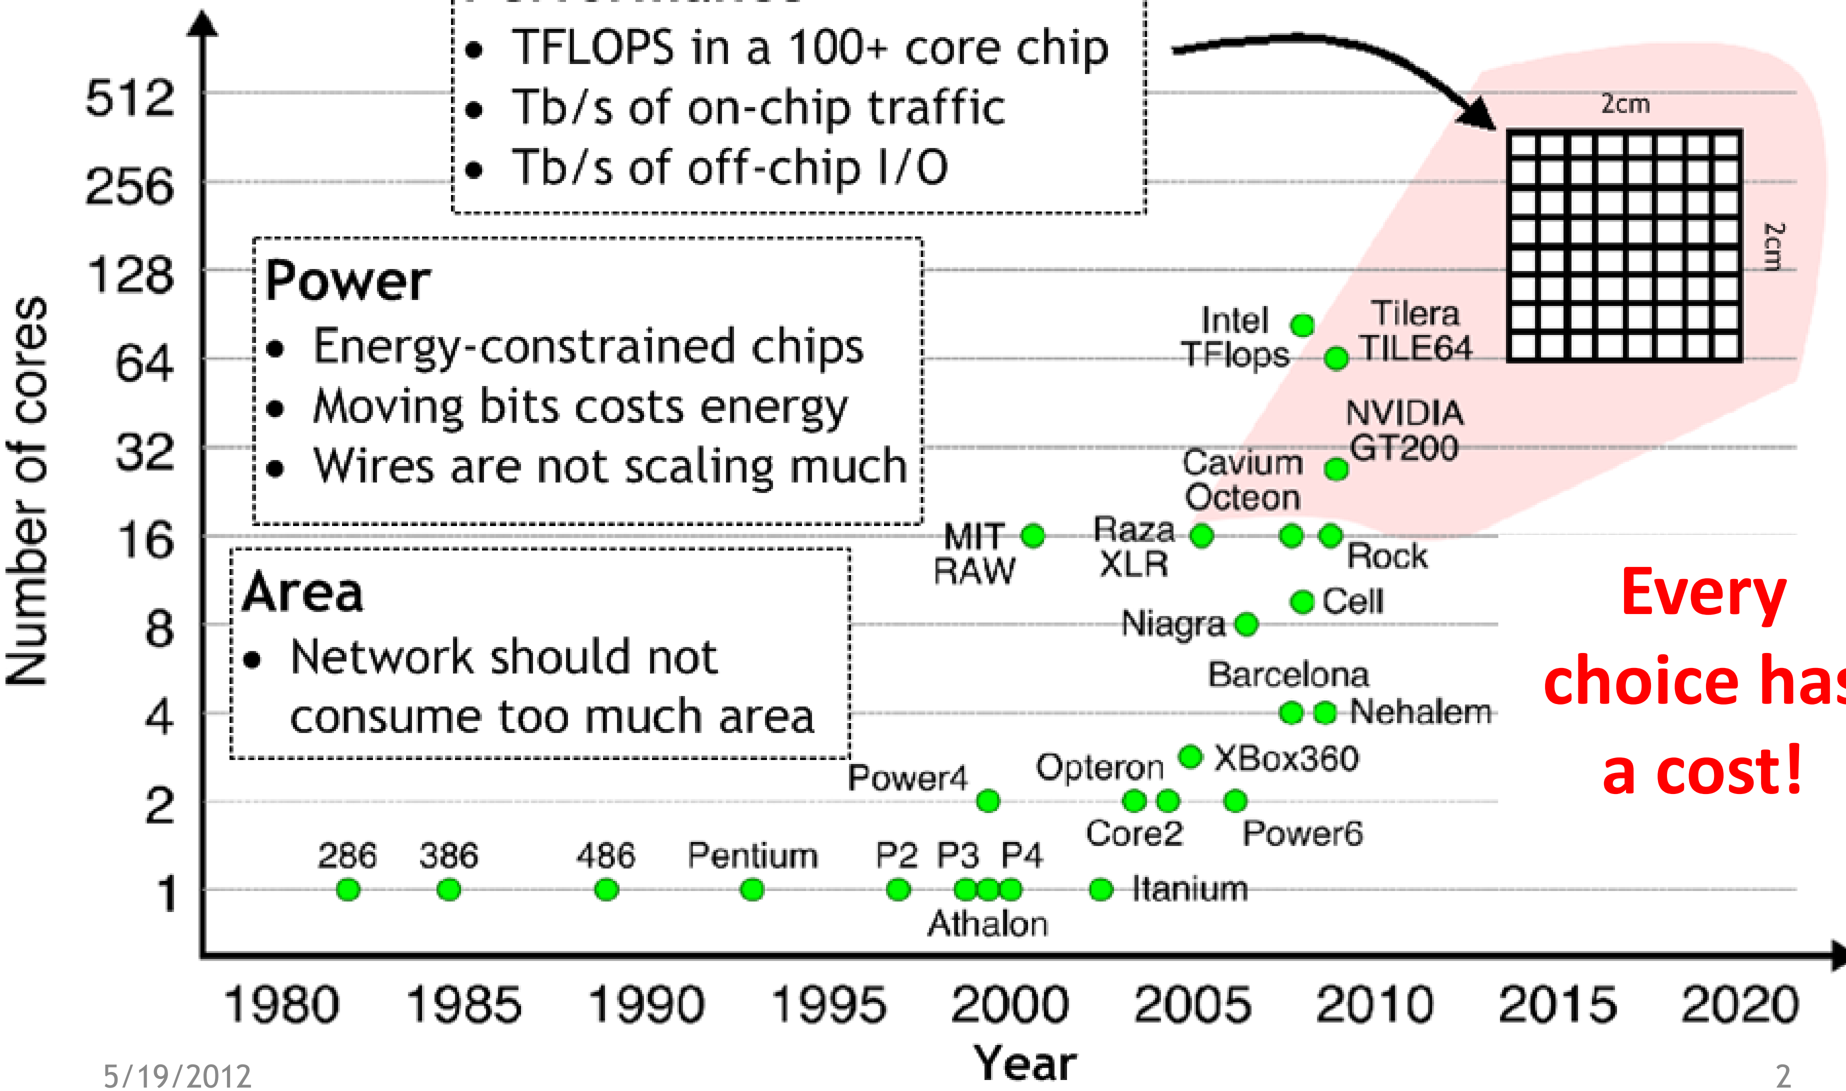


**Every choice has a cost!**

# Potential for Photonics

- Many recent works utilize photonics

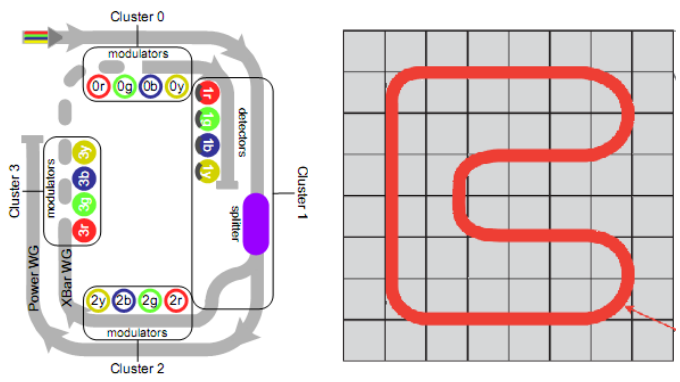


Photonics on-chip [Vantrease ’08, Kurian ‘10]

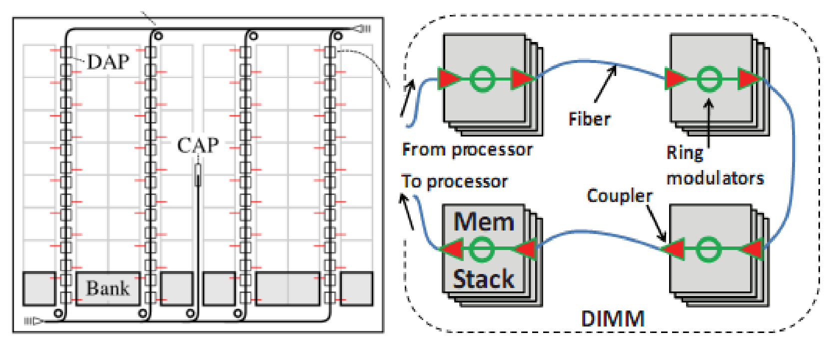


Photonics to DRAM [Beamer ‘10, Udipi ‘11]

# Potential for Photonics

- Many recent works utilize photonics

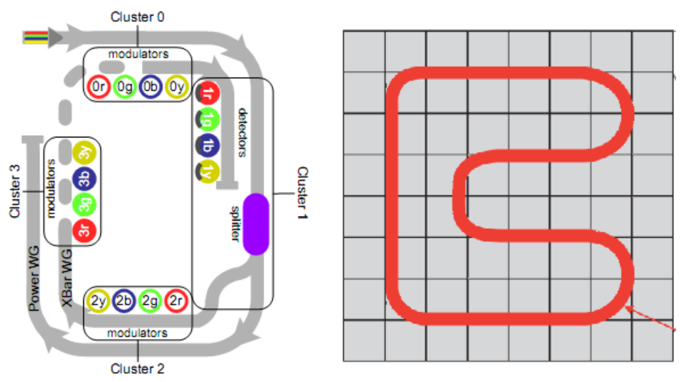


Photonics on-chip [Vantrease '08, Kurian '10]

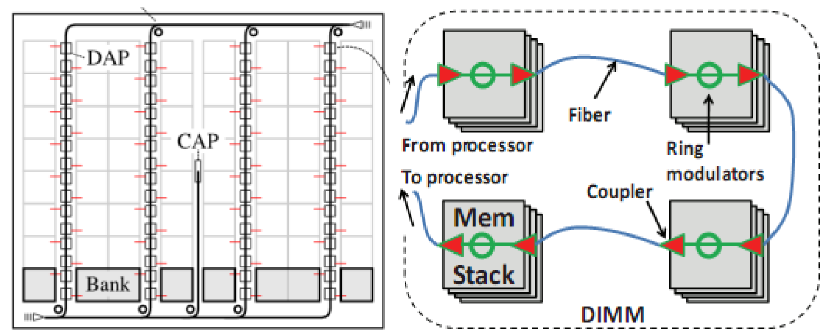


Photonics to DRAM [Beamer '10, Udipi '11]

- Tradeoffs of photonics not well explored

# Potential for Photonics

- Many recent works utilize photonics

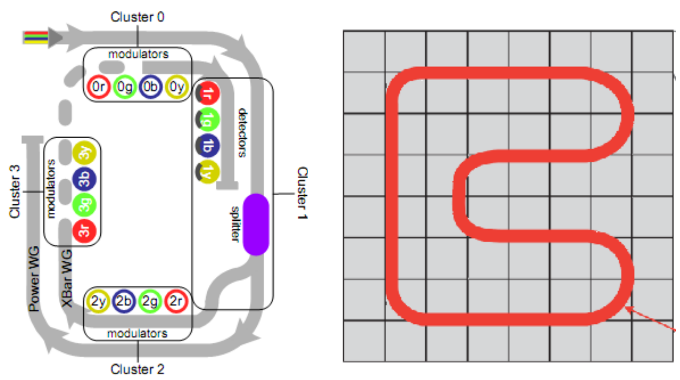


Photonics on-chip [Vantrease '08, Kurian '10]

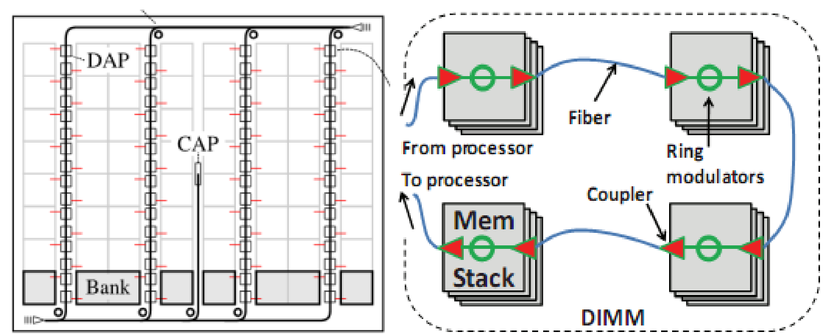


Photonics to DRAM [Beamer '10, Udipi '11]

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

# Potential for Photonics

- Many recent works utilize photonics

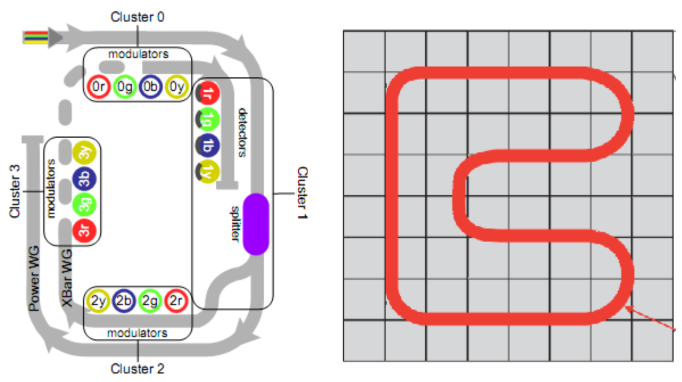


Photonics on-chip [Vantrease '08, Kurian '10]

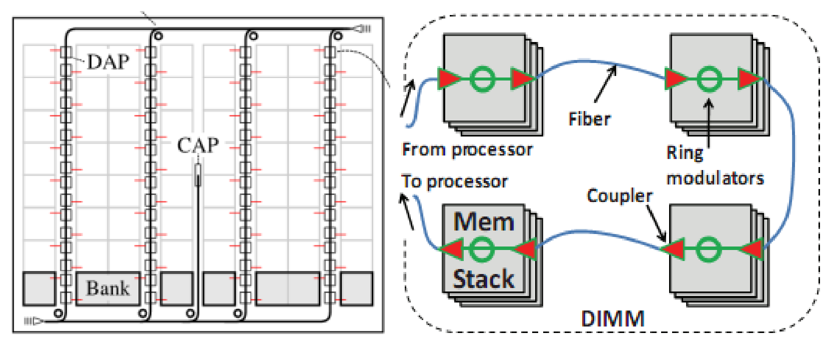


Photonics to DRAM [Beamer '10, Udipi '11]

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

# What does a NoC Cost?

**Network**

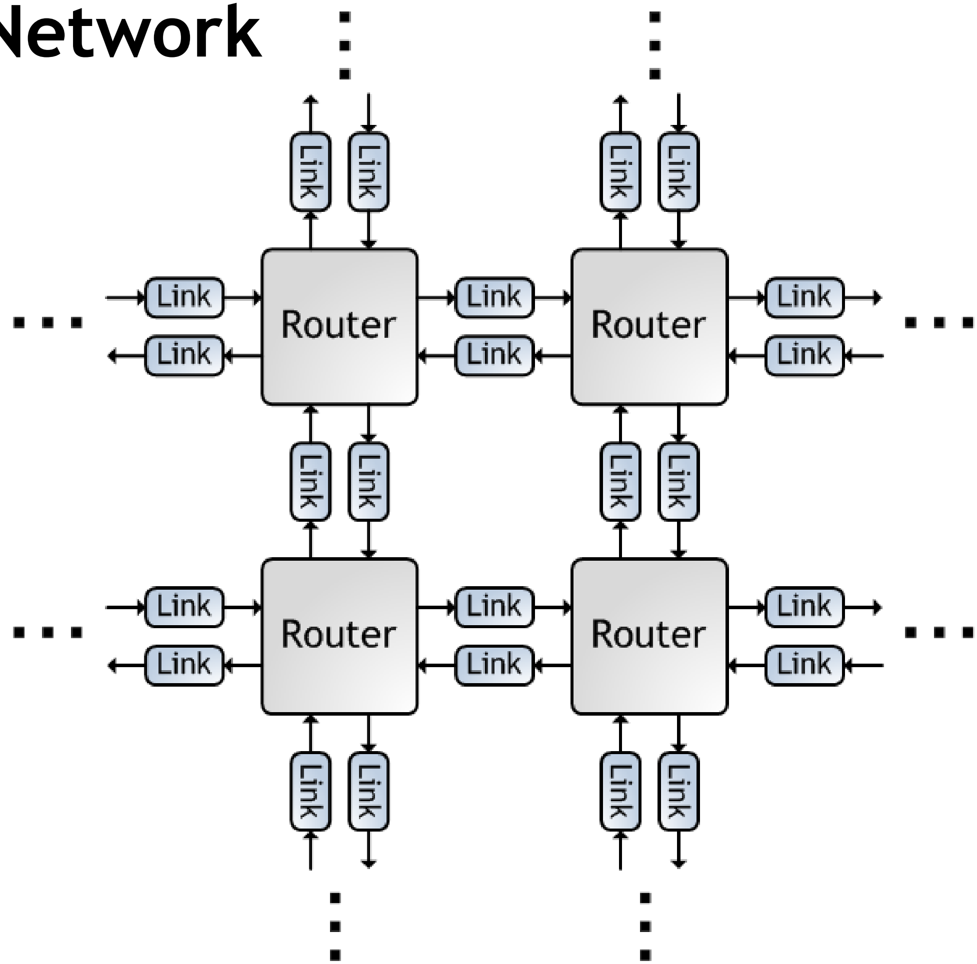

# What does an NoC Cost?

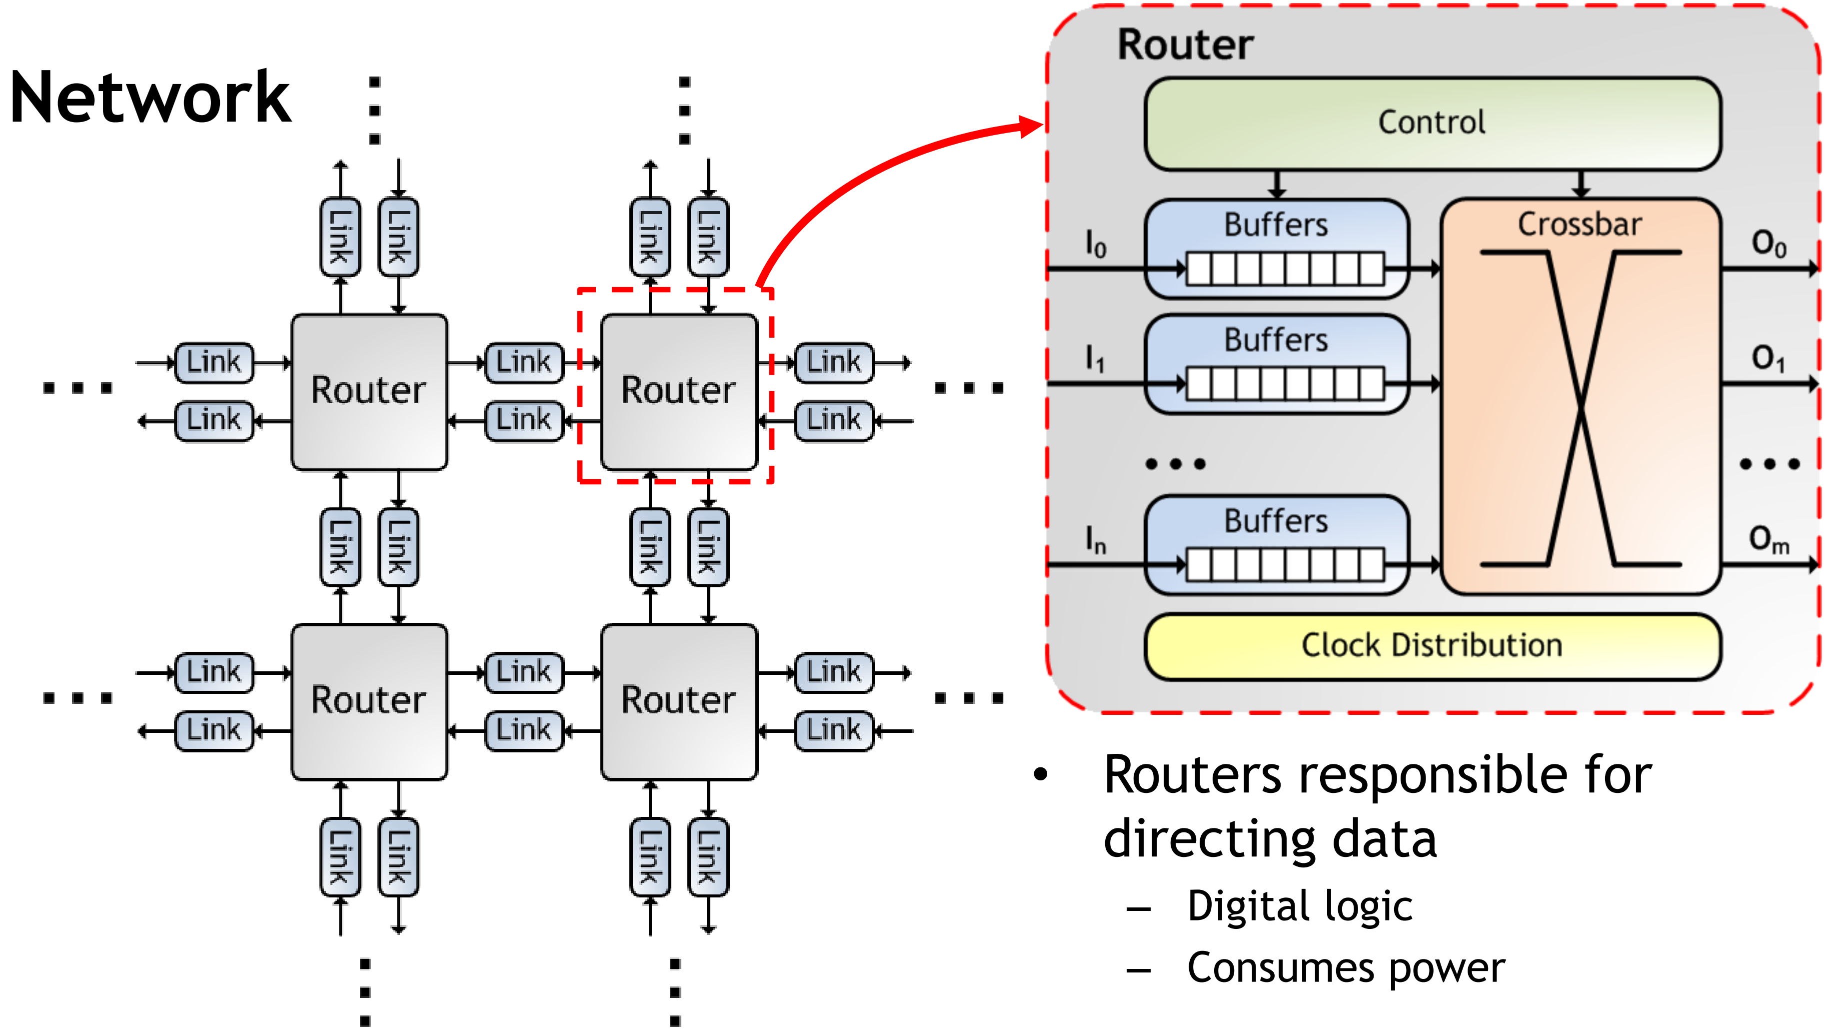


- Routers responsible for directing data
  - Digital logic
  - Consumes power

# What does a NoC Cost?

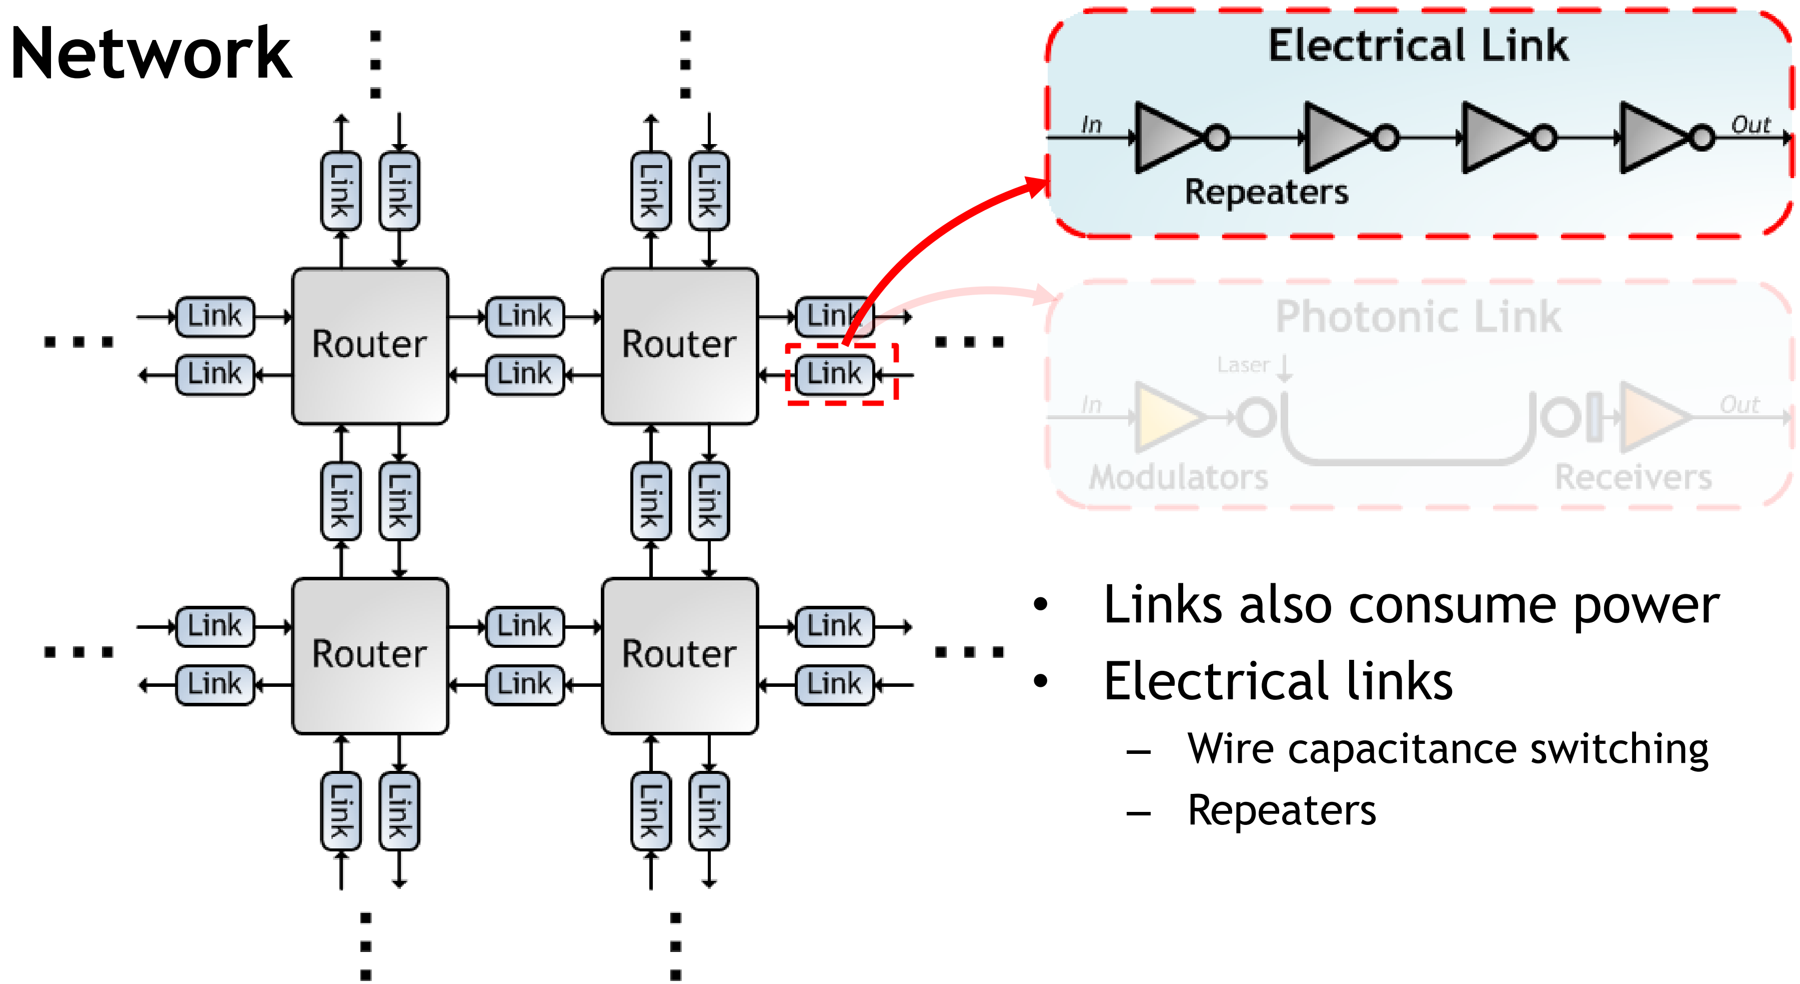


- Links also consume power
- Electrical links
  - Wire capacitance switching
  - Repeaters

# What does a NoC Cost?

Network

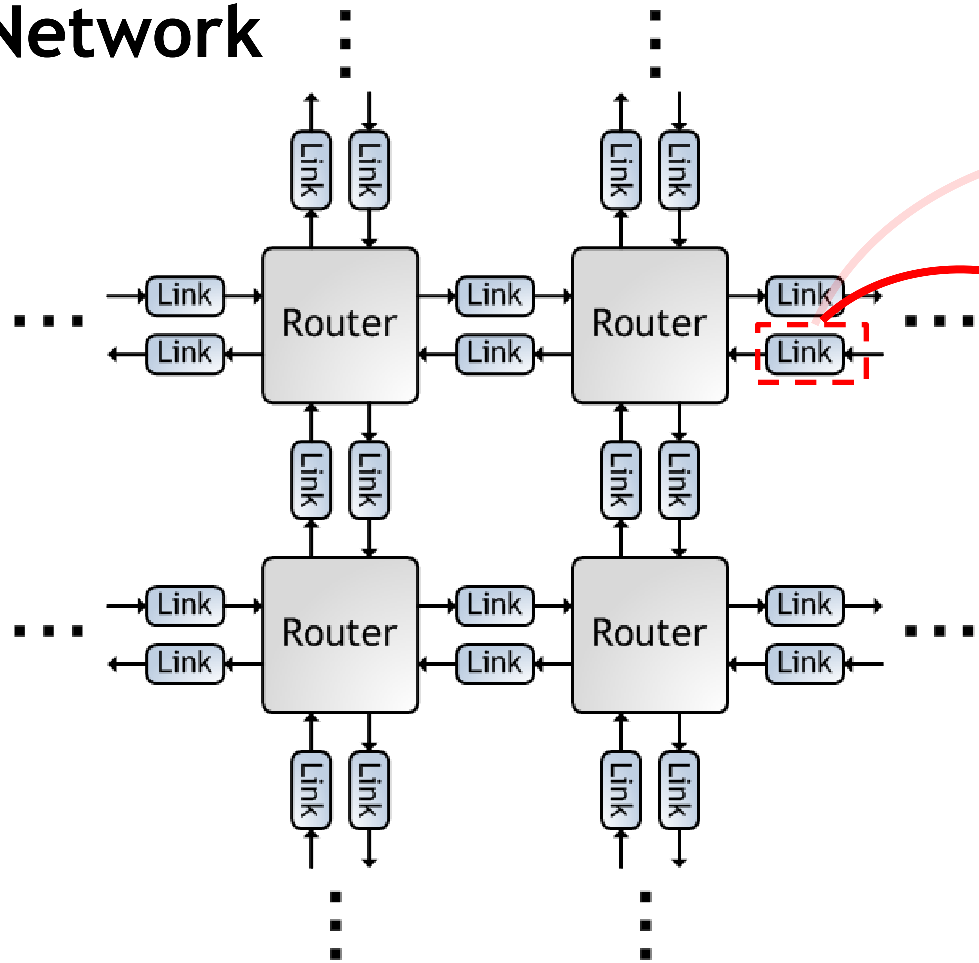


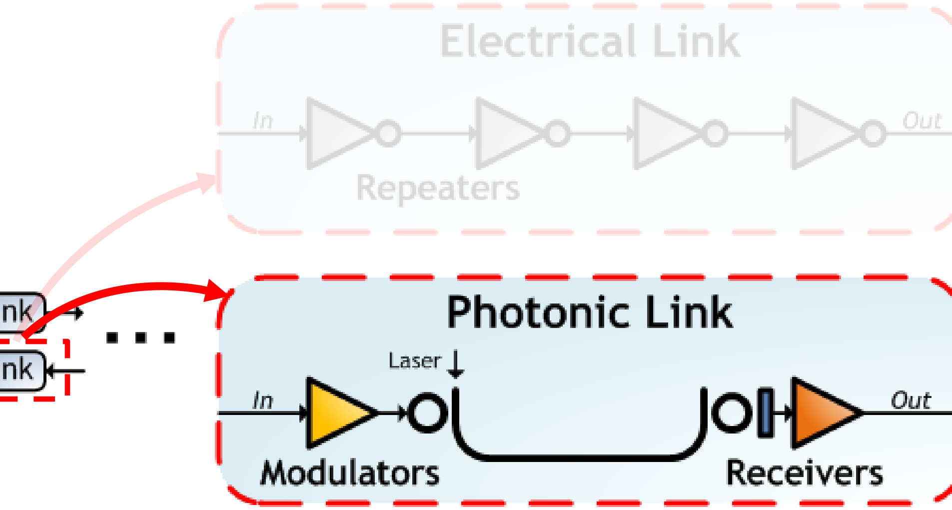


- Photonic links
  - Receivers, Modulators
  - Laser
  - Ring thermal tuning
  - Serialize/Deserialize

# What does a NoC Cost?

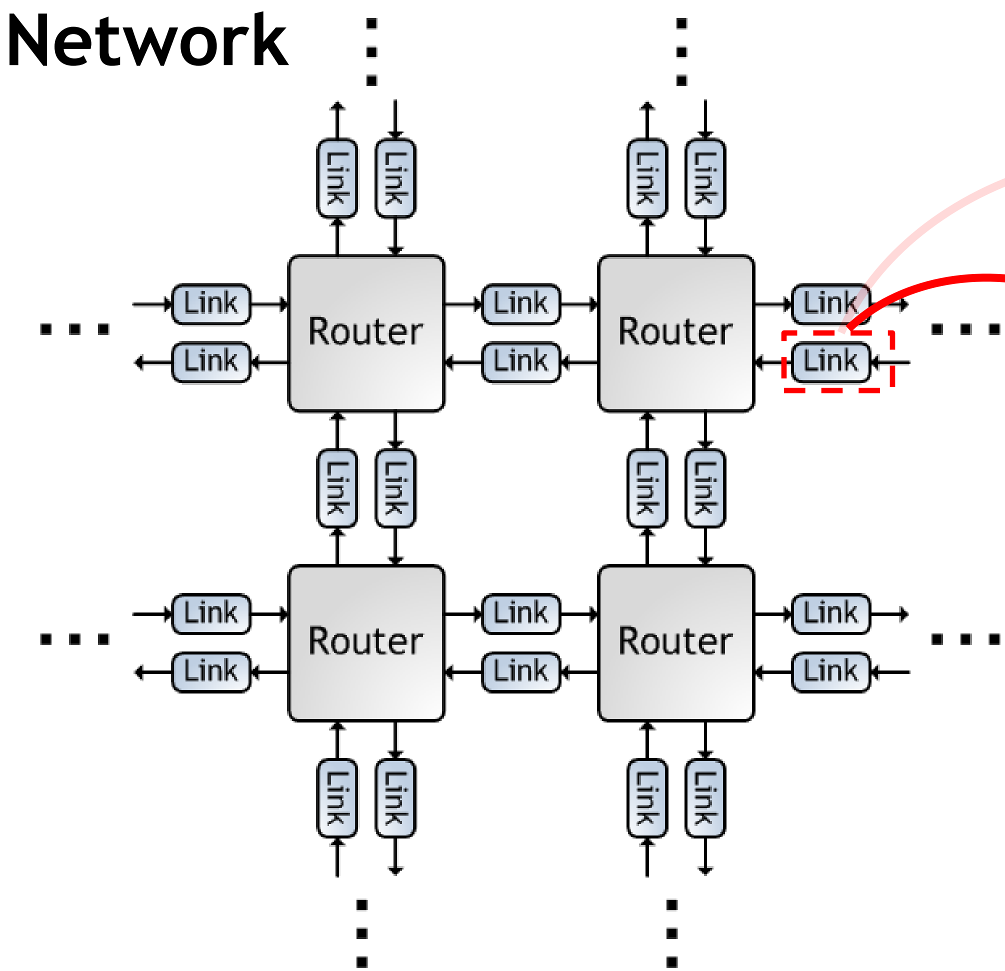


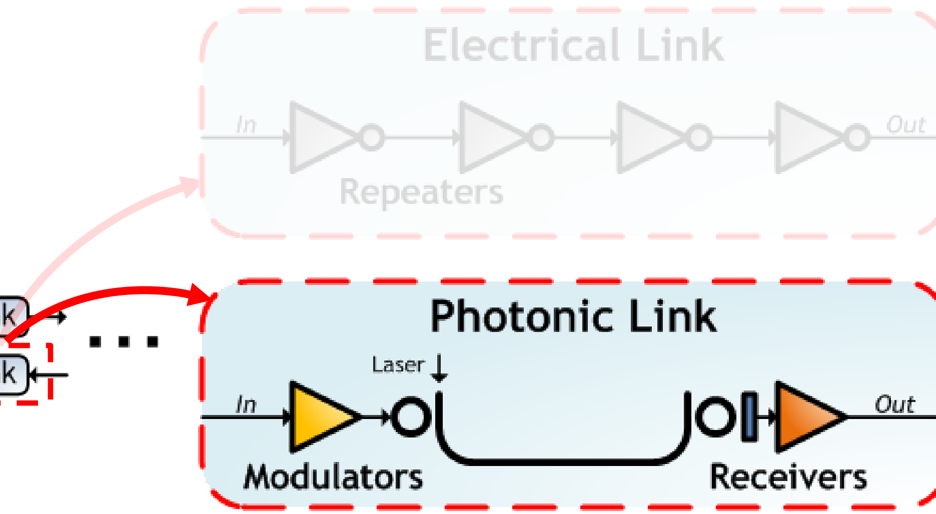


- Photonic links
  - Receivers, Modulators
  - Laser
  - Ring thermal tuning
  - Serialize/Deserialize

**All these costs need to be visible to the network architect!**

# Existing Architectural Tools

# Existing Architectural Tools

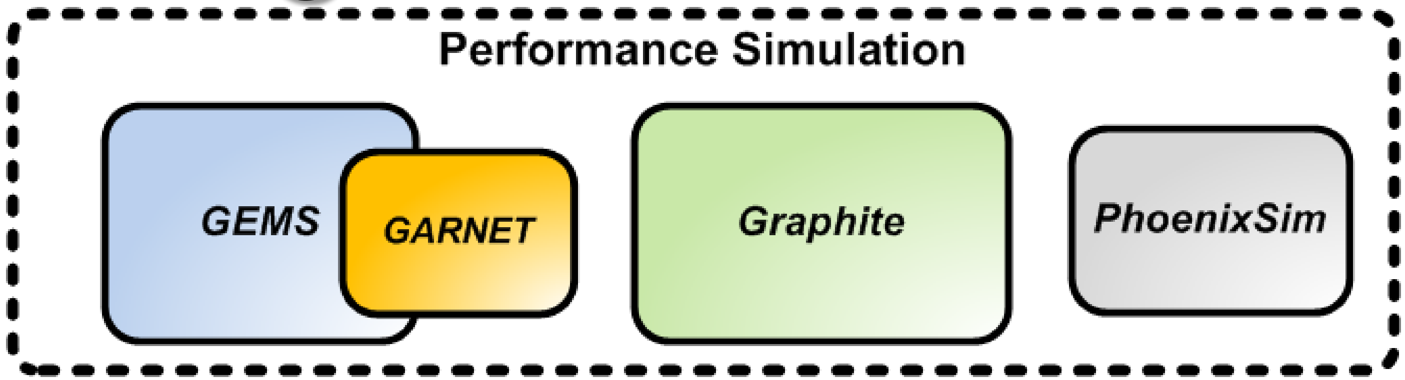

# Existing Architectural Tools

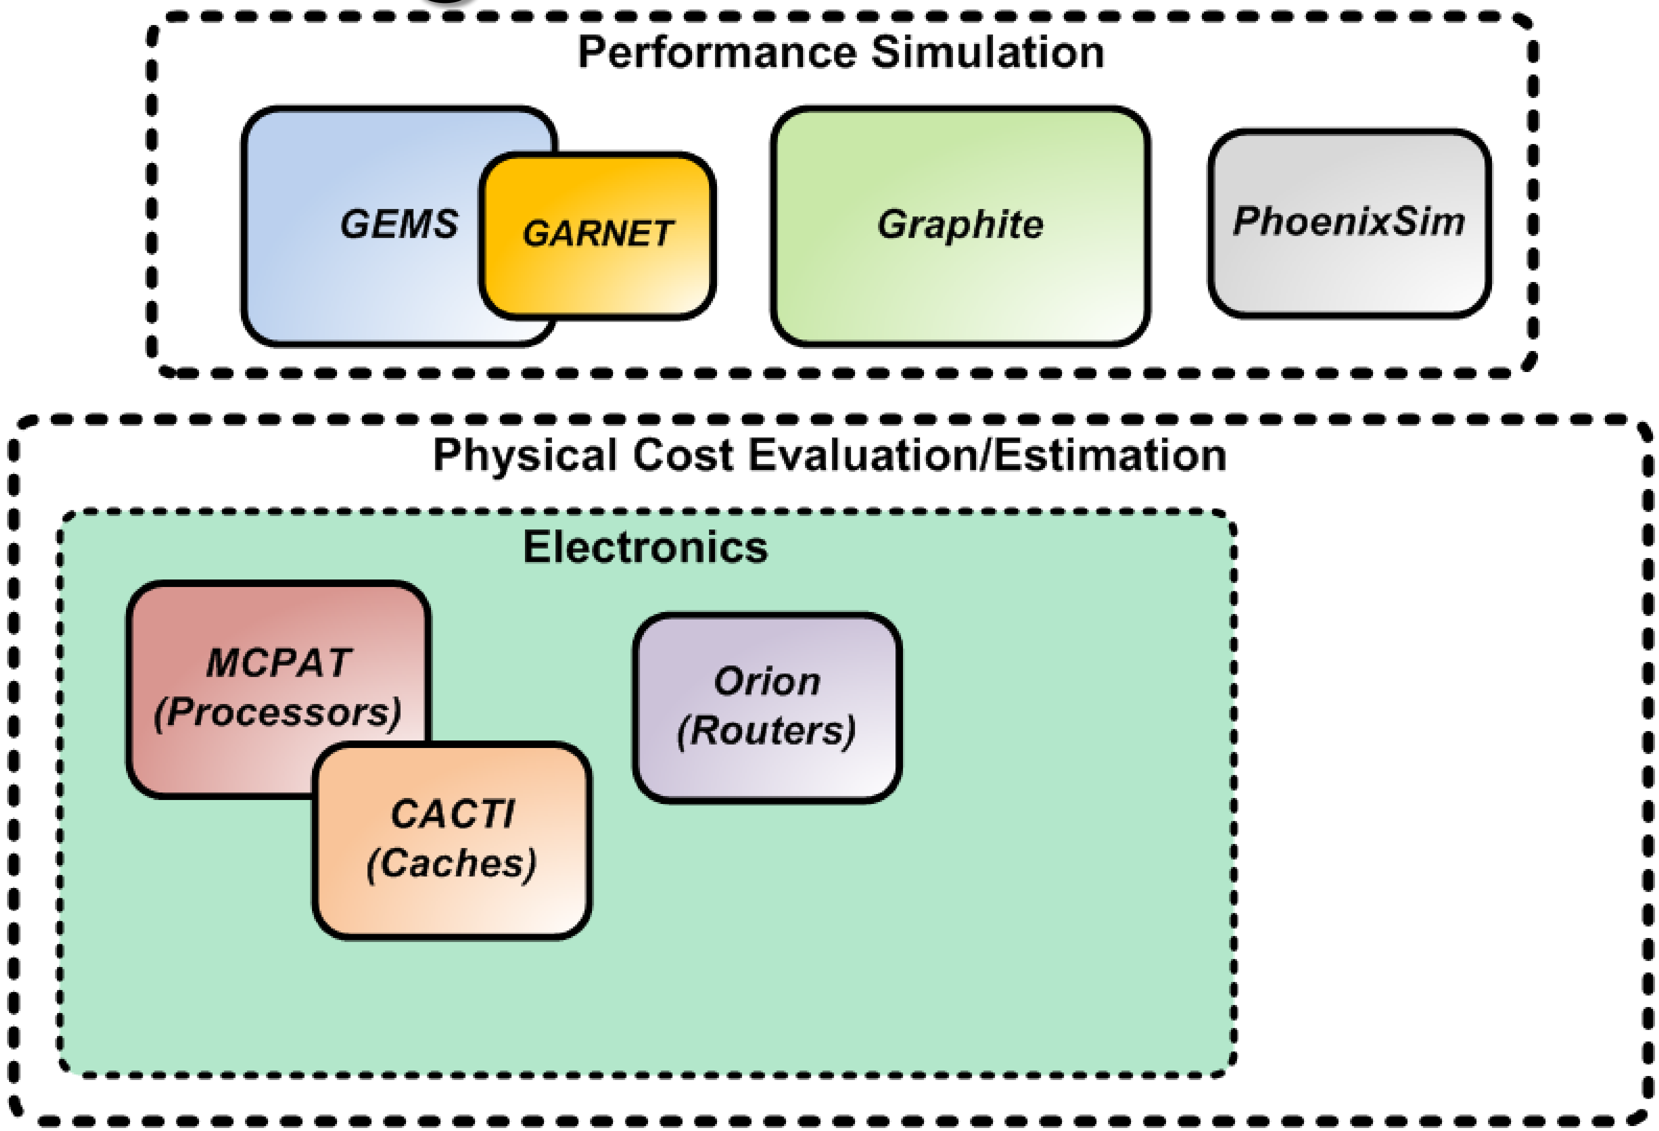

# Existing Architectural Tools

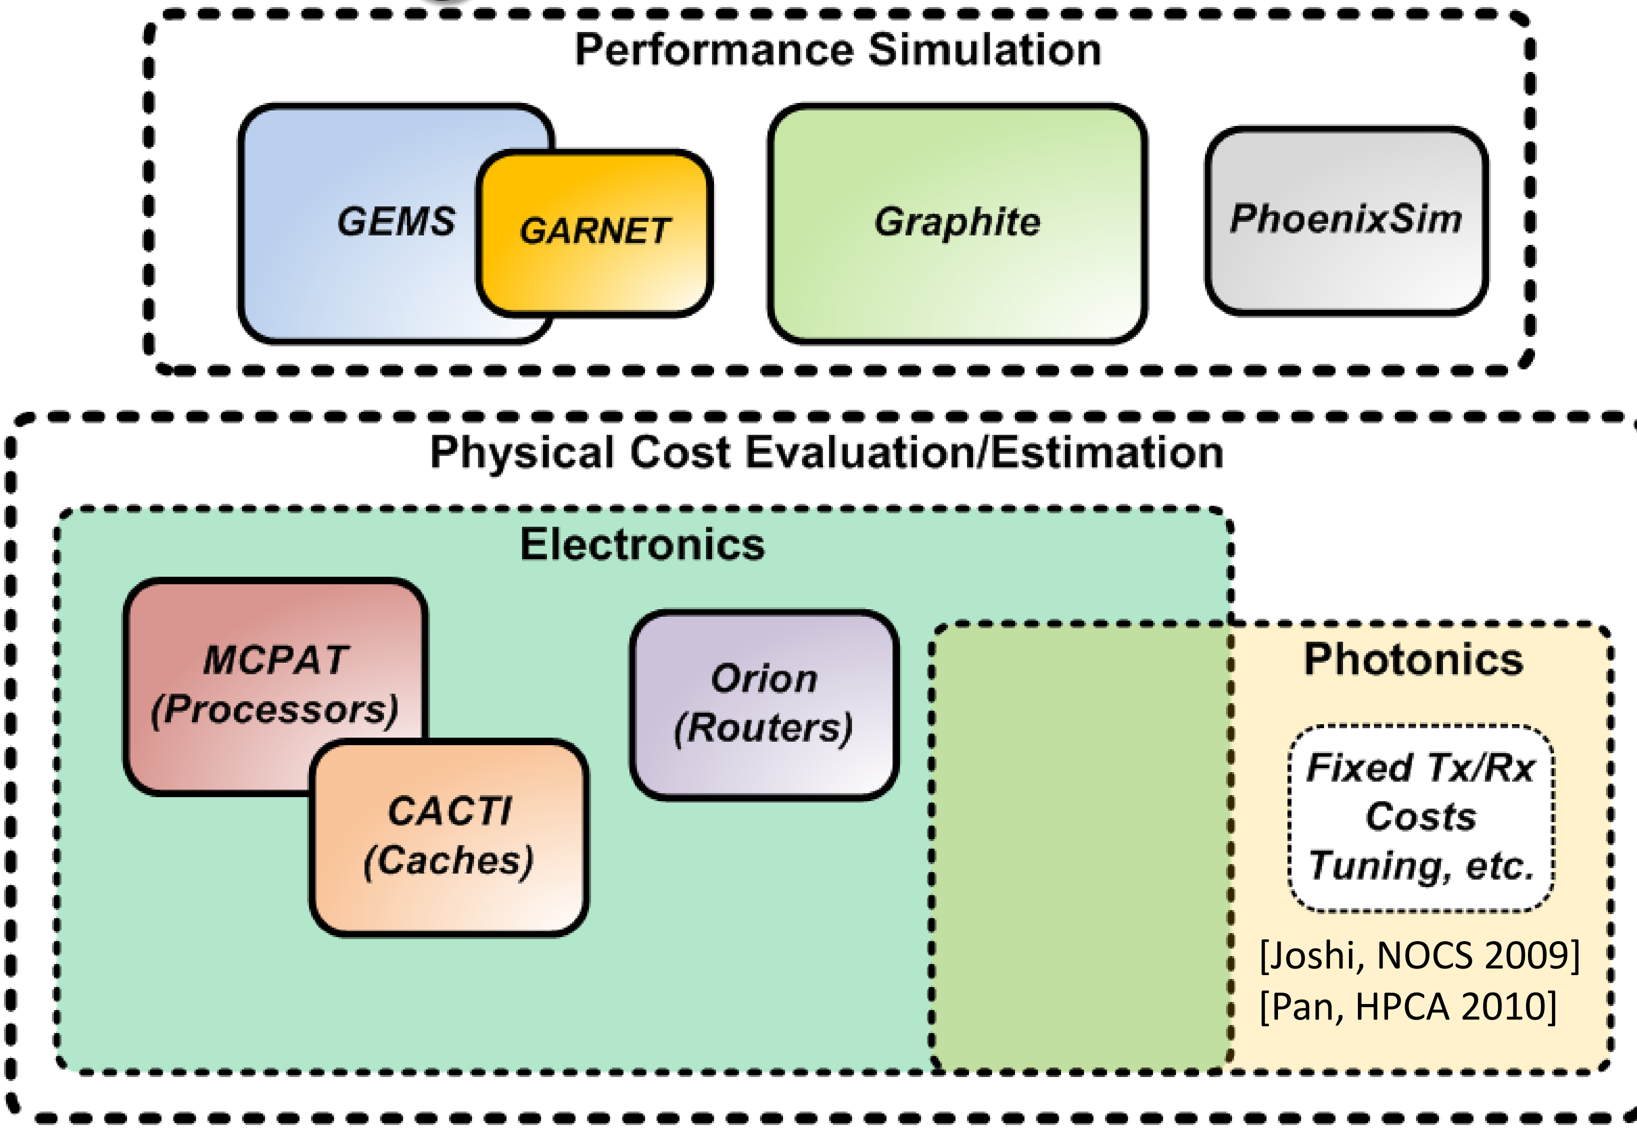

# Existing Architectural Tools

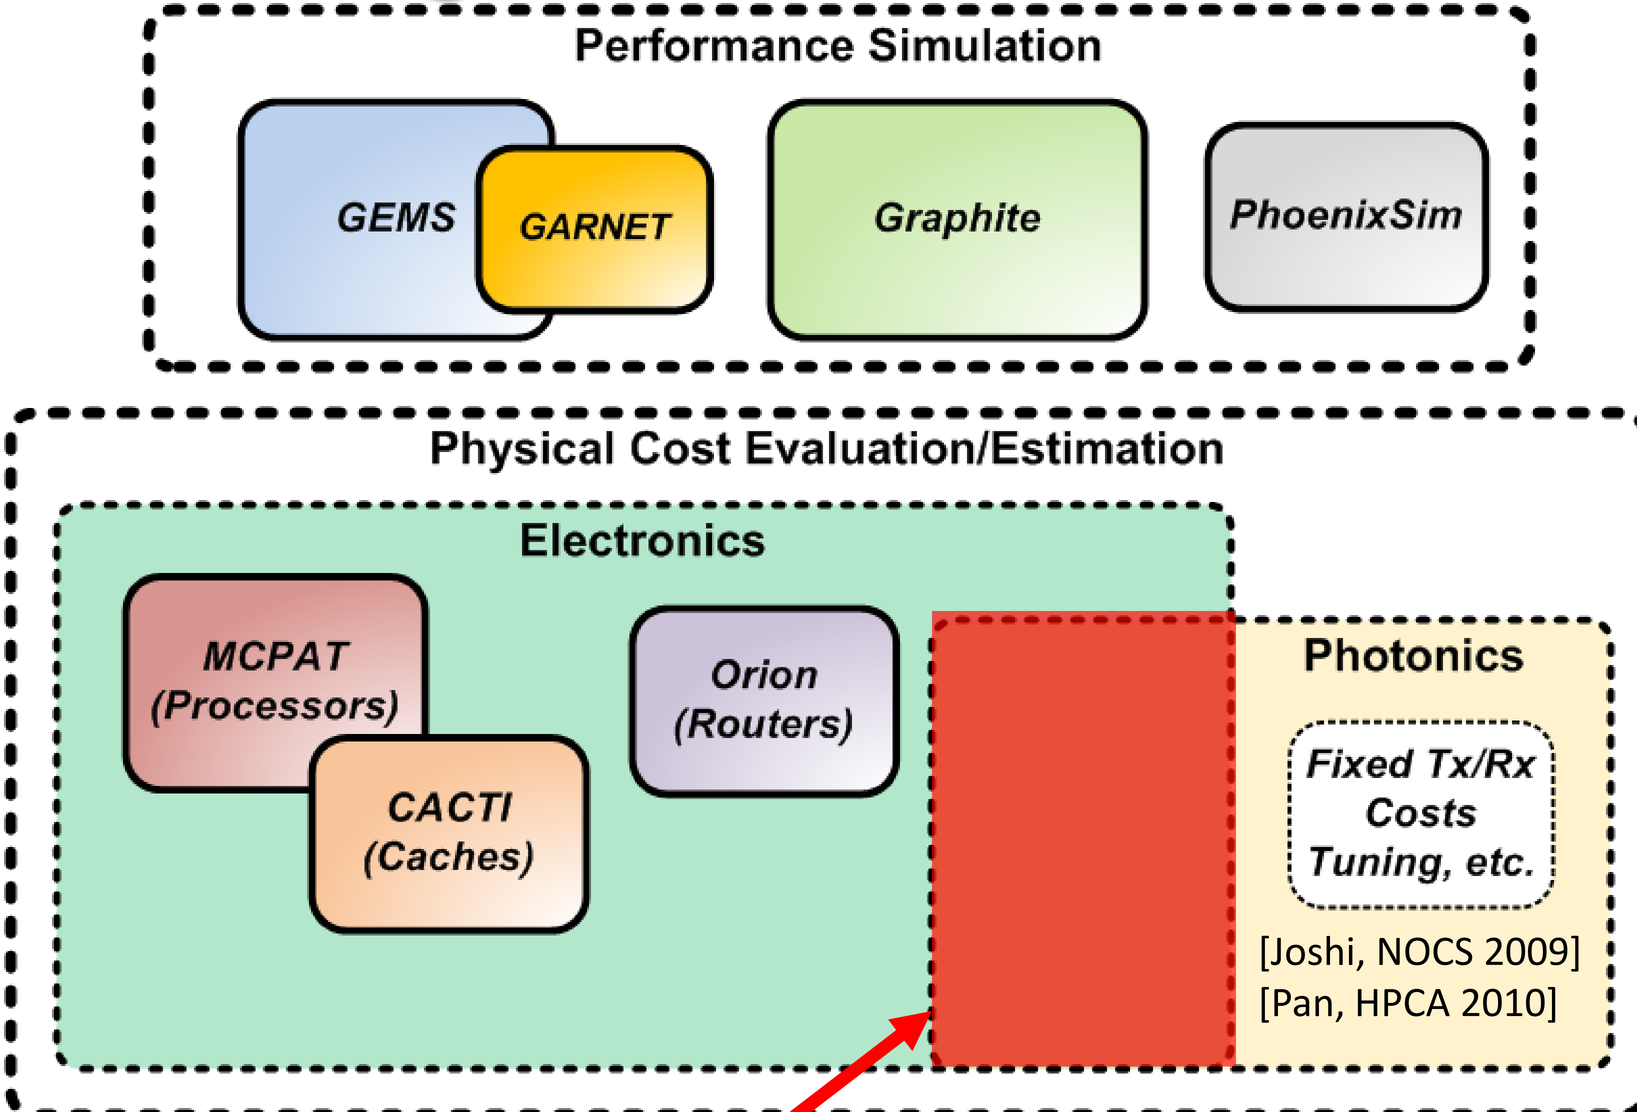


**Nothing currently models the interface between electronics and photonics**

# Why Not Just Photonics?

# Why Not Just Photonics?

- Original plan for DSENT, but…

# Why Not Just Photonics?

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

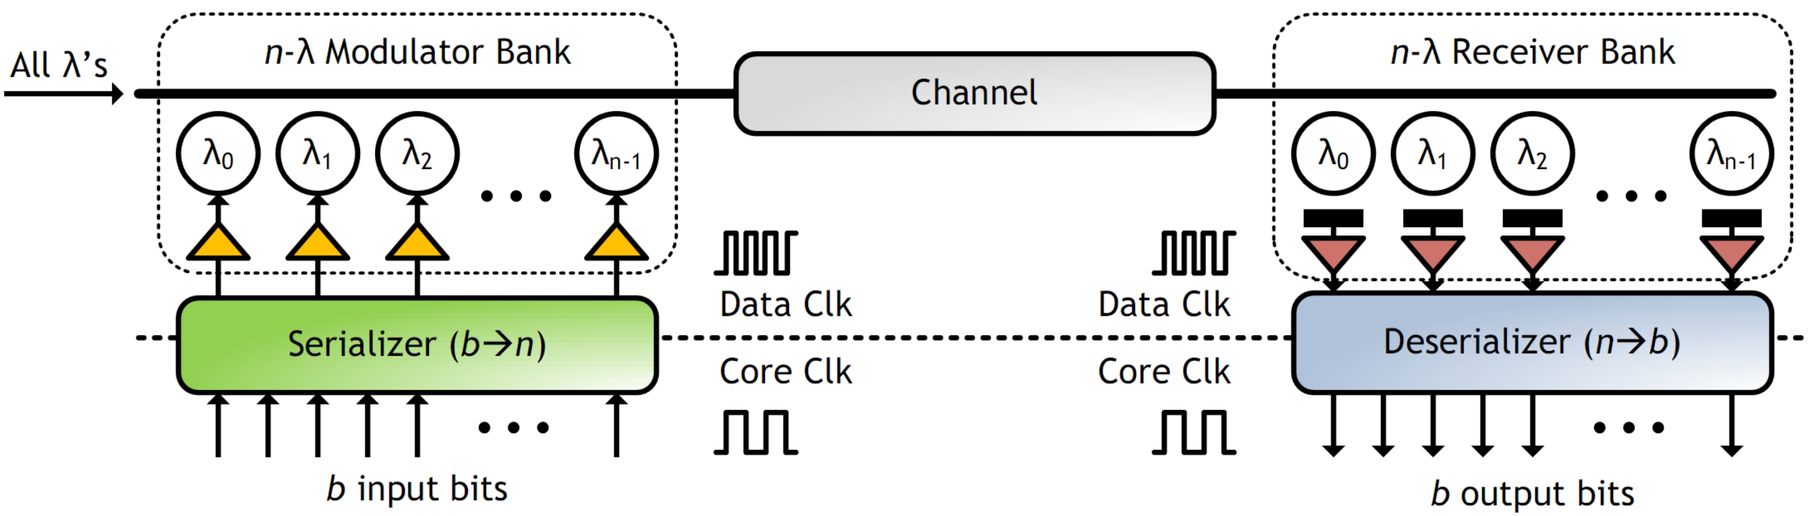

# Why Not Just Photonics?

- 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 fairly with photonics…

# Orion 2.0 Issues

# Orion 2.0 Issues

- ✗ Scaling factors no longer valid for advanced processes

# Orion 2.0 Issues

- Scaling factors no longer valid for advanced processes
- Very difficult to add technology or extend existing models

# Orion 2.0 Issues

- Scaling factors no longer valid for advanced processes
- Very difficult to add technology or extend existing models
- Incomplete architectural models and timing for the router

# Orion 2.0 Issues

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

# Orion 2.0 Issues

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

# Orion 2.0 Issues

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

# DSENT

## Design Space Exploration of Networks Tool

# DSENT

Design Space Exploration of Networks Tool

- Overview

# DSENT

Design Space Exploration of Networks Tool

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

# DSENT

Design Space Exploration of Networks Tool

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

# DSENT

Design Space Exploration of Networks Tool

- 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

**Technology File**

```
# This file contains the model for SOI 45nm LVT
Name = SOI45LVT

# Supply voltage used across the circuit (V)
Vdd = 1.0
# Temperature (K)
Temperature = 340

# ==============================================
# Parameters for transistors
# ==============================================

# Gate length (m)
Gate->Length = 0.040e-6
# Contacted gate pitch (m)
Gate->PitchContacted = 0.20e-6
# Non-contacted gate pitch (m)
Gate->PitchNonContacted = 0.20e-6
```

**Network Parameter File**

```
# Clos Parameters
ClockFrequency                = 1e9
NumberInputSites              = 64
NumberOutputSites             = 64
NumberBitsPerFlit             = 64
NumberIngressRouters          = 8
NumberMiddleRouters           = 8
NumberEgressRouters           = 8


# Router-specific parameters
Router->NumberVirtualNetworks                      = 3
Router->NumberVirtualChannelsPerVirtualNetwork     = [1,1,1]
Router->NumberBuffersPerVirtualChannel             = [4,1,1]
Router->InputPort->BufferModel                     = DFFRAM
Router->CrossbarModel                              = MultiplexerCrossbar
Router->SwitchAllocator->ArbiterModel              = MatrixArbiter
Router->ClockTreeModel                             = BroadcastHTree
Router->ClockTree->NumberLevels                    = 6
Router->ClockTree->WireLayer                       = Intermediate
Router->ClockTree->WireWidthMultiplier             = 1.0
```

# Two Ways to Use DSENT

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

**Technology File**

```
# This file contains the model for SOI 45nm LVT
Name = SOI45LVT

# Supply voltage used across the circuit (V)
Vdd = 1.0
# Temperature (K)
Temperature = 340

# ==============================================
# Parameters for transistors
# ==============================================

# Gate length (m)
Gate->Length = 0.040e-6
# Contacted gate pitch (m)
Gate->PitchContacted = 0.20e-6
# Non-contacted gate pitch (m)
Gate->PitchNonContacted = 0.20e-6
```

**Network Parameter File**

```
# Clos Parameters
ClockFrequency                 = 1e9
NumberInputSites               = 64
NumberOutputSites              = 64
NumberBitsPerFlit              = 64
NumberIngressRouters           = 8
NumberMiddleRouters            = 8
NumberEgressRouters            = 8


# Router-specific parameters
Router->NumberVirtualNetworks                          = 3
Router->NumberVirtualChannelsPerVirtualNetwork         = [1,1,1]
Router->NumberBuffersPerVirtualChannel                 = [4,1,1]
Router->InputPort->BufferModel                         = DFFRAM
Router->CrossbarModel                                  = MultiplexerCrossbar
Router->SwitchAllocator->ArbiterModel                  = MatrixArbiter
Router->ClockTreeModel                                 = BroadcastHTree
Router->ClockTree->NumberLevels                        = 6
Router->ClockTree->WireLayer                           = Intermediate
Router->ClockTree->WireWidthMultiplier                 = 1.0
```

**Run DSENT**

```
$ ./dsent -cfg configs/dsent.cfg.pclos
```

**Results**

```
NddPower>>PhotonicClos:RingTuning = 1.01789 (1.01789 * 1)
NddPower>>PhotonicClos:Laser = 0.212644 (0.212644 * 1)
NddPower>>PhotonicClos:Leakage = 0.385034 (0.385034 * 1)
Area>>PhotonicClos:Active = 1.38268e-06 (1.38268e-06 * 1)
```

# Two Ways to Use DSENT

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

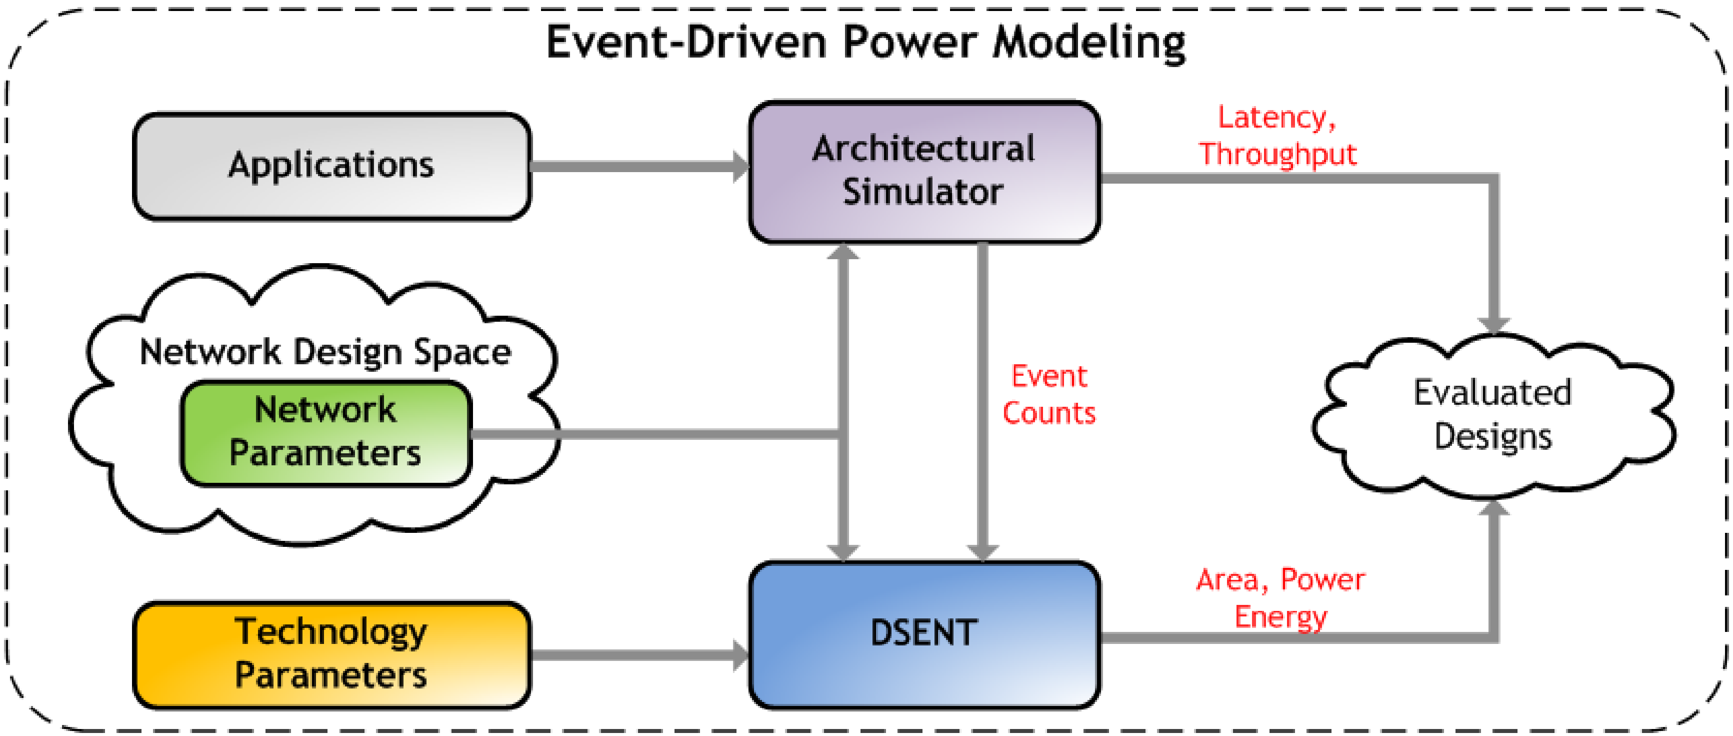

# DSENT

## Design Space Exploration of Networks Tool

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

# Electrical Model

- ASIC-like modeling flow, generates primitives/standard cells

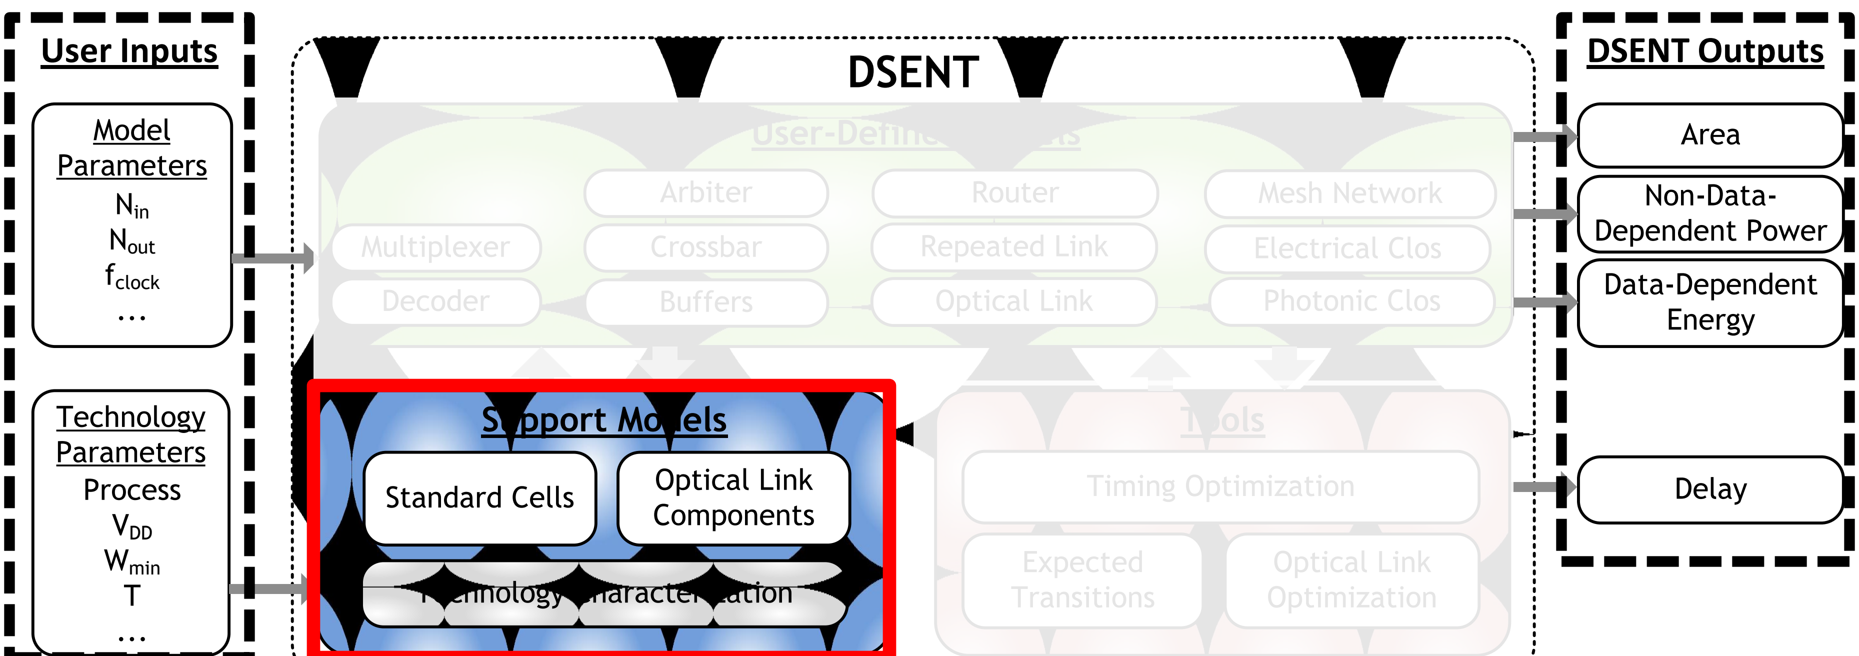

# Electrical Model

- ASIC-like modeling flow, generates primitives/standard cells
- Keep relevant tech parameters, simplify technology entry

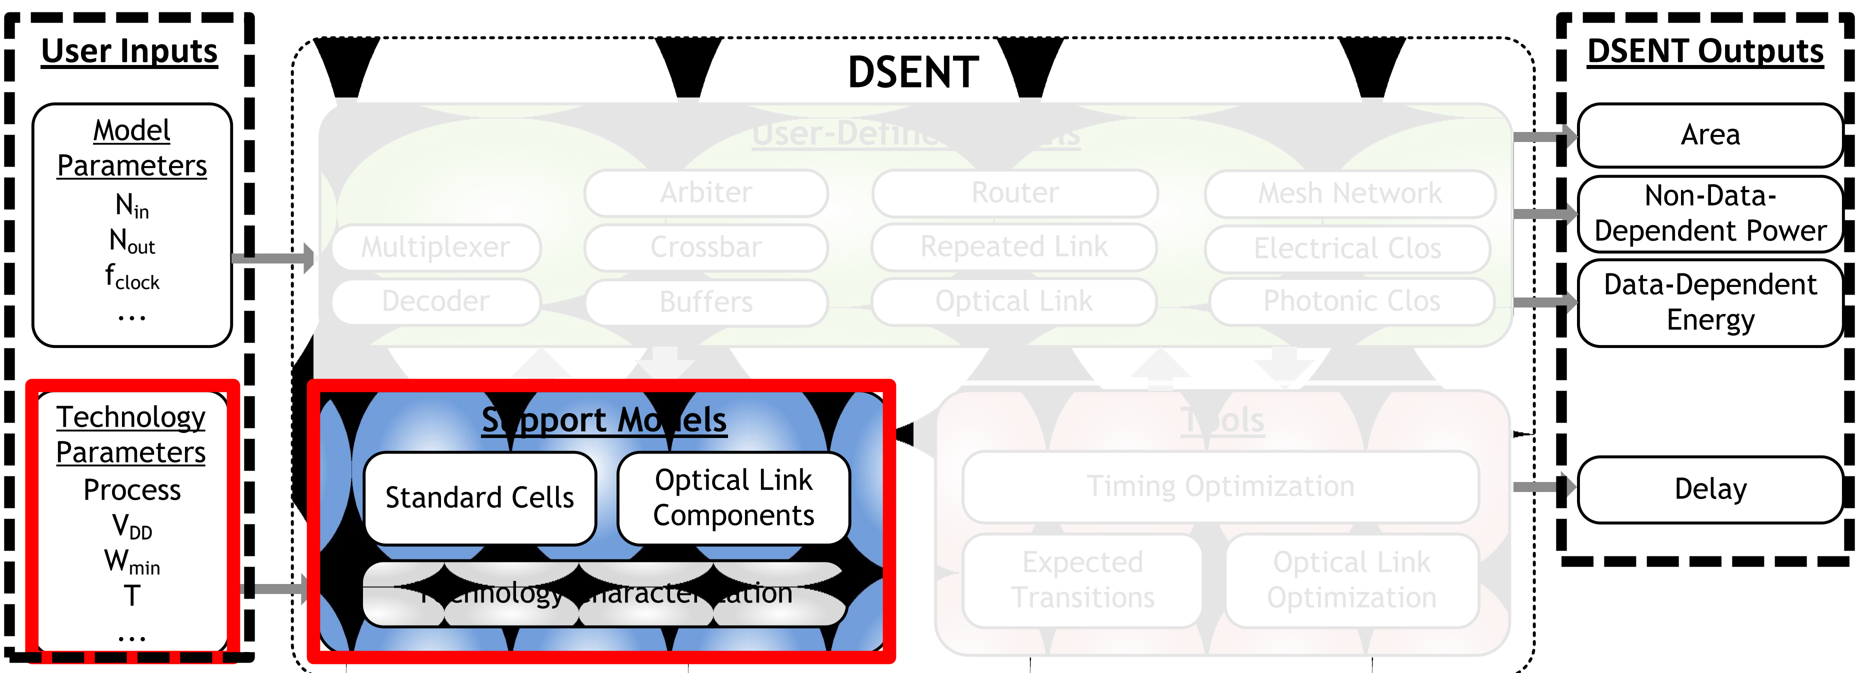

# Electrical Model

- ASIC-like modeling flow, generates primitives/standard cells
- Keep relevant tech parameters, simplify technology entry
- Delay model, timing-constrained cell sizing, electrical links

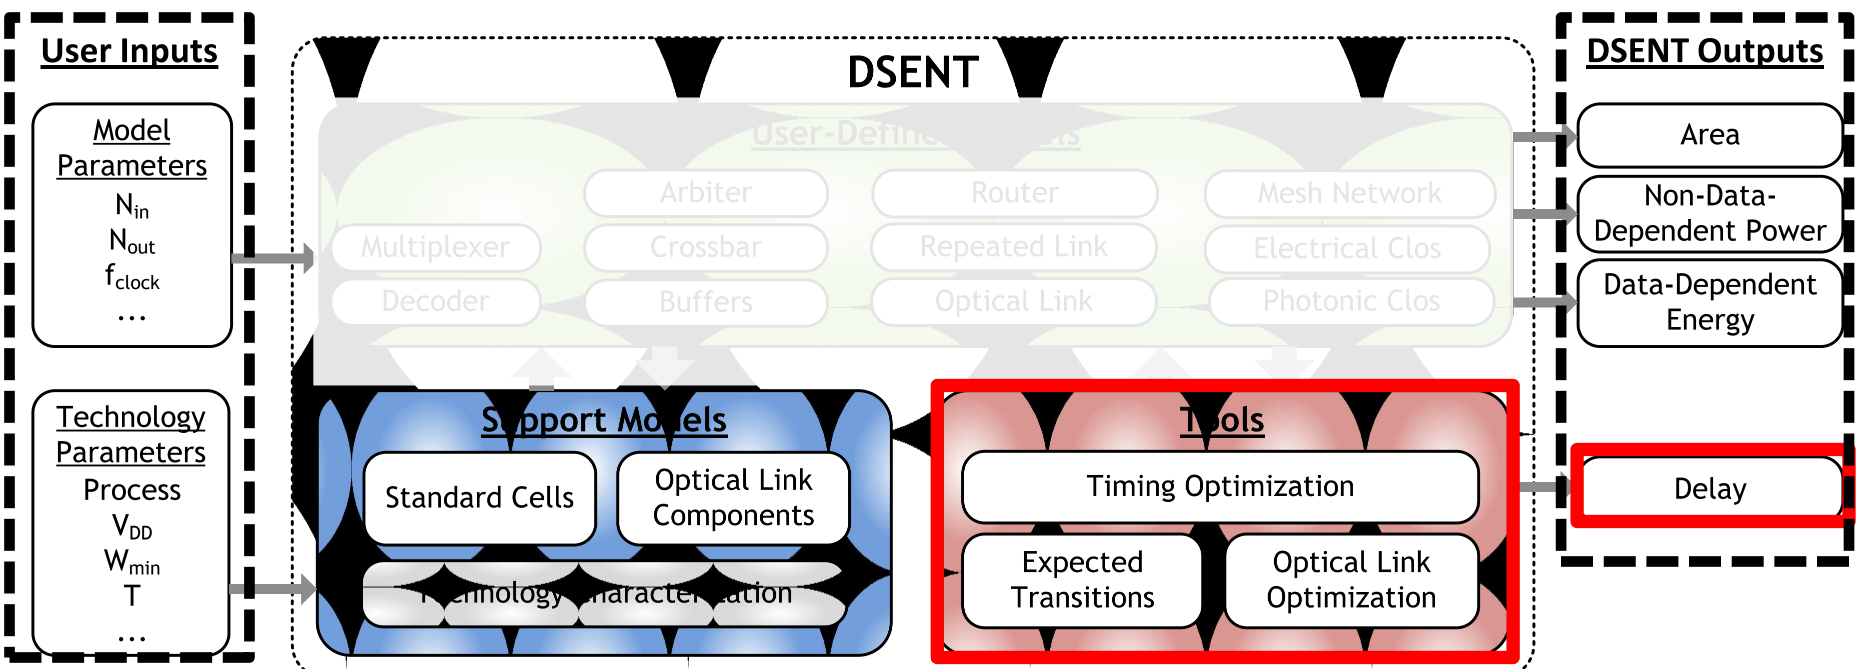

# Electrical Model

- ASIC-like modeling flow, generates primitives/standard cells
- Keep relevant tech parameters, simplify technology entry
- Delay model, timing-constrained cell sizing, electrical links
- Able to model more generic digital, beyond just routers

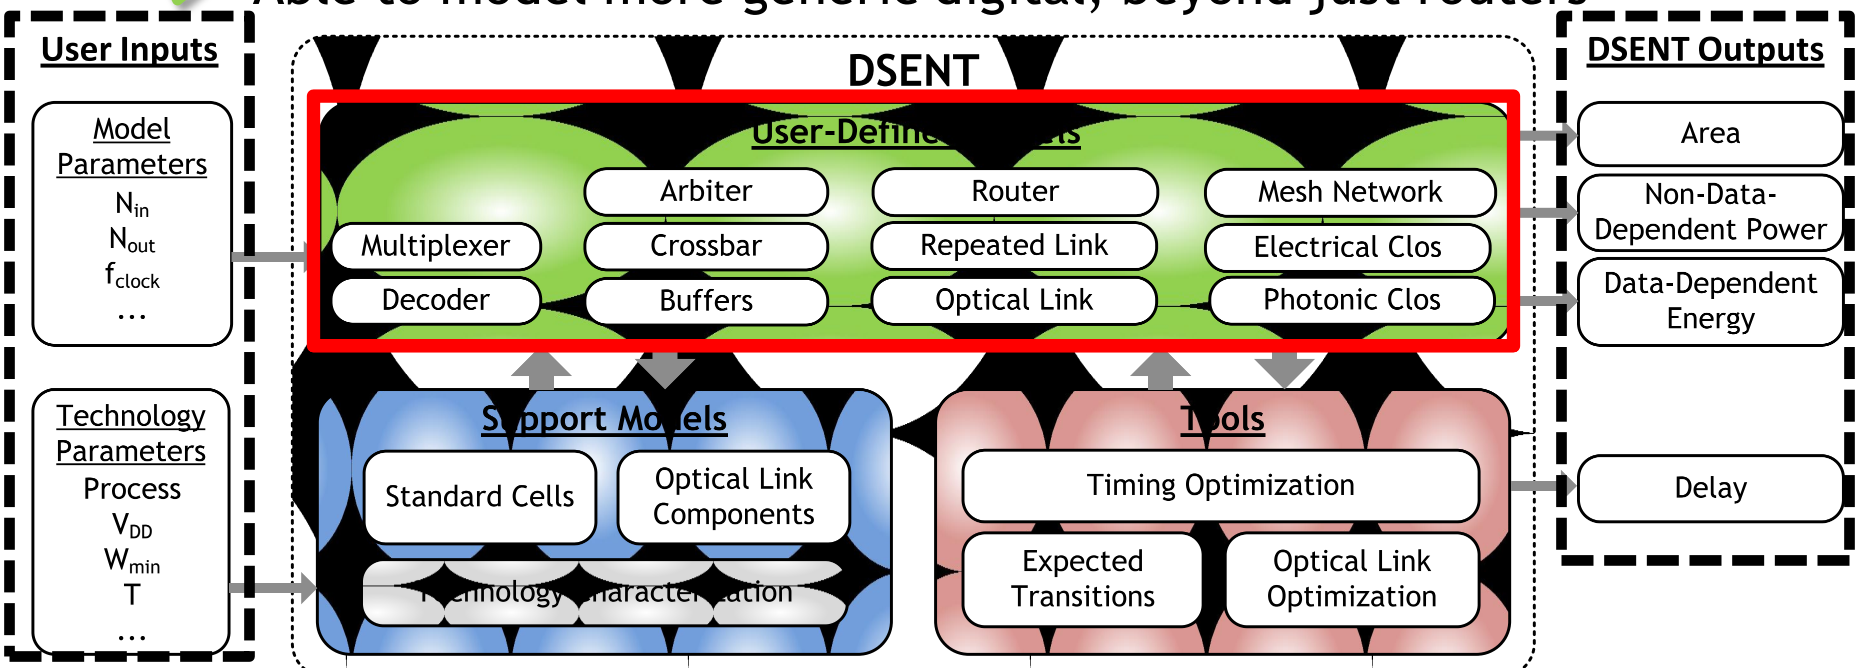

# Electrical Model

- 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

# Electrical Model

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

# Photonics Model

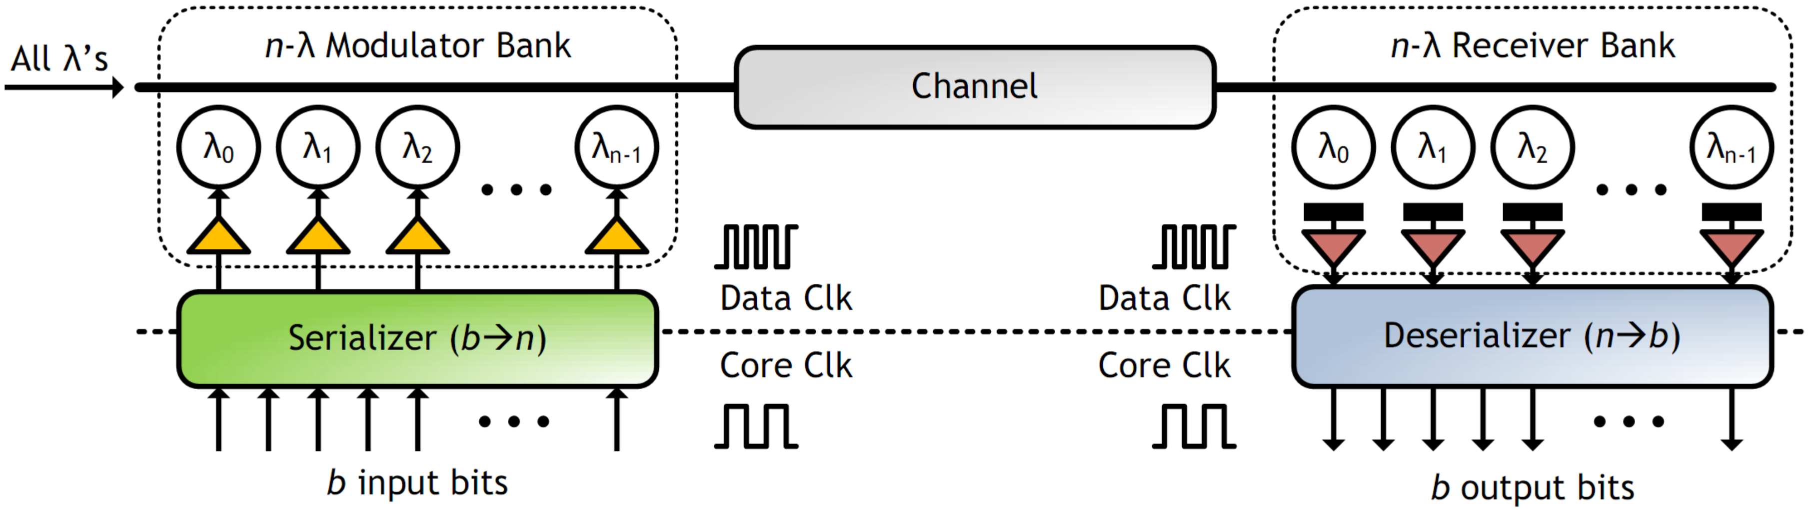


- Four different sources of power consumption
  - Modulator, receivers
  - Laser power
  - Thermal tuning
  - Serialize, deserialize backends

# Photonics Model

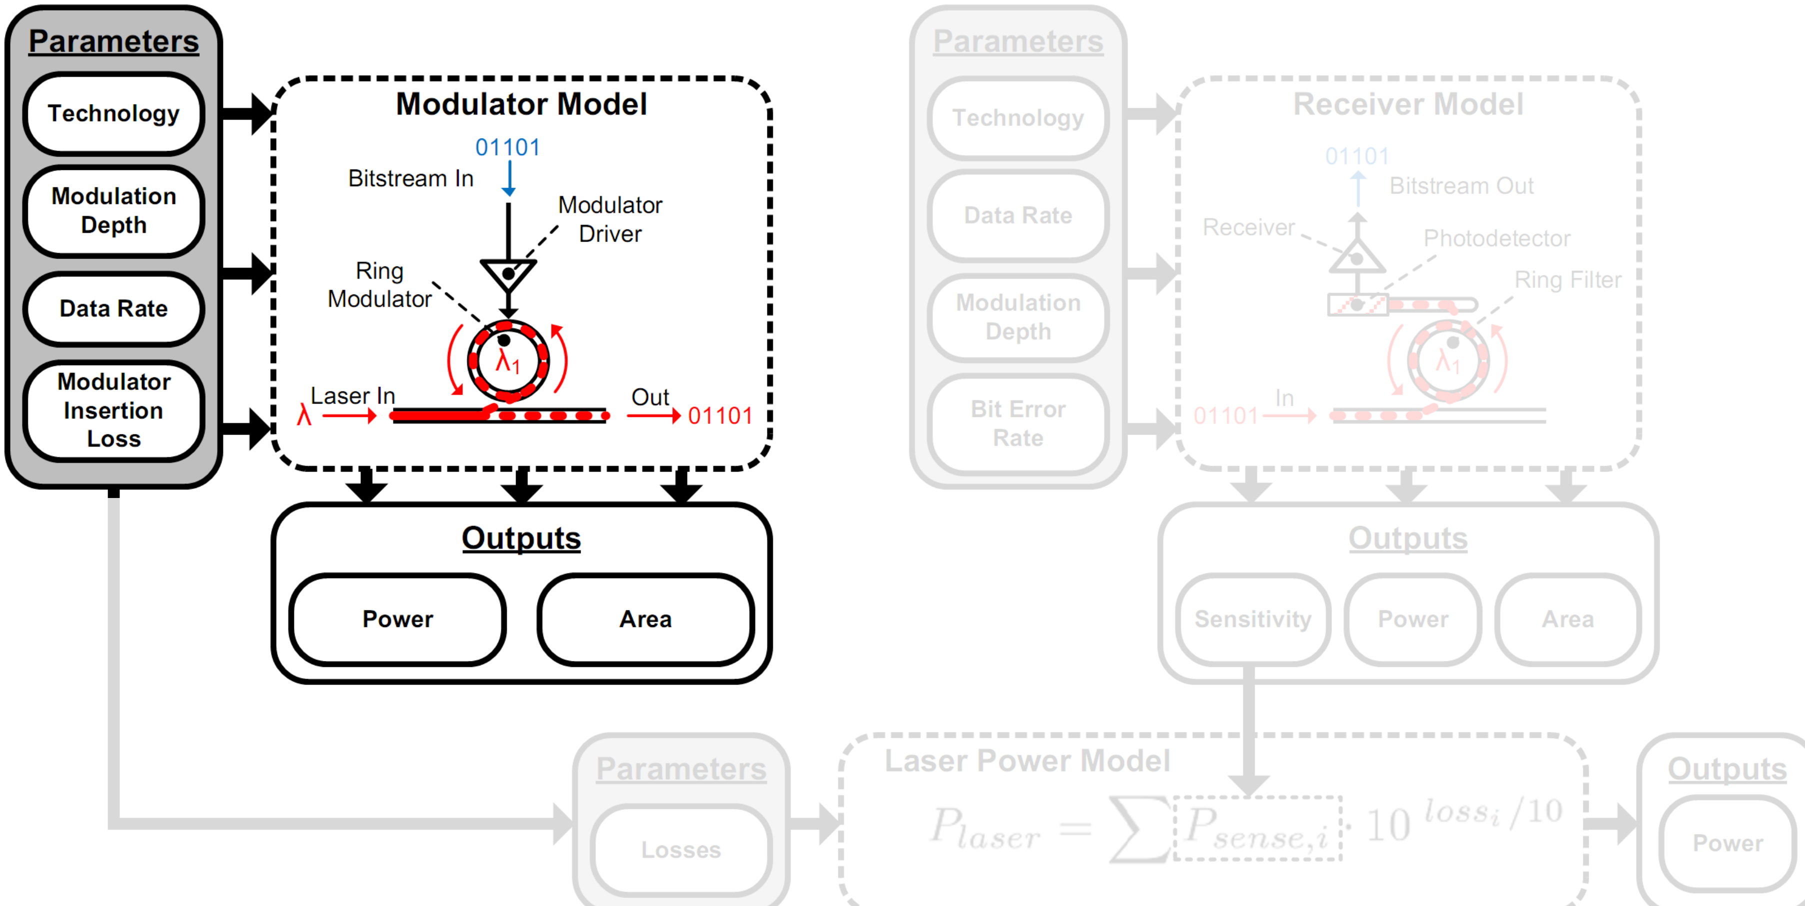


- Modulator becomes more expensive with:
  - High data-rate
  - Higher modulation depth (extinction ratio)
  - Lower insertion loss

# Photonics Model

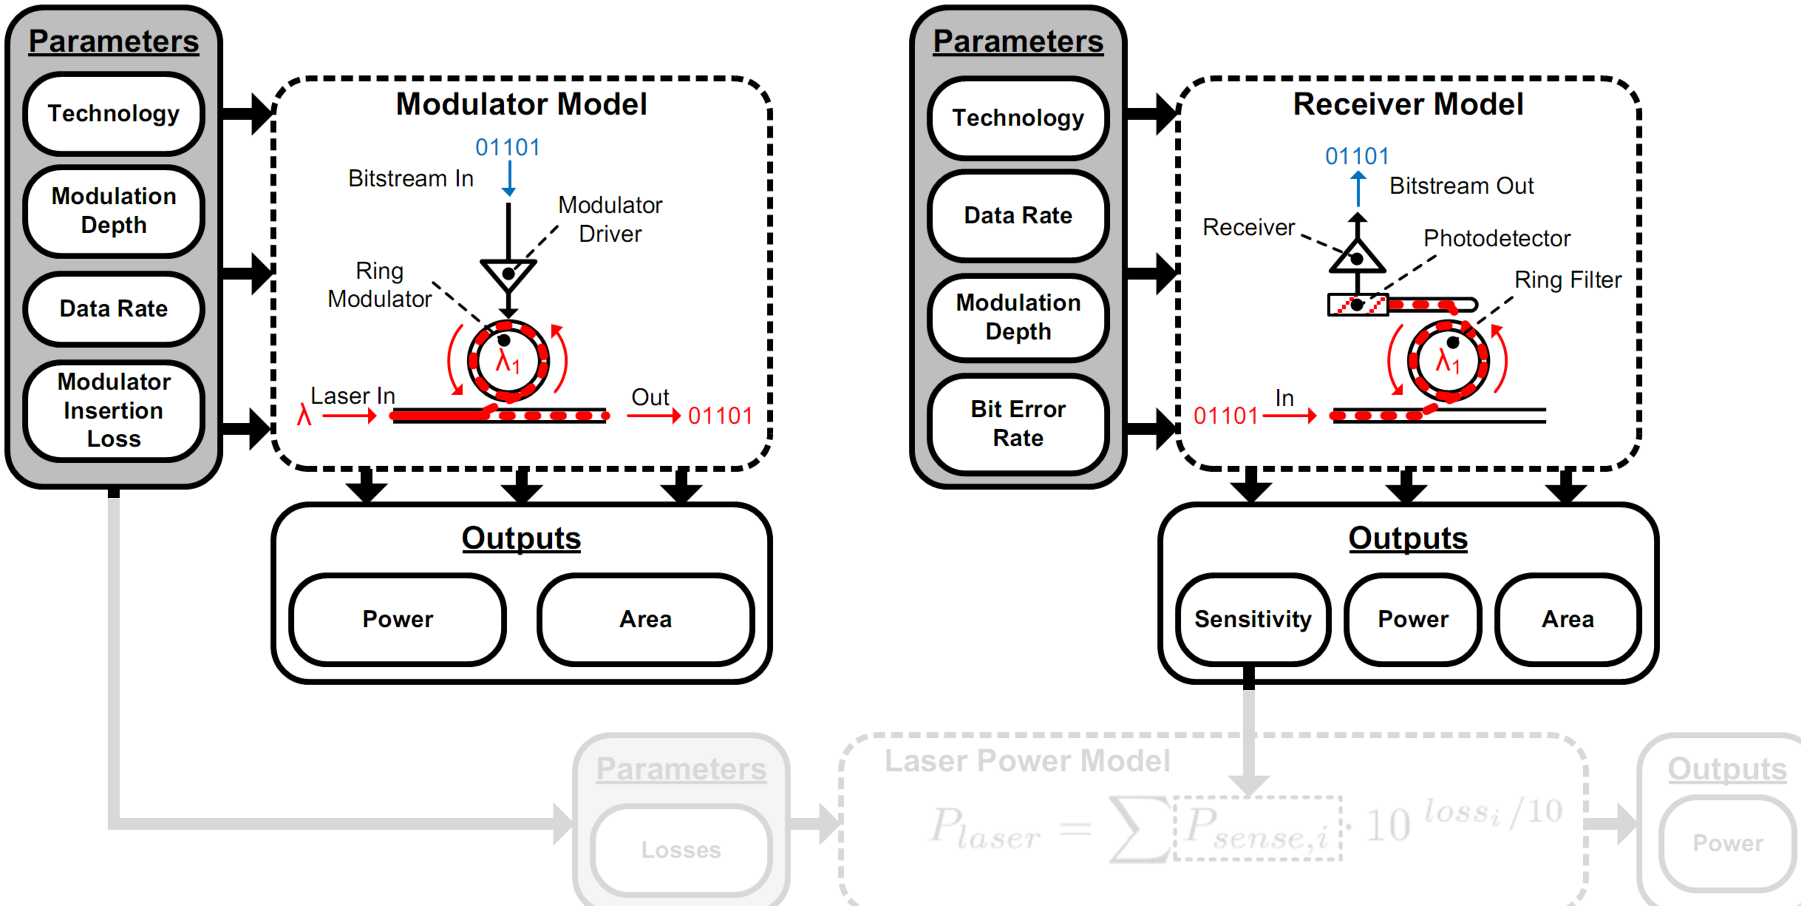


- Receiver becomes more expensive with:
  - High data-rate
- Receiver sensitivity degrades with:
  - High data-rate
  - Lower modulation depth
  - Higher bit error rate requirement

# Photonics Model

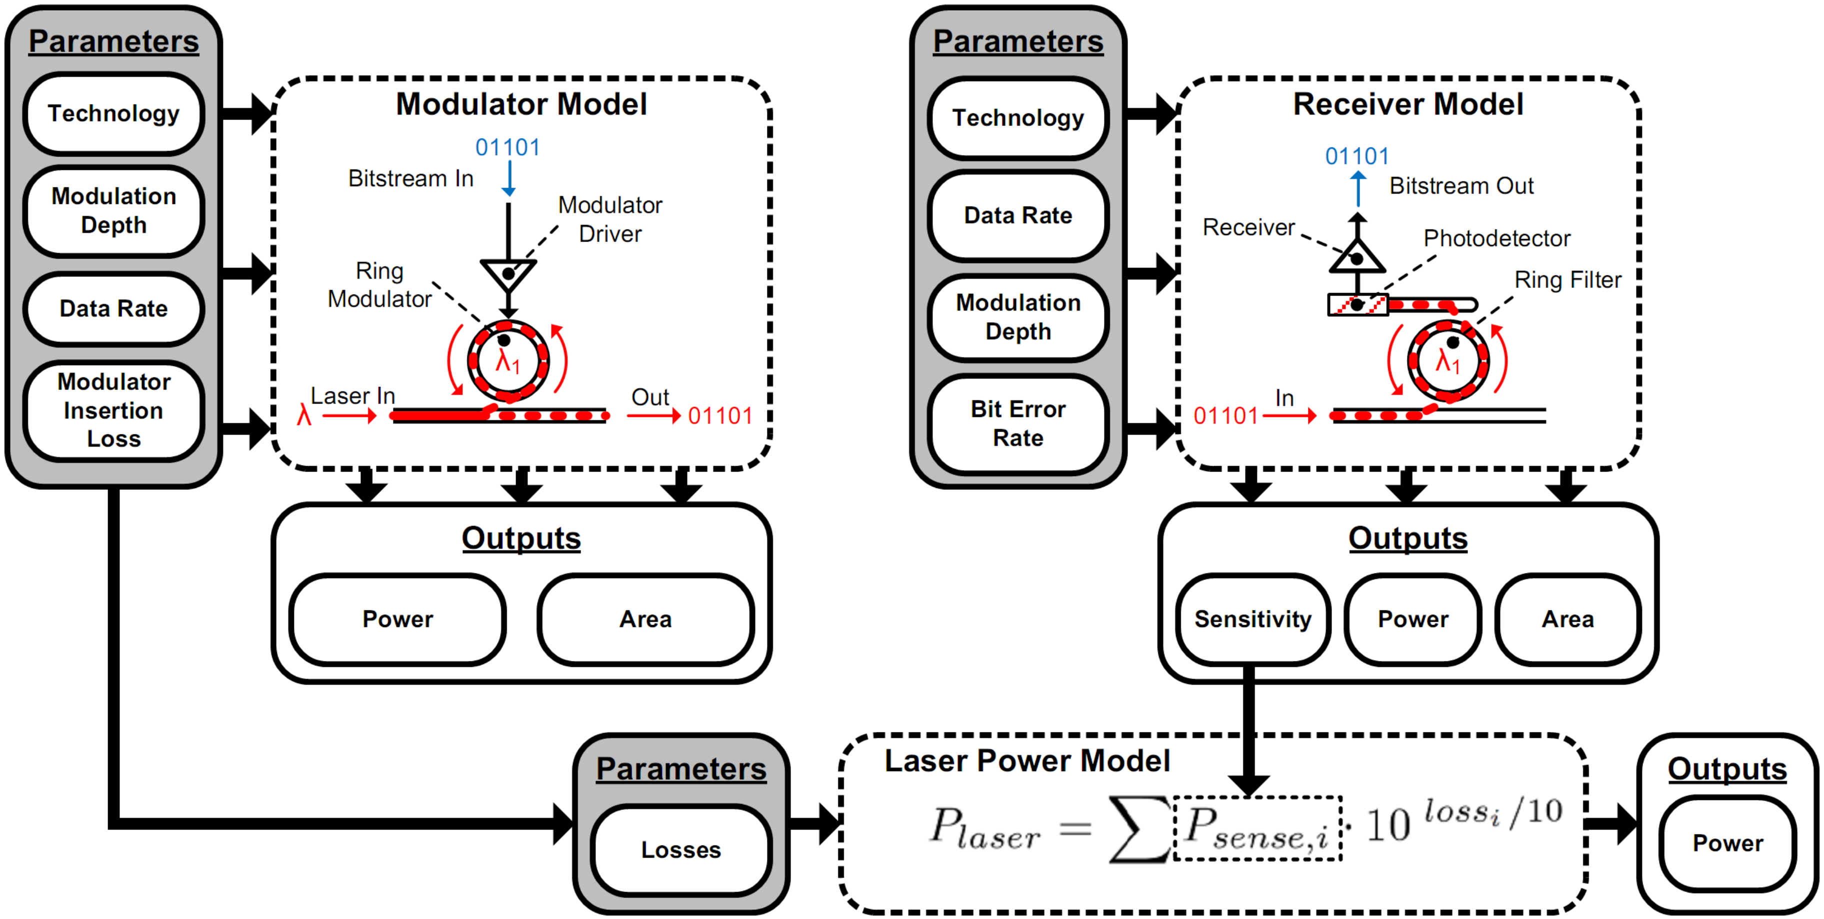


- Laser power requirement gets worse with:
  - Higher receiver sensitivity requirement
  - Higher channel losses, e.g. higher modulator insertion loss

# Photonics Model

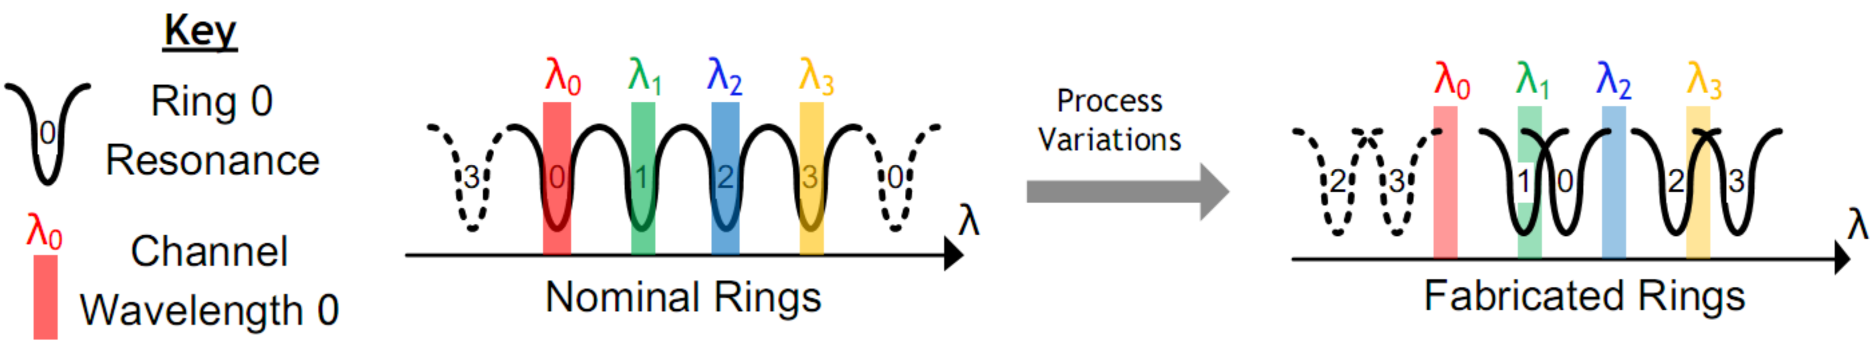


- Ring resonator devices are sensitive to process, temperature, active tuning is needed

# Photonics Model

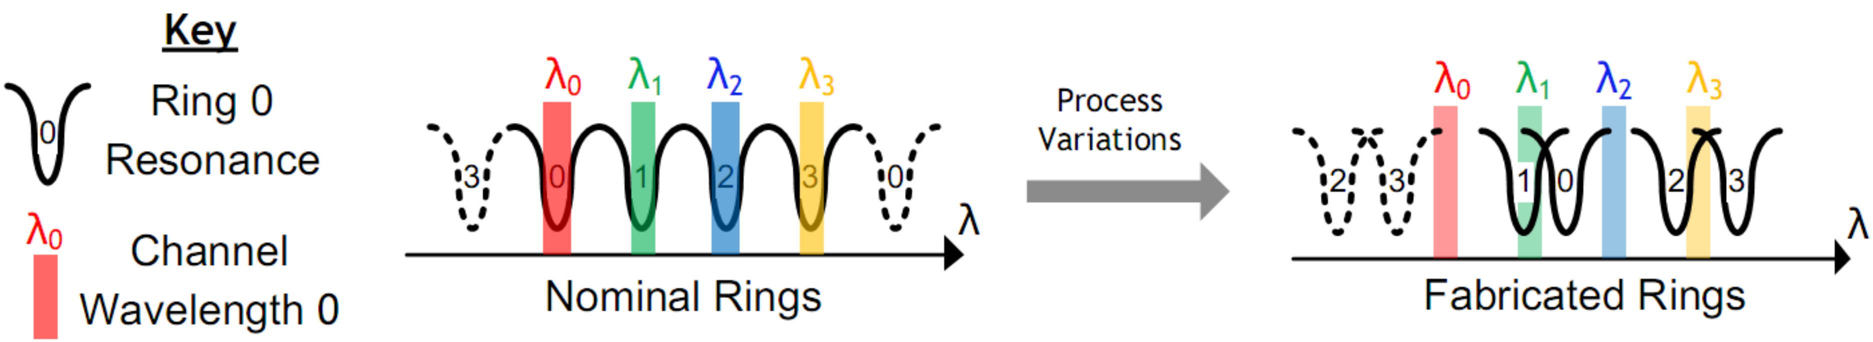


- Ring resonator devices are sensitive to process, temperature, active tuning is needed
- Not necessarily a fixed cost per ring!
  - [Georgas CICC 2011, Nitta HPCA 2011]

# Photonics Model

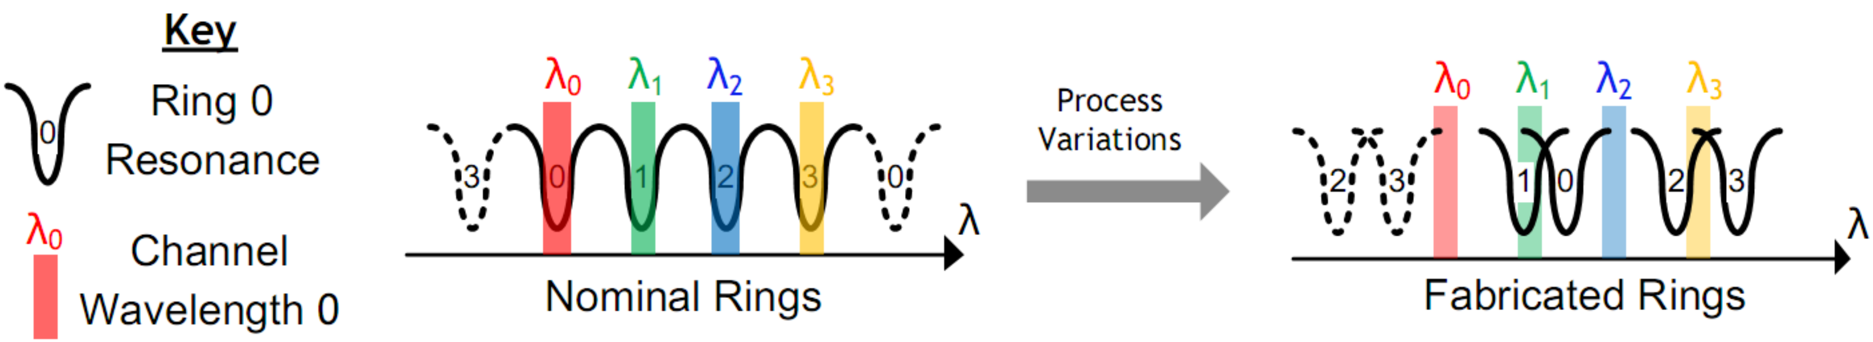


- Ring resonator devices are sensitive to process, temperature, active tuning is needed
- Not necessarily a fixed cost per ring!
  - [Georgas CICC 2011, Nitta HPCA 2011]

✓ DSENT models schemes for tuning, impact of process sigmas

# Photonics Model

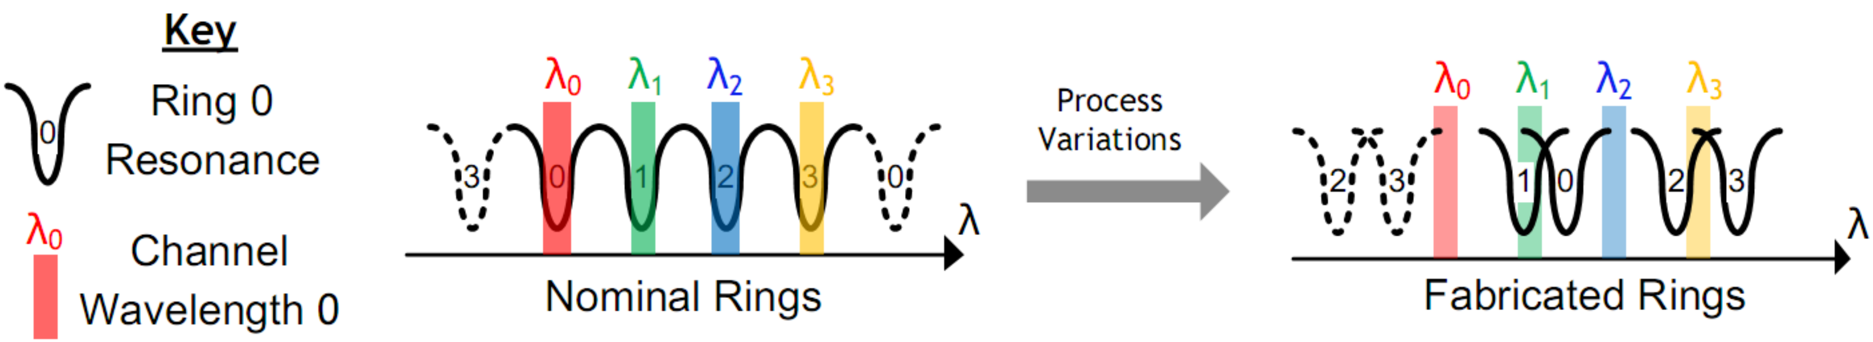


- Ring resonator devices are sensitive to process, temperature, active tuning is needed
- Not necessarily a fixed cost per ring!
  - [Georgas CICC 2011, Nitta HPCA 2011]

✔ 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

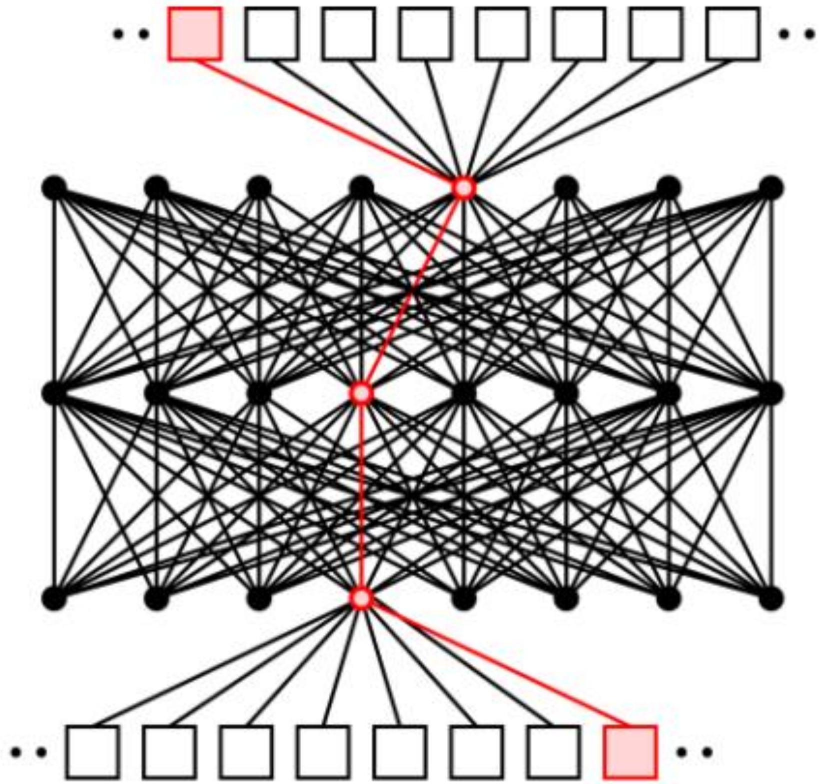

[Joshi, NOCS 2009]

- 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)

# 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

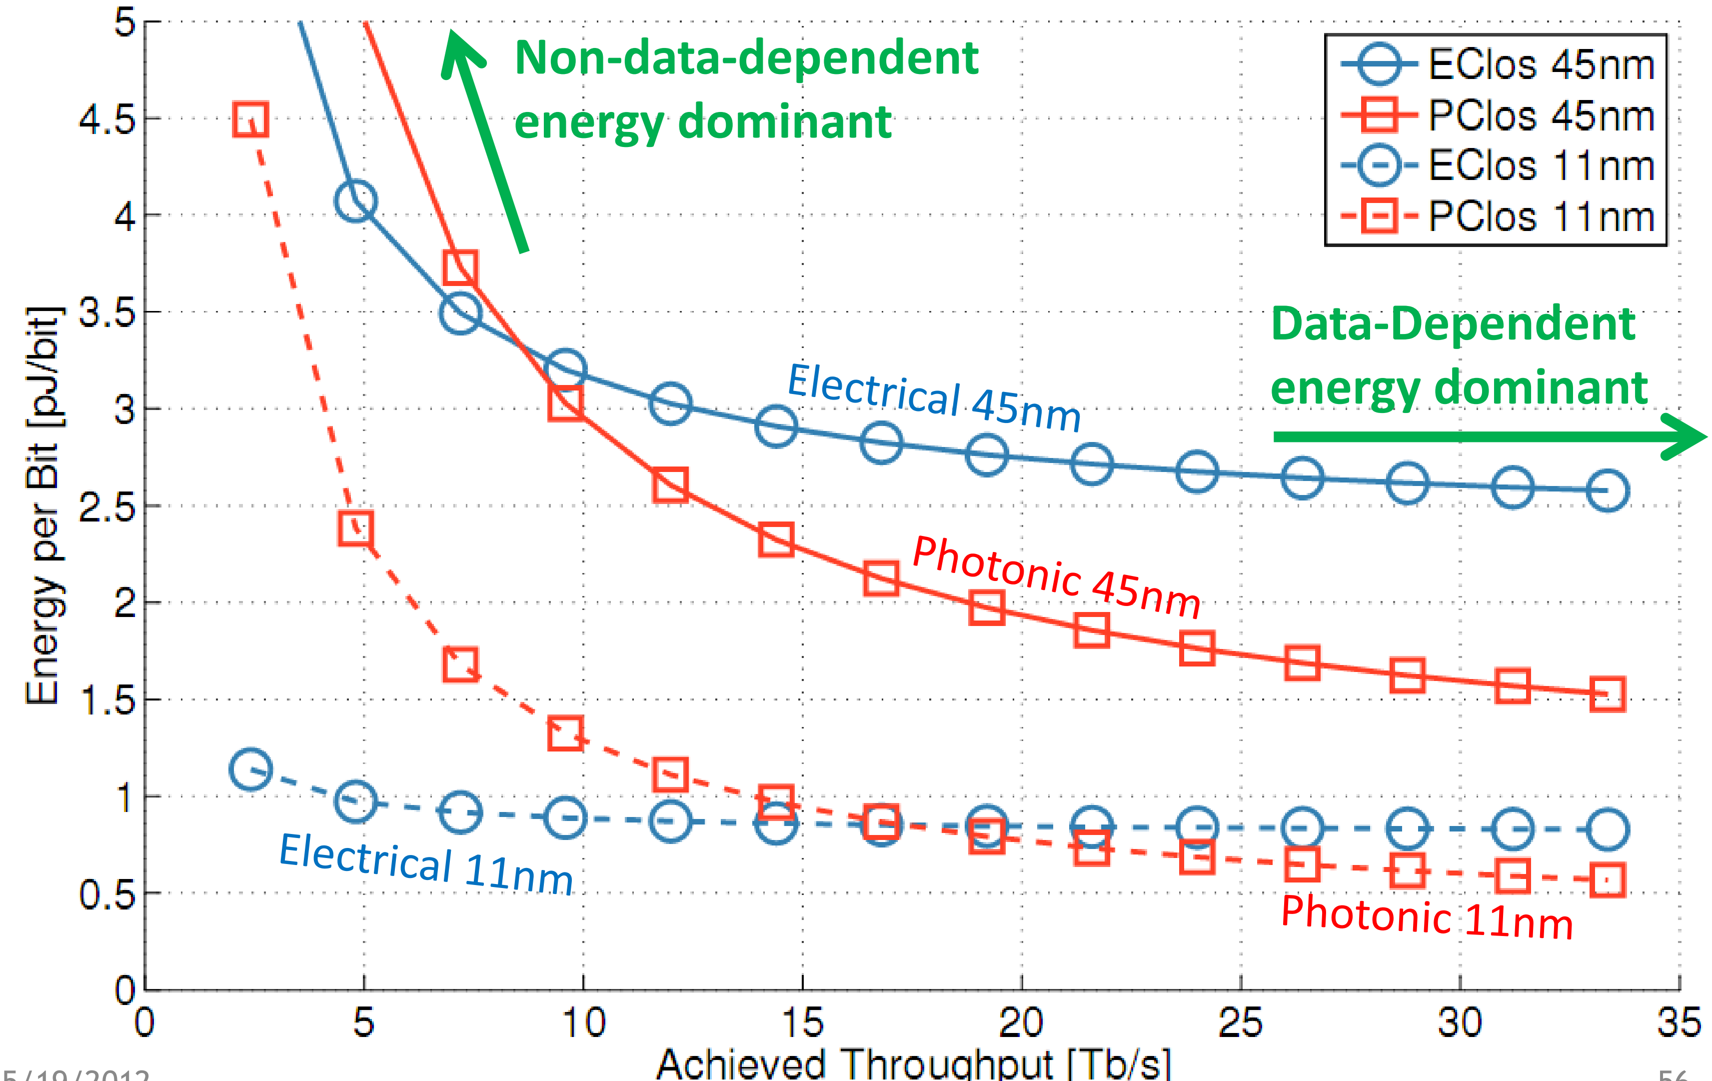

# Effect of Utilization

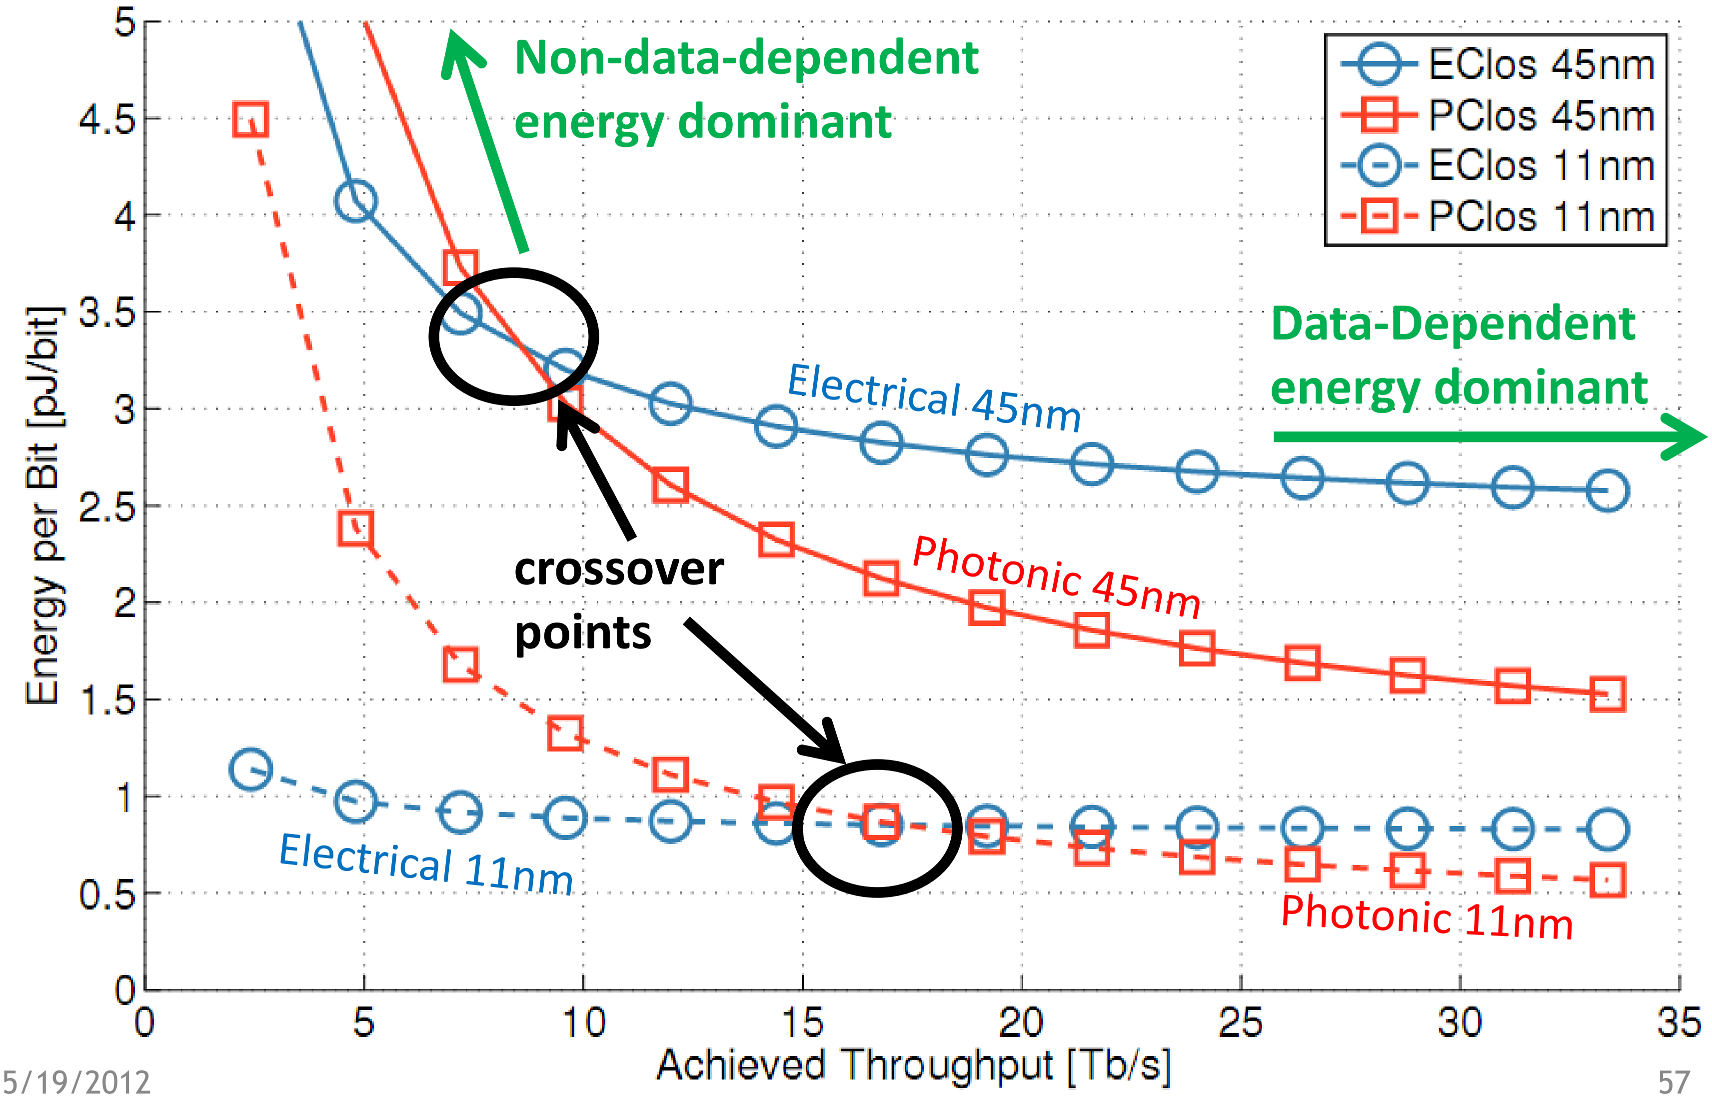

# Effect of Utilization

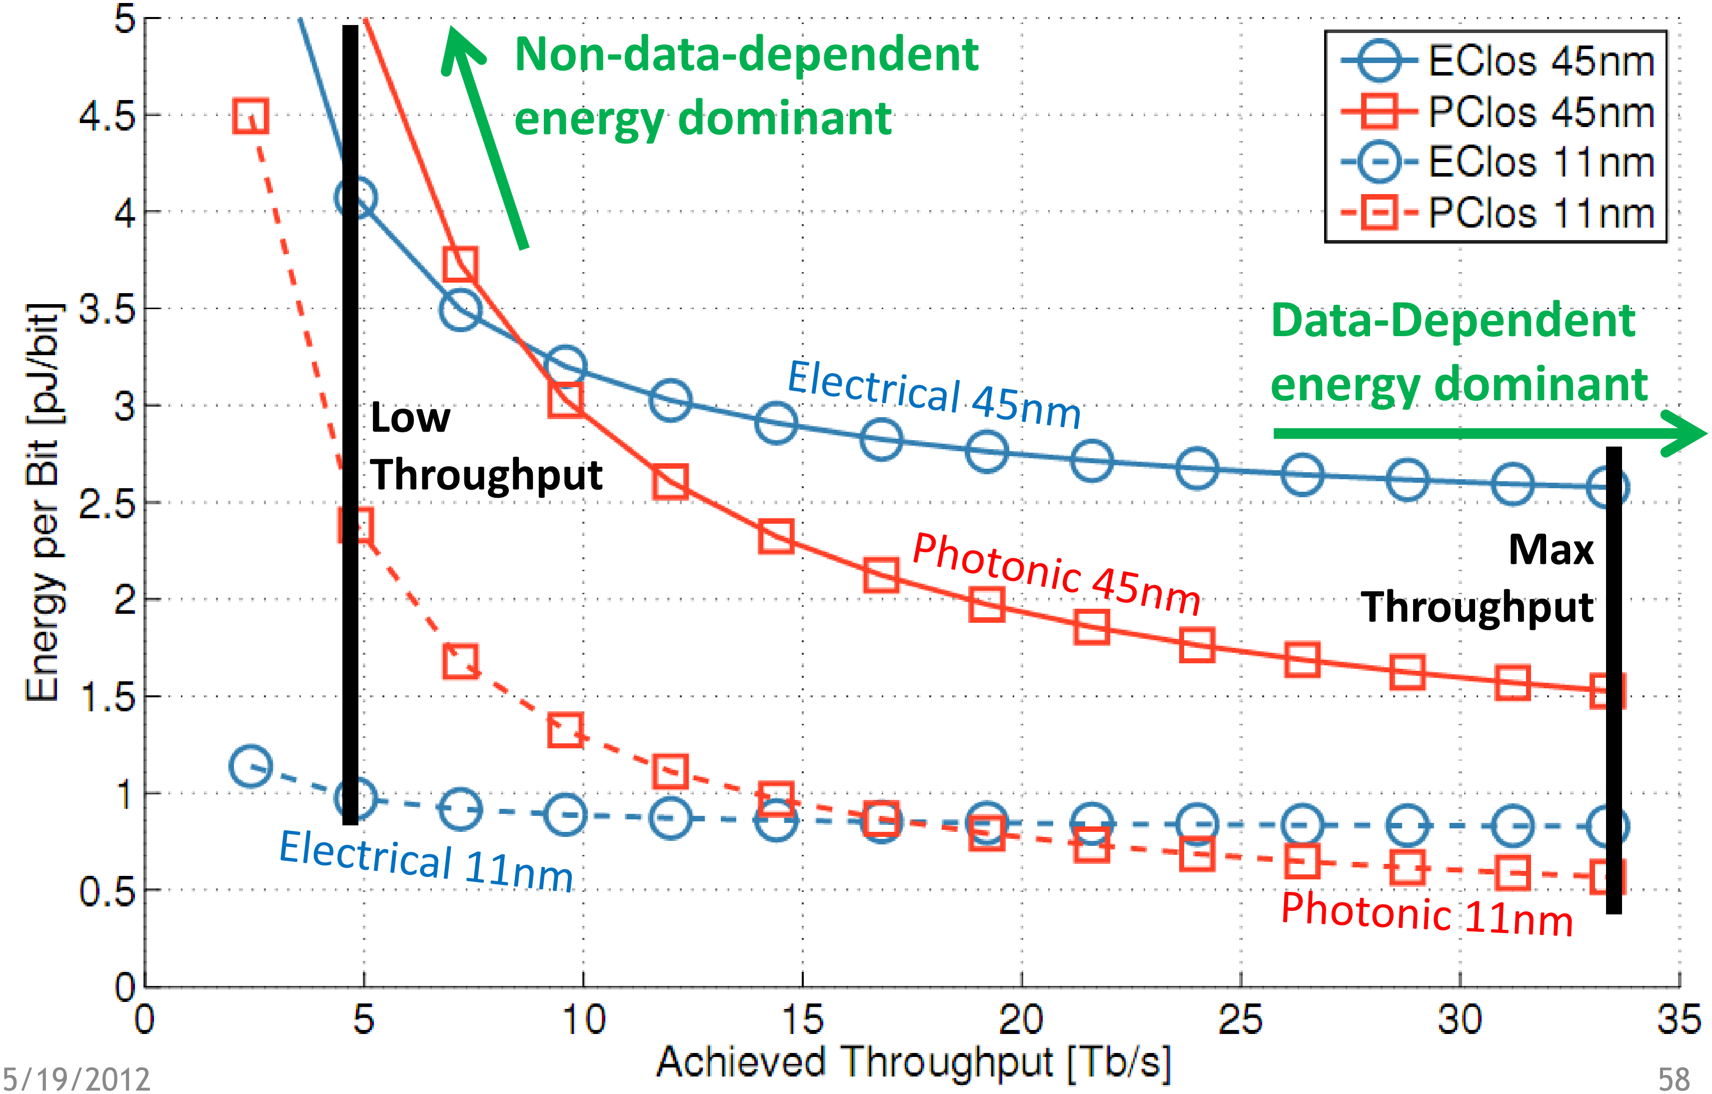

# Energy per Bit Breakdown

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

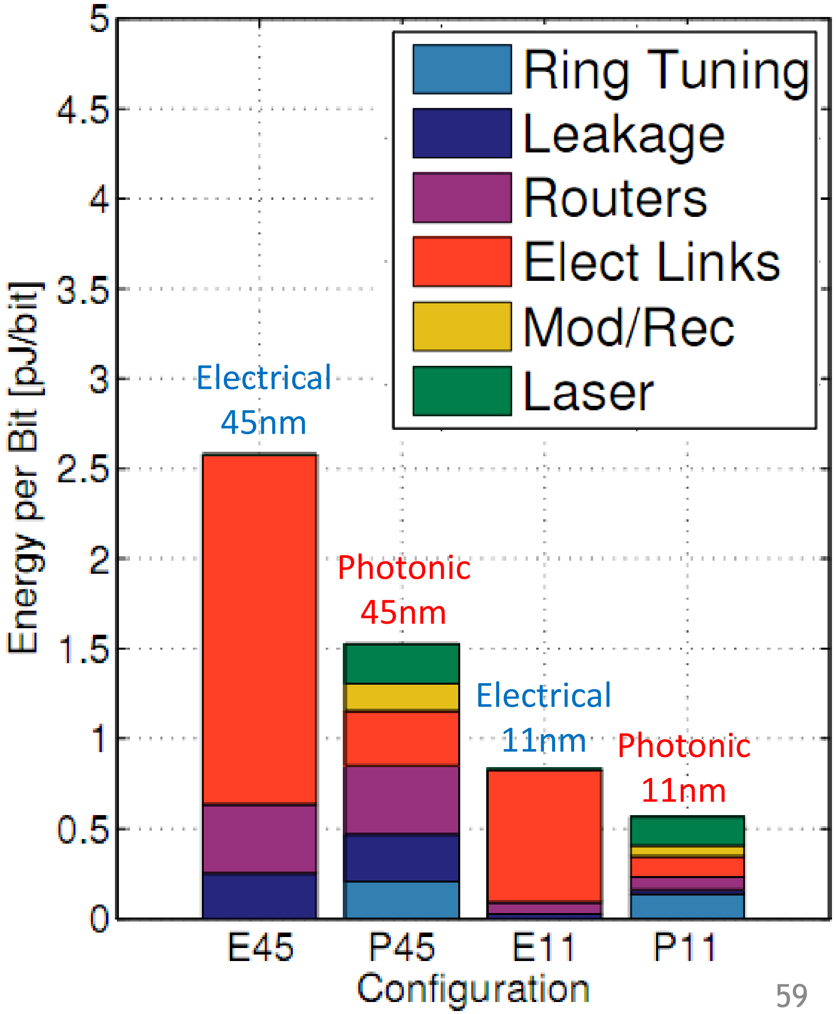

# Energy per Bit Breakdown

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

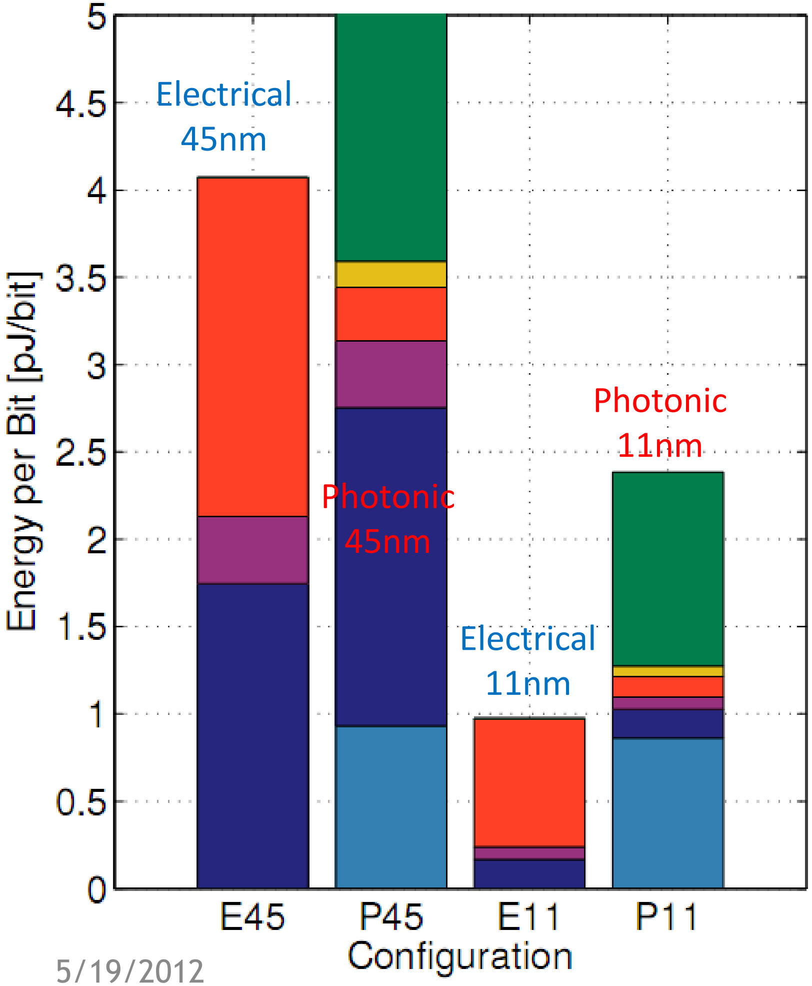


## Energy Breakdown at Max Network Throughput (33 Tb/s)

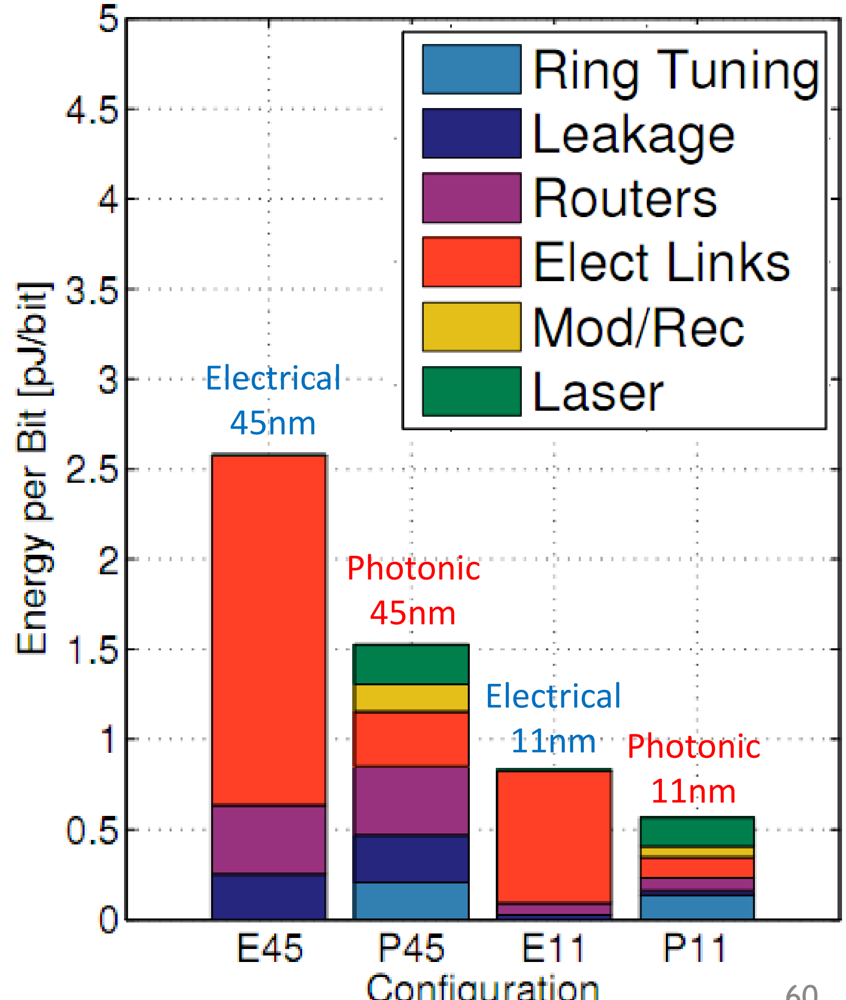

# Energy per Bit Breakdown

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

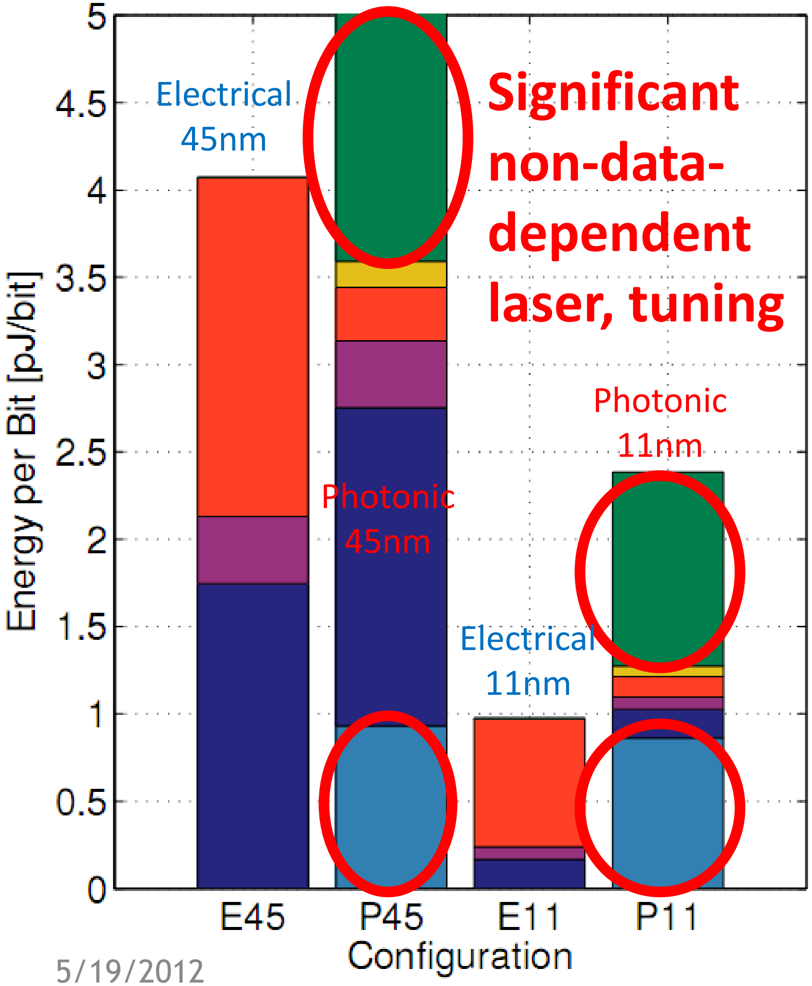


**Energy Breakdown at Max Network Throughput (33 Tb/s)**

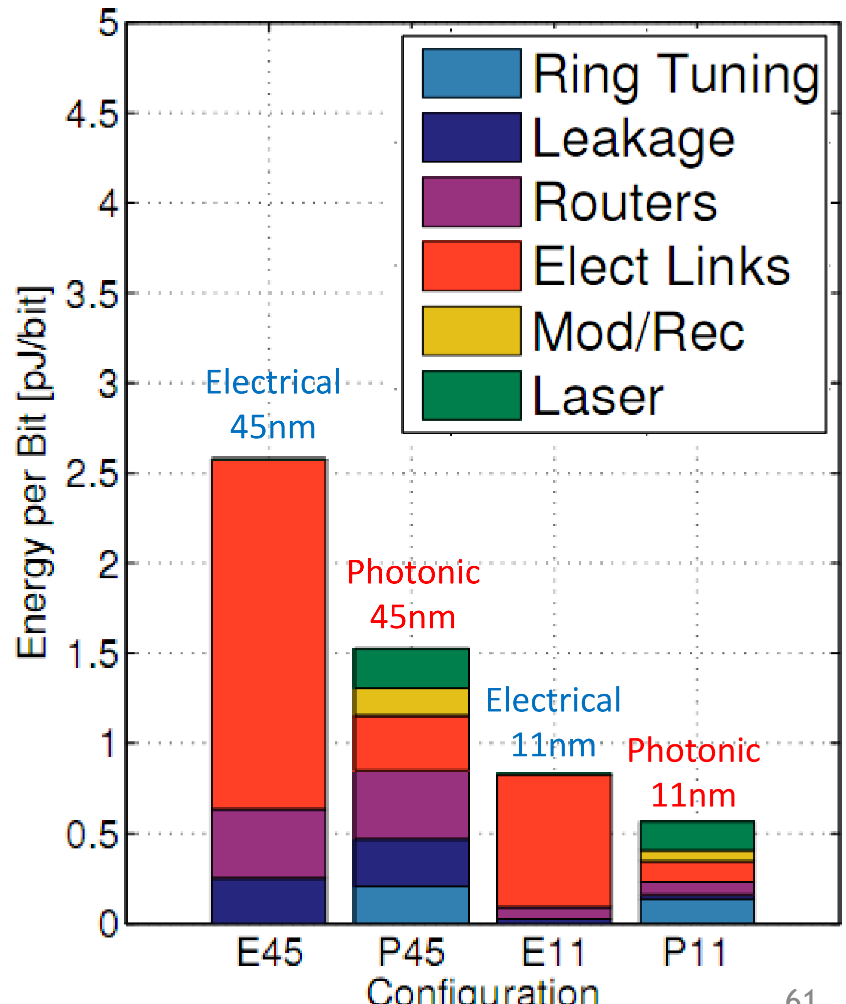

# Energy per Bit Breakdown

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

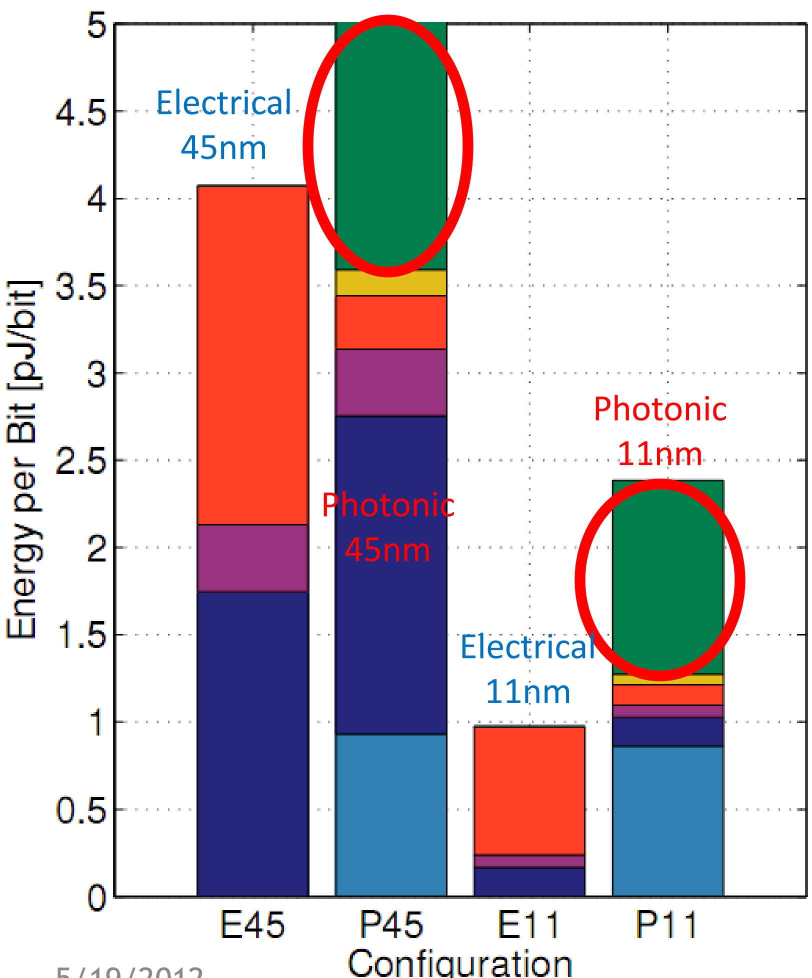


**"Wow non-data-dependent laser really hurts, can I make it better?"**

# Energy per Bit Breakdown

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

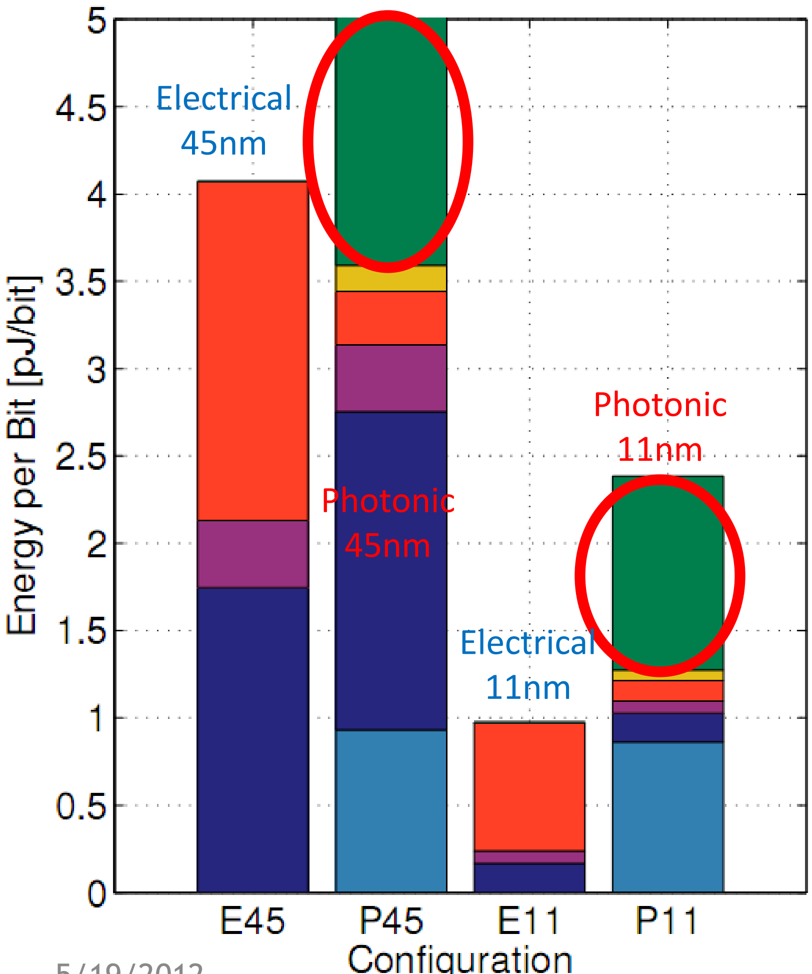


**"Wow non-data-dependent laser really hurts, can I make it better?"**

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

# Tech Parameter Study

**Evaluate the effect of waveguide losses**

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

# Tech Parameter Study

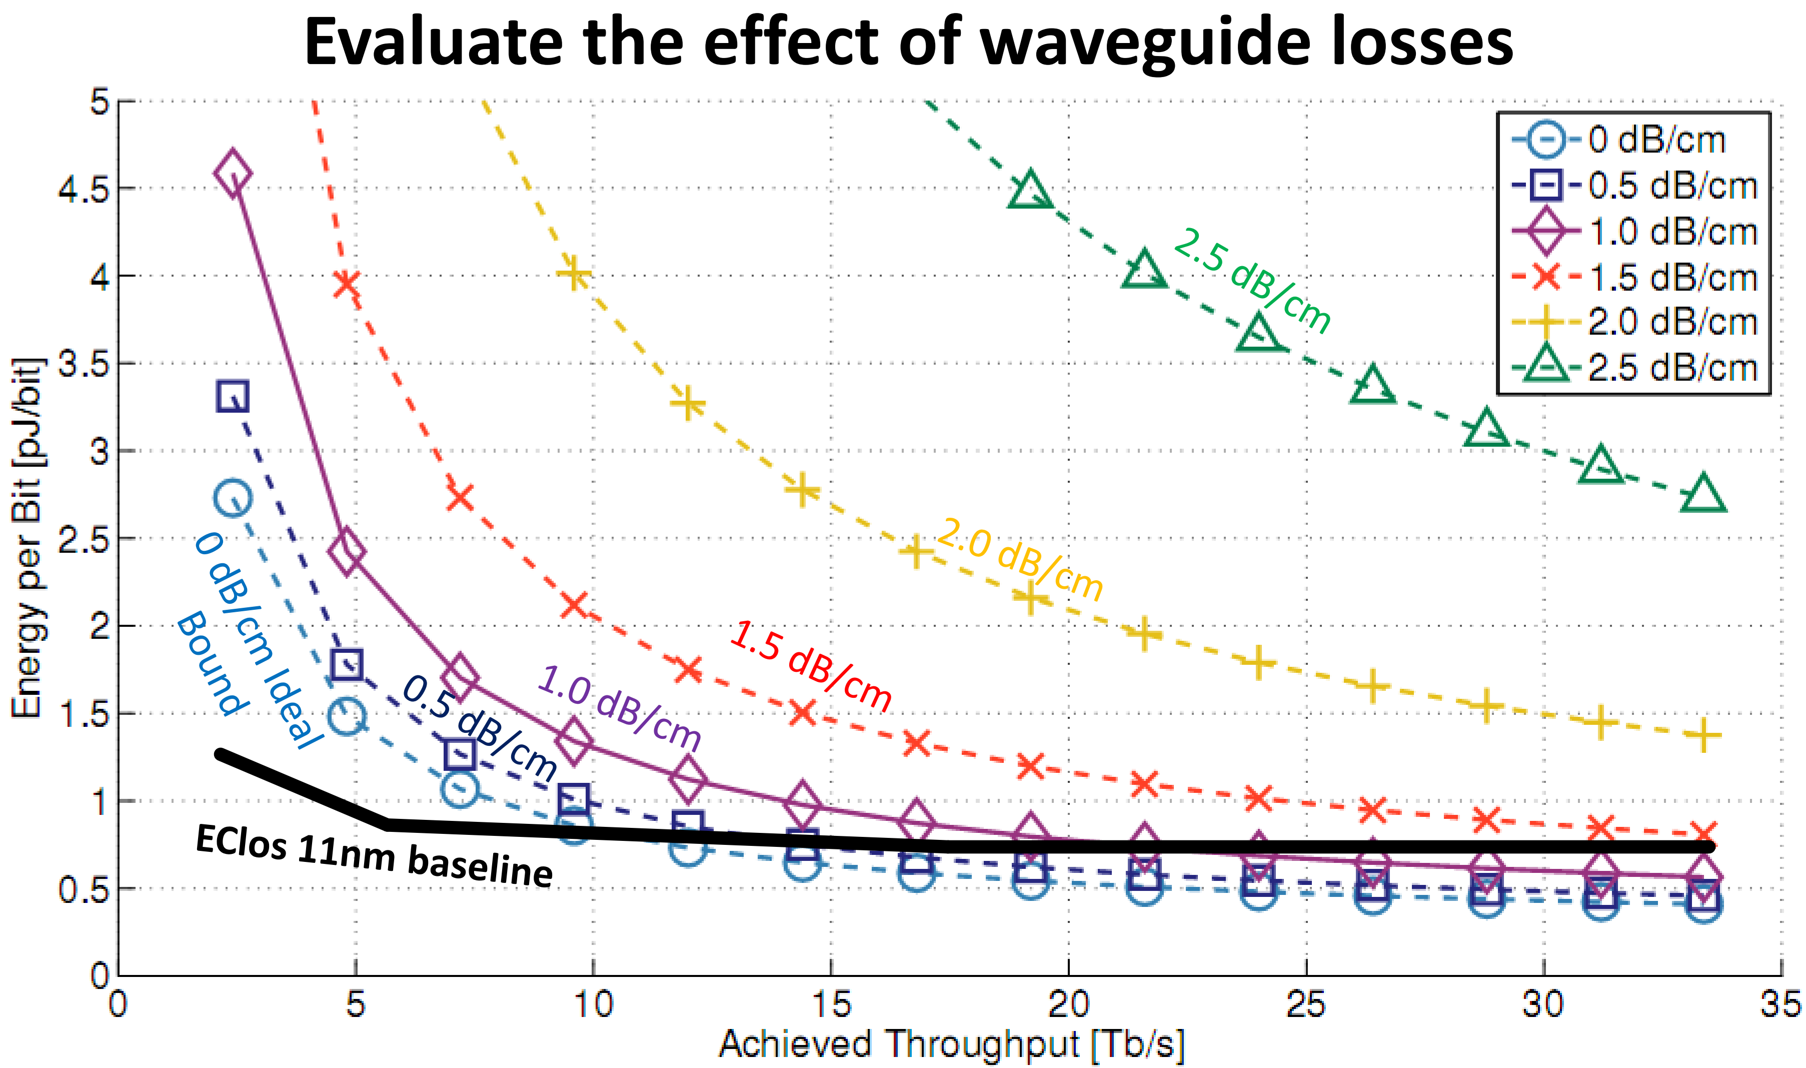

# Tech Parameter Study

## Evaluate the effect of waveguide losses

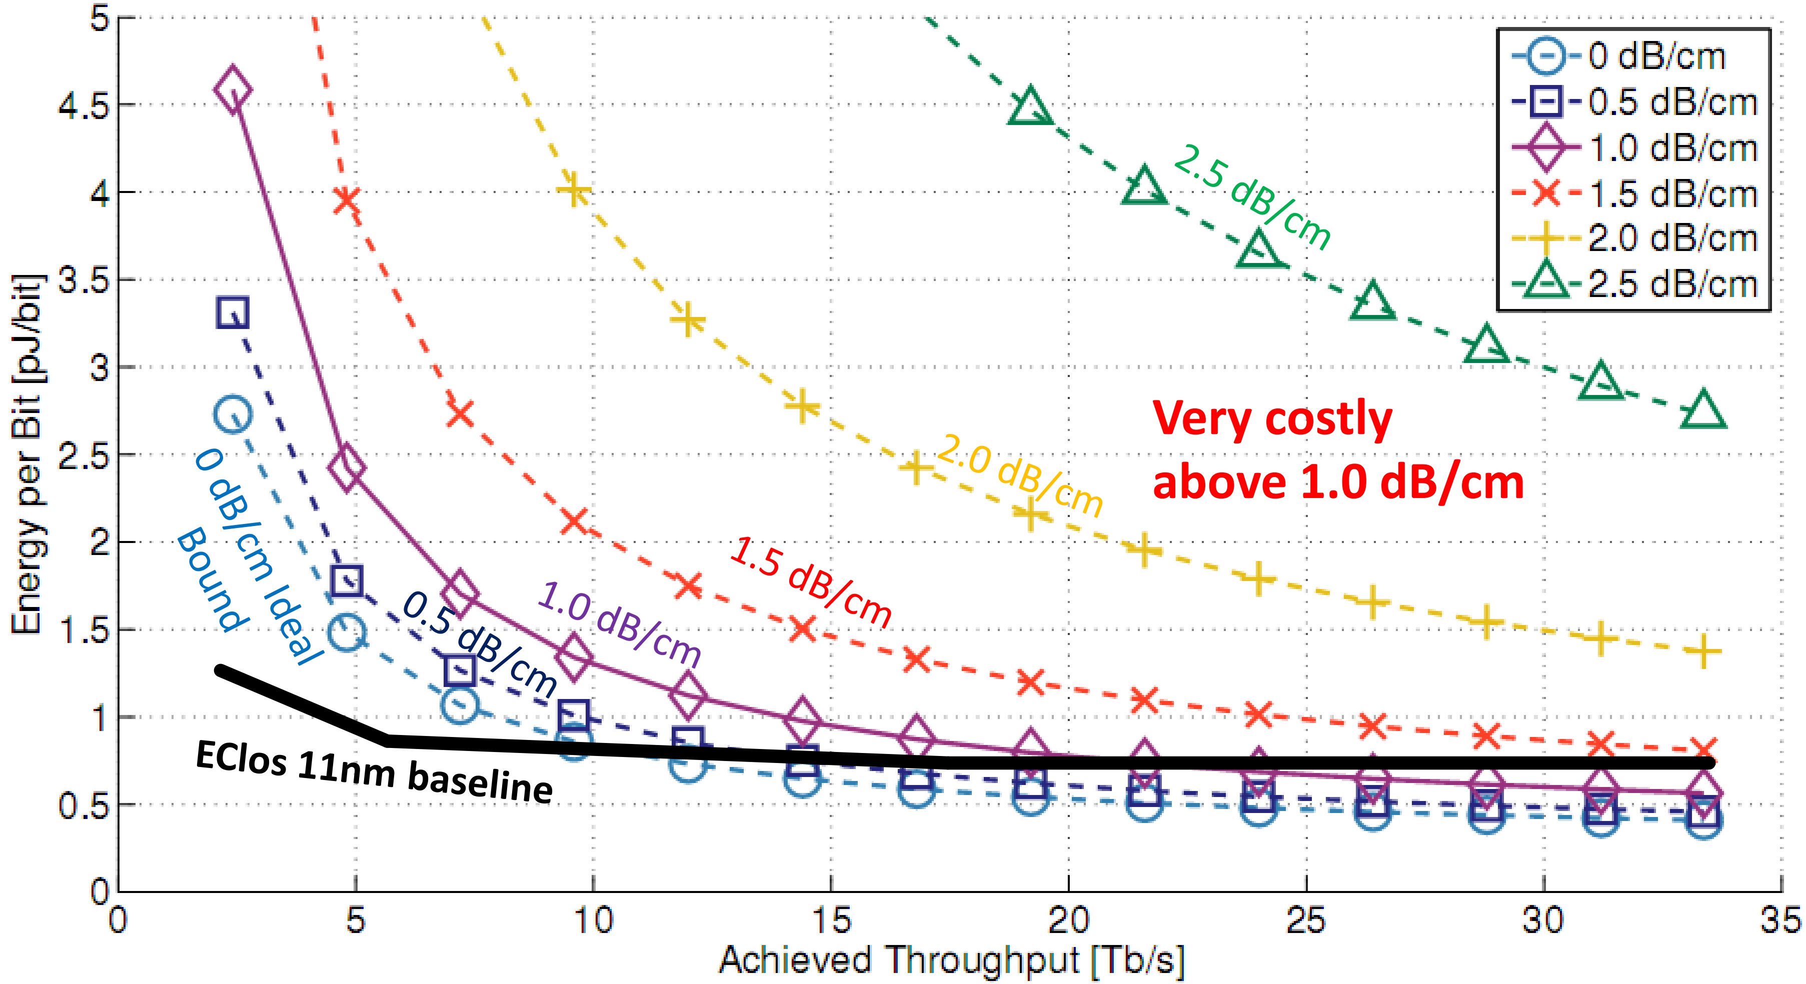

# Tech Parameter Study

## Evaluate the effect of waveguide losses

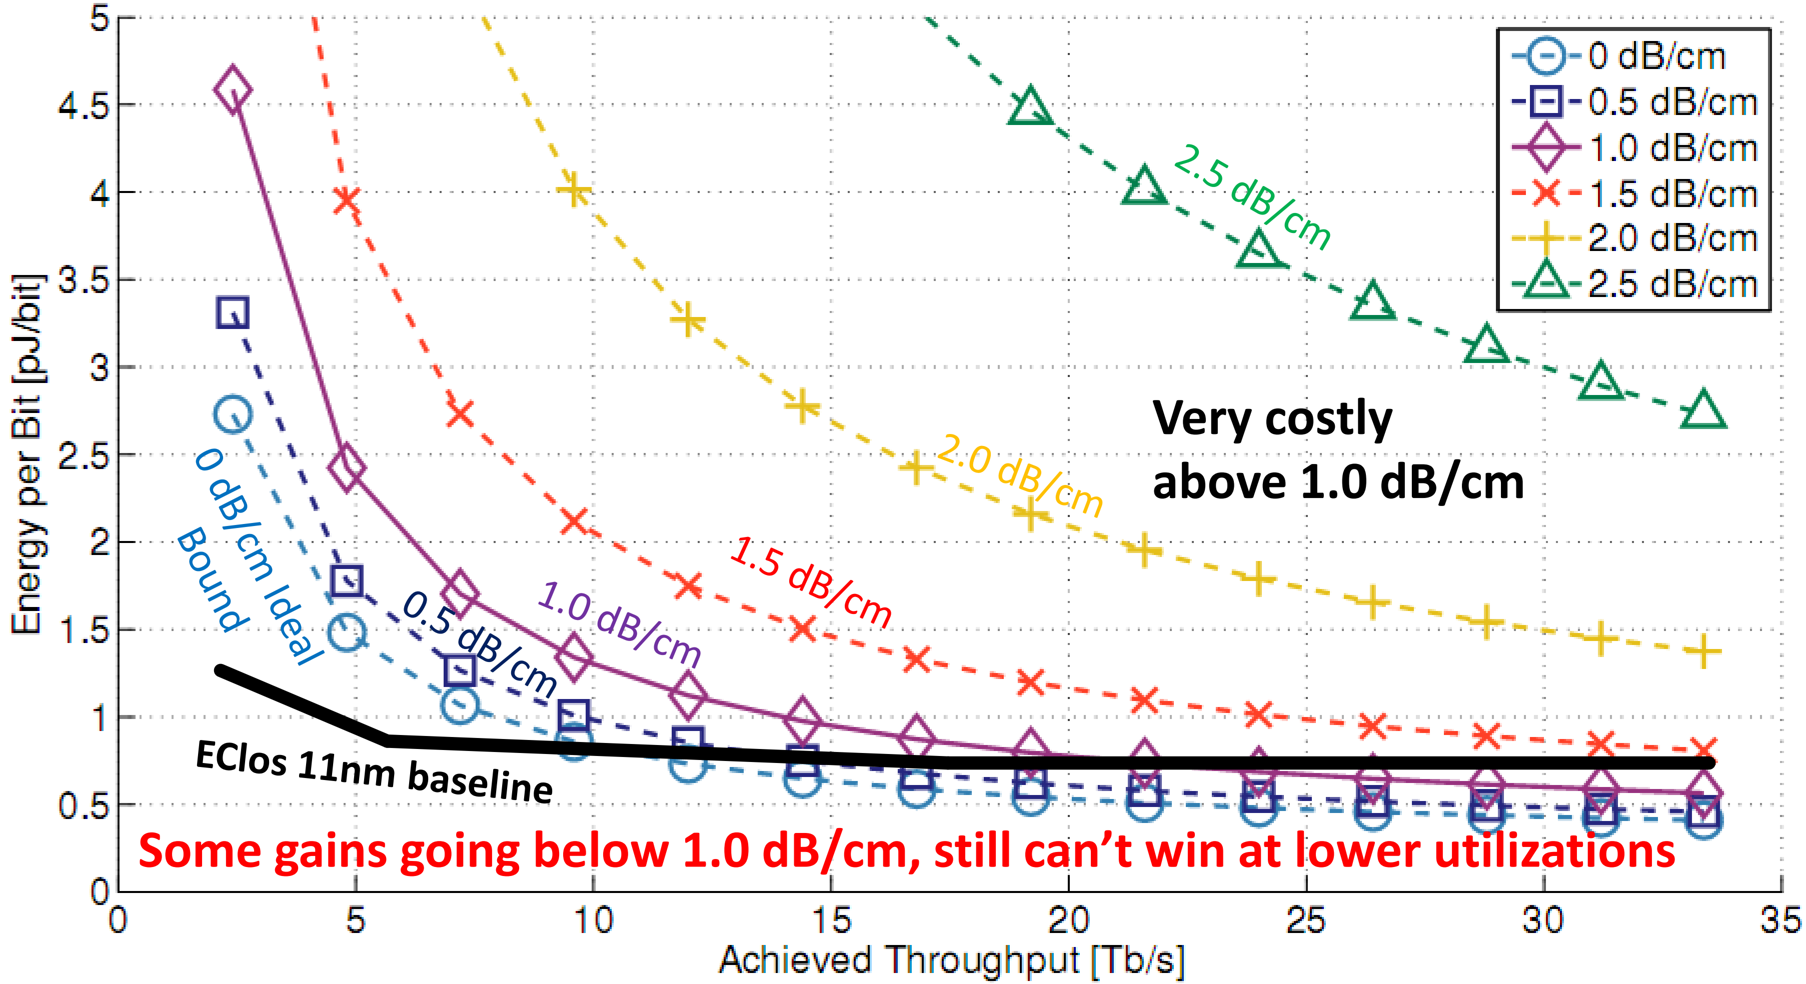

# Tech Parameter Study

## Evaluate the effect of waveguide losses

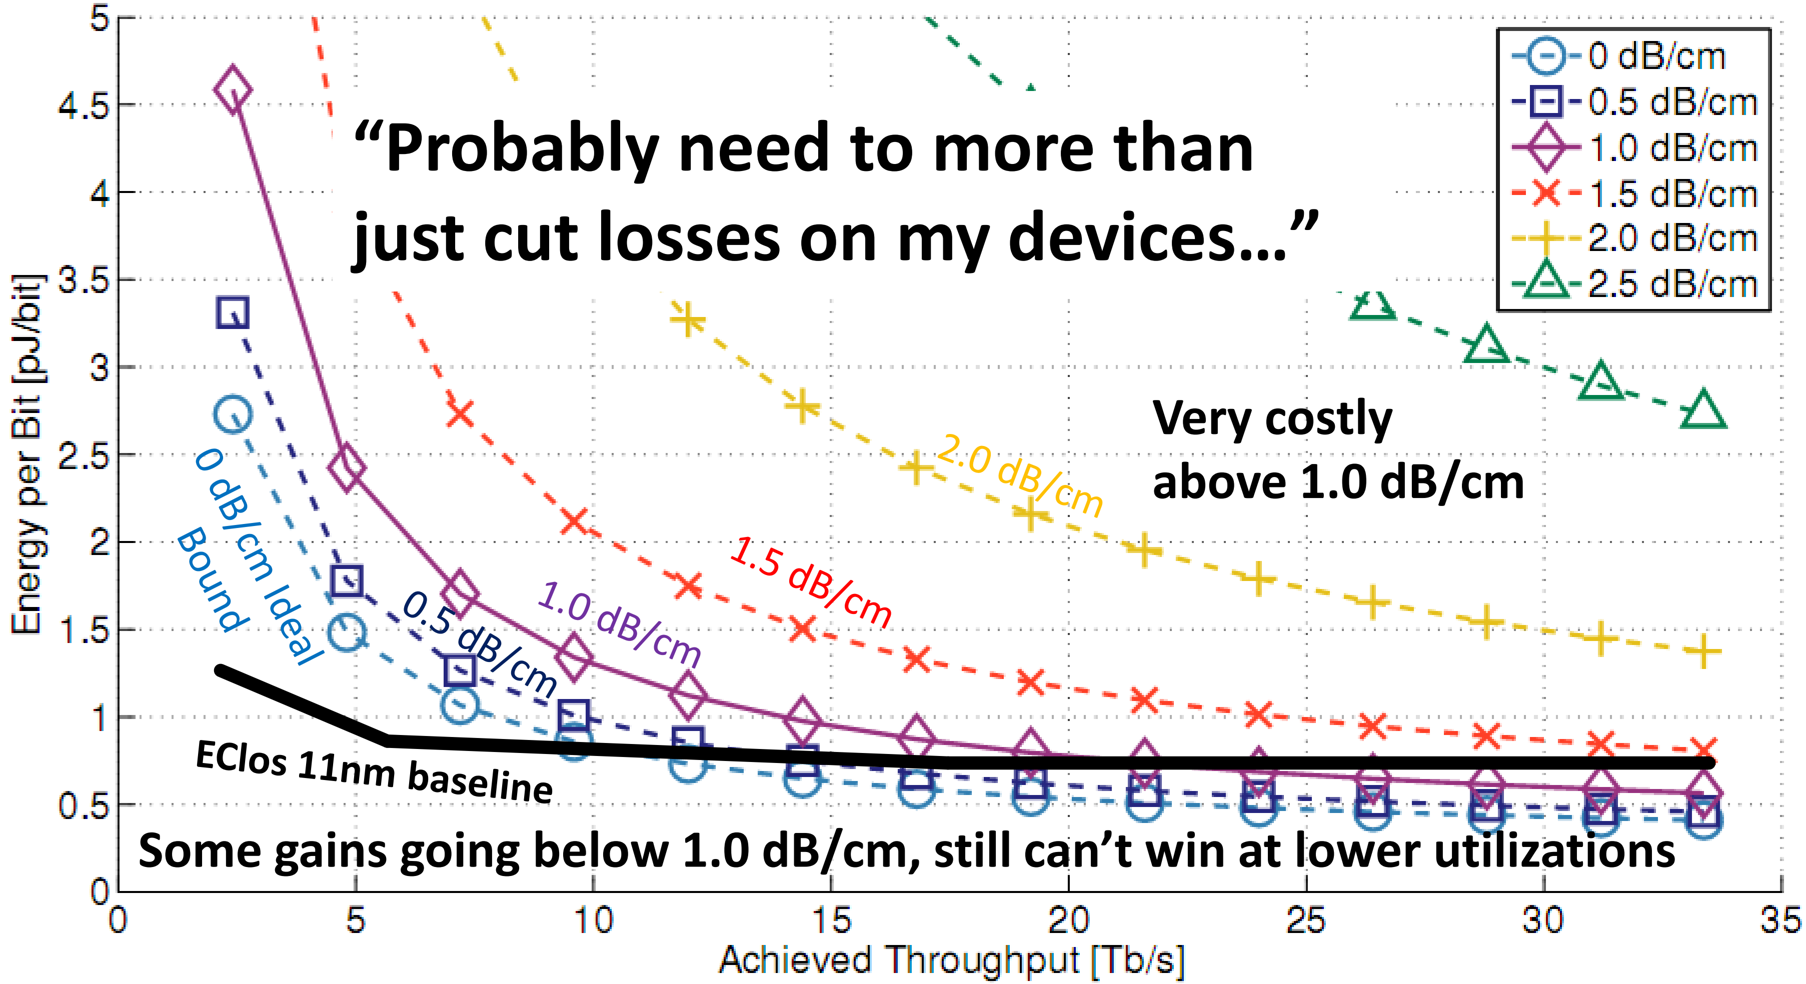


**"Probably need to more than just cut losses on my devices…"**

**Very costly above 1.0 dB/cm**

**Some gains going below 1.0 dB/cm, still can't win at lower utilizations**

# Tech Parameter Study

- Story doesn't end here…

# Tech Parameter Study

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

# Tech Parameter Study

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

# 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

# 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

# Conclusion

- Design decisions in NoCs require evaluation

# Conclusion

- 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

# Conclusion

- 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

# Conclusion

- 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 $mm^2$ |
| 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 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

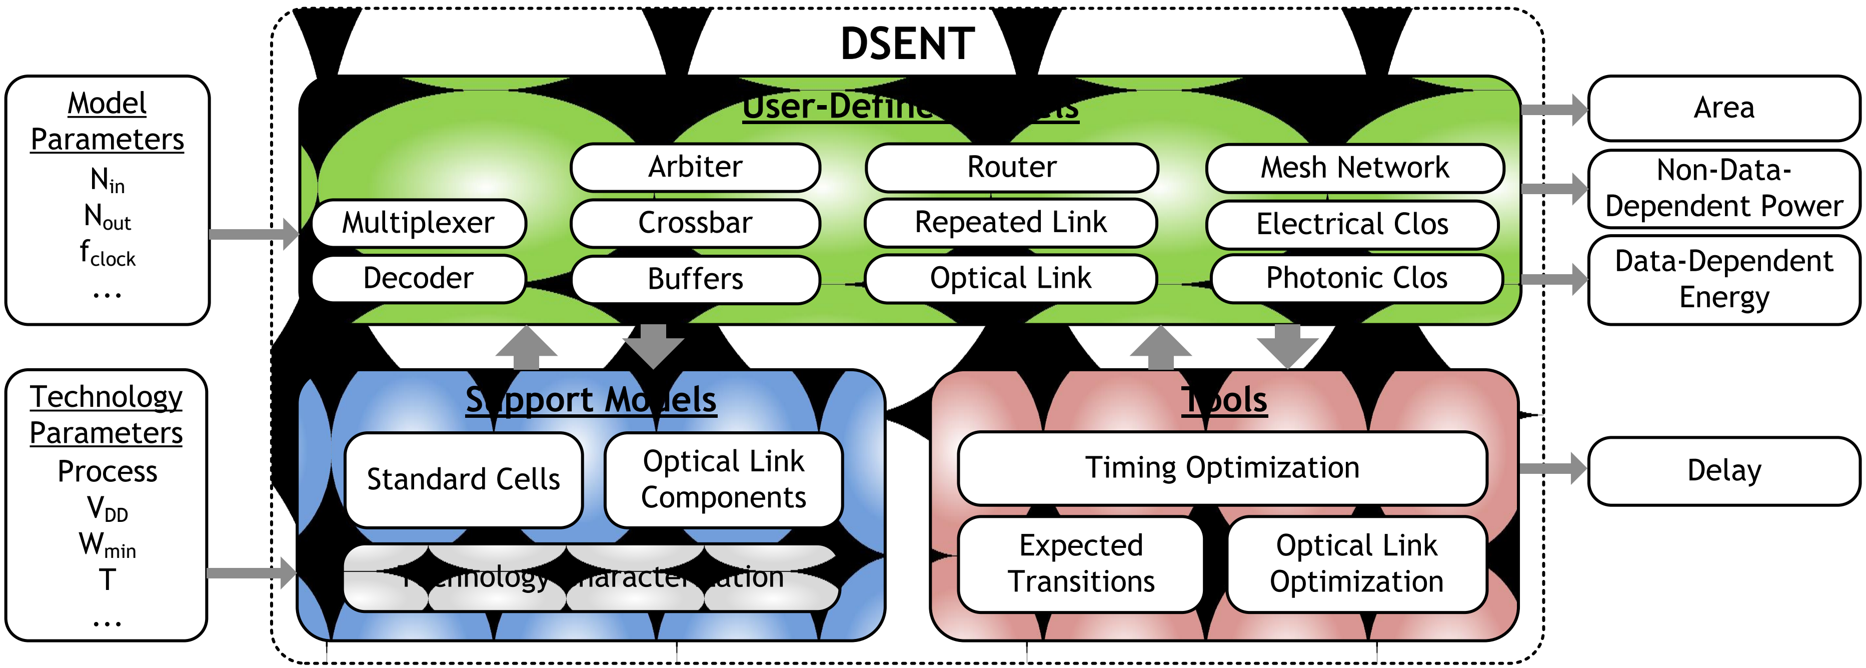

# Technology Characterization

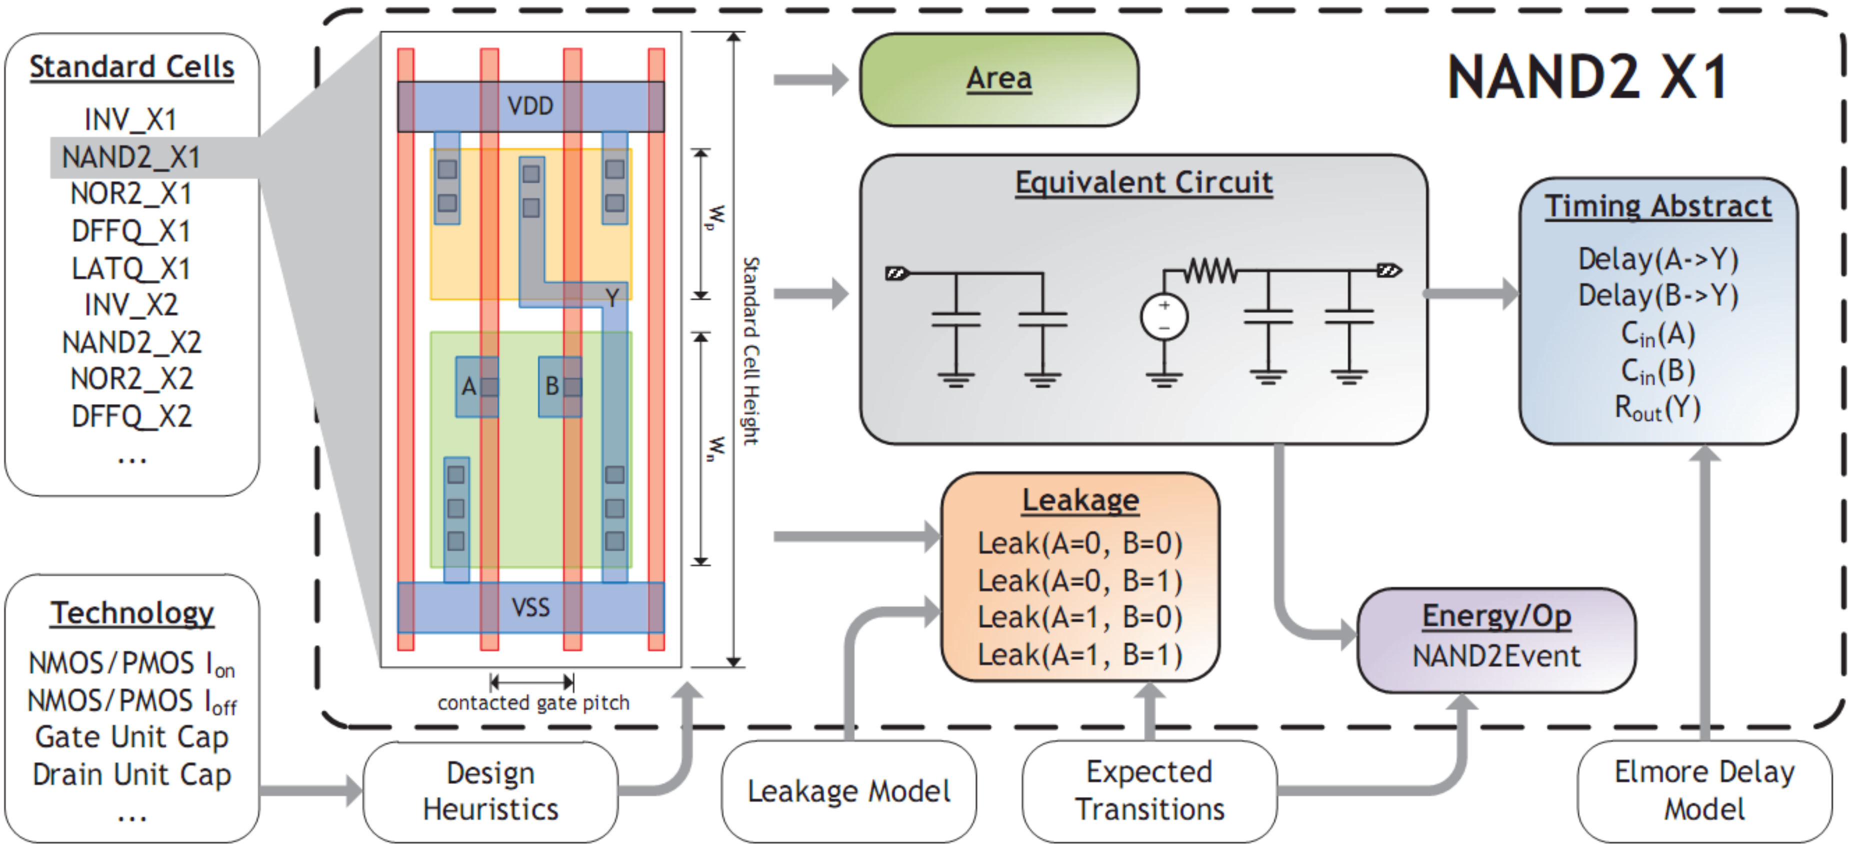

# Optical Models

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

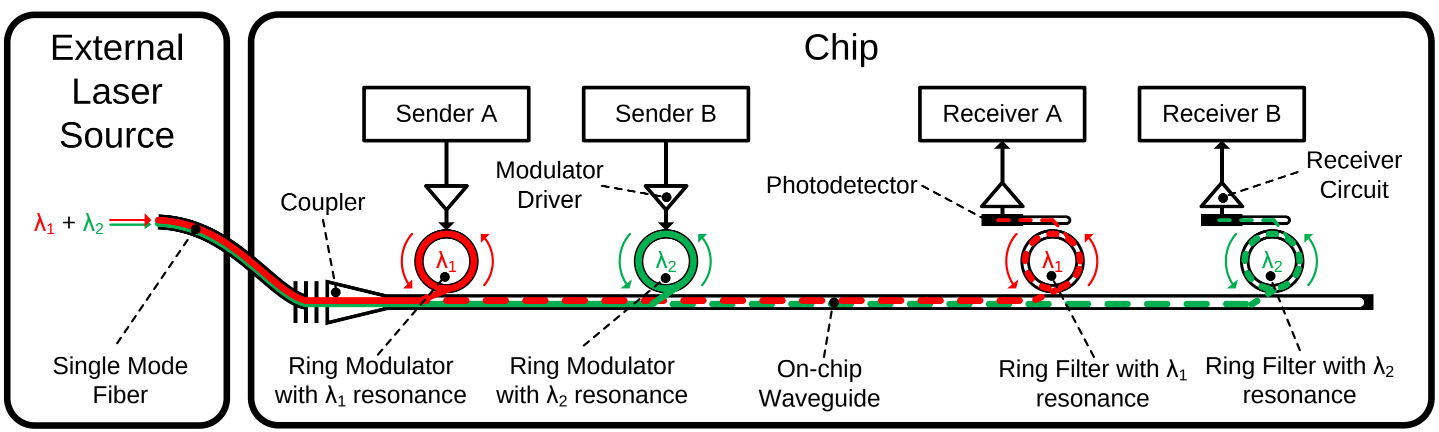

# Timing Optimization

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

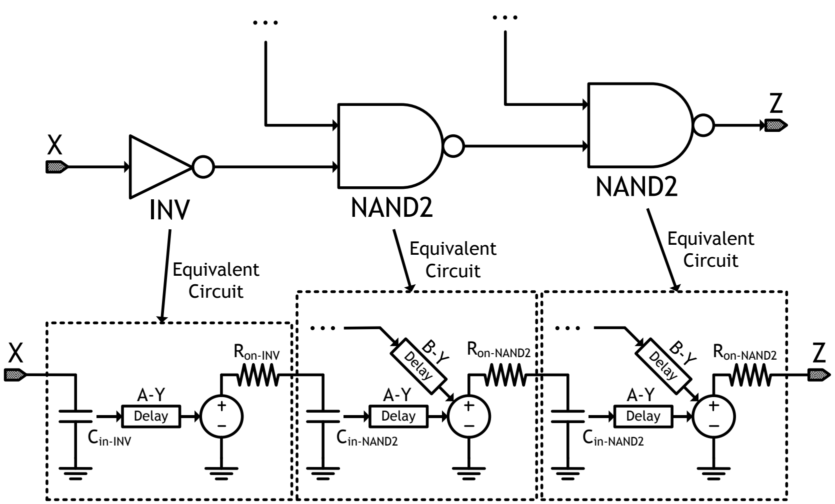


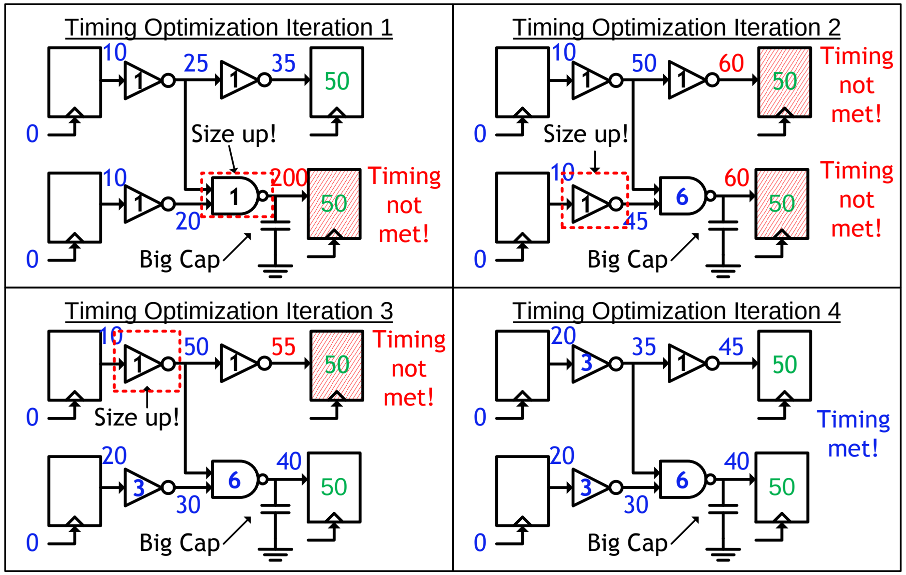

# Expected Transitions

- A simplified expected transition probability model

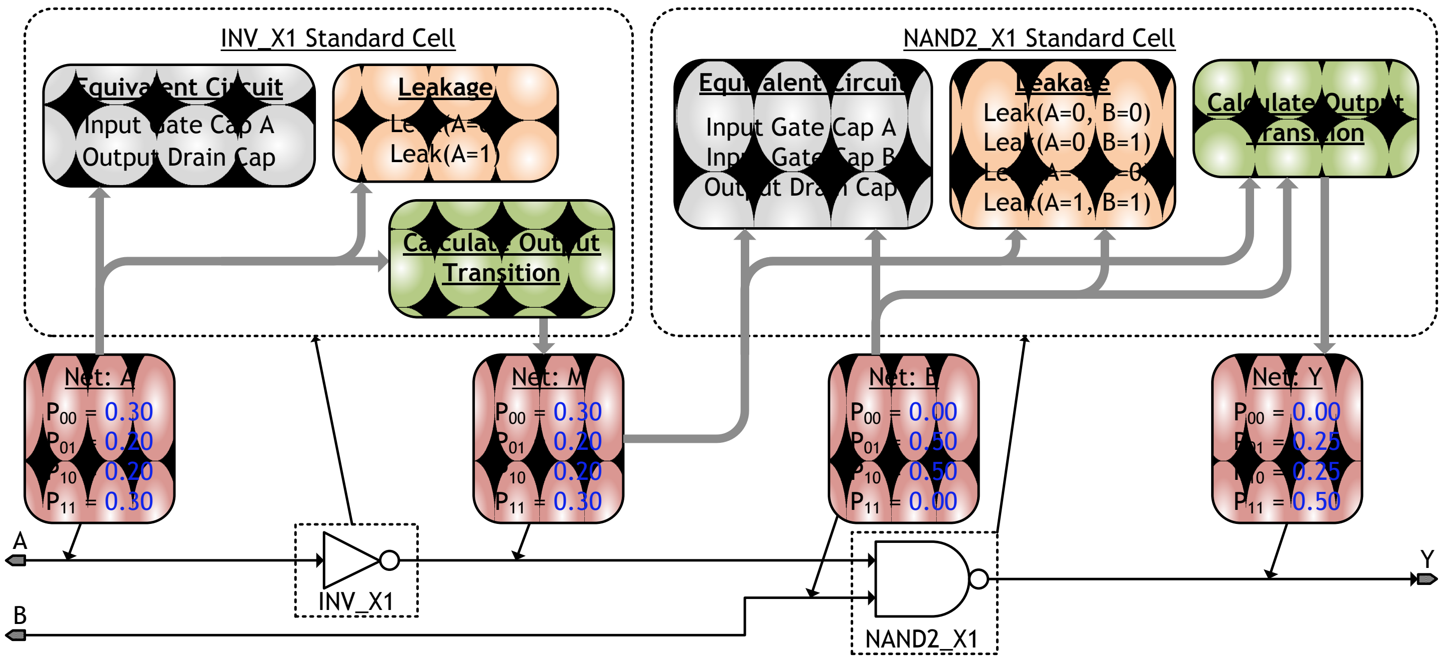

# Power Breakdown (Half)

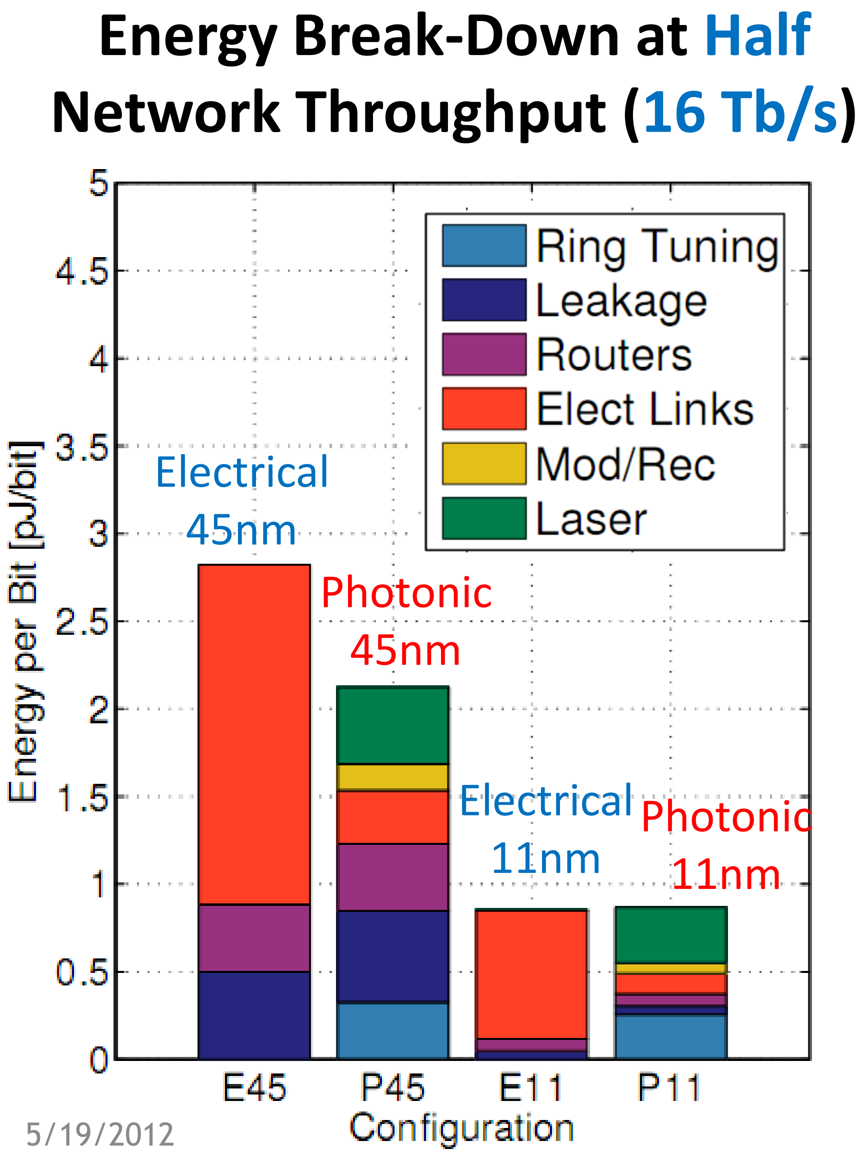


- Photonics (P45, P11) are roughly even with electronics

# Network Case Study

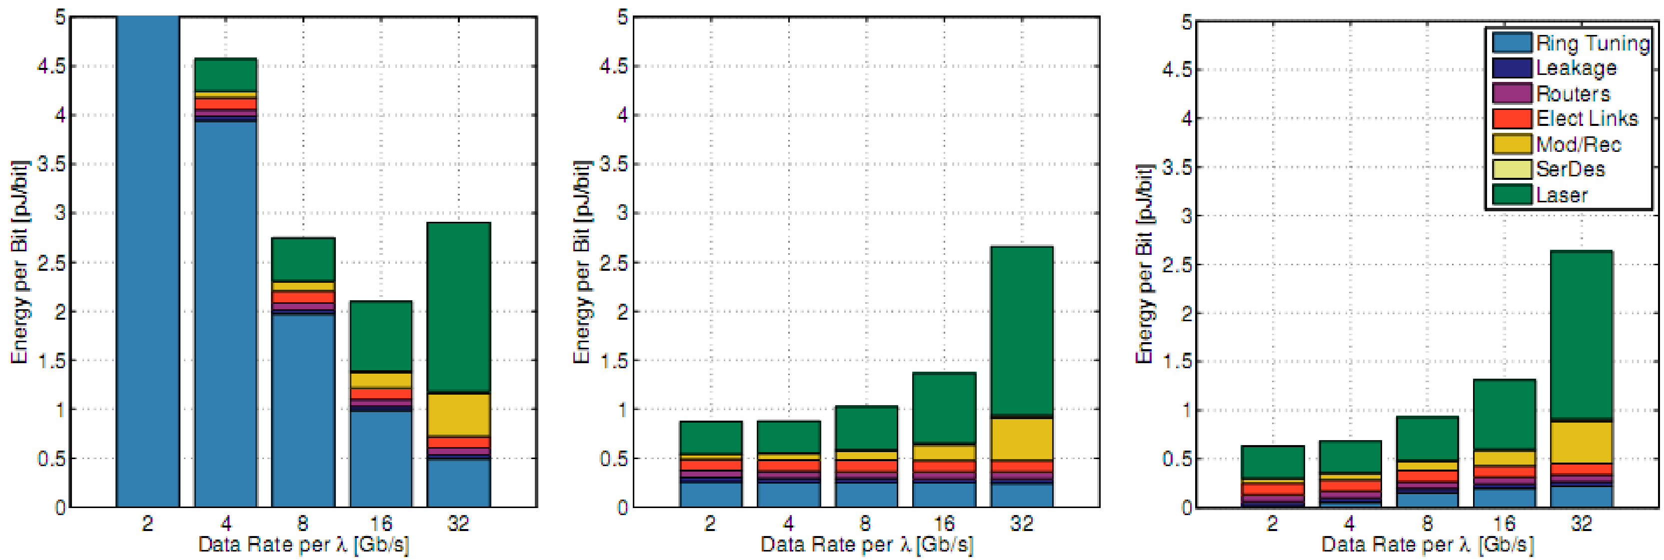


(a) Full-Thermal Tuning (conservative) (b) Bit Reshuffled Tuning (default) (c) Electrically-Assisted Tuning (optimistic)

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 2 GHz core clock, a data-rate of 2, 4, 8, 16, 32 Gb/s per wavelength ($\lambda$) 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

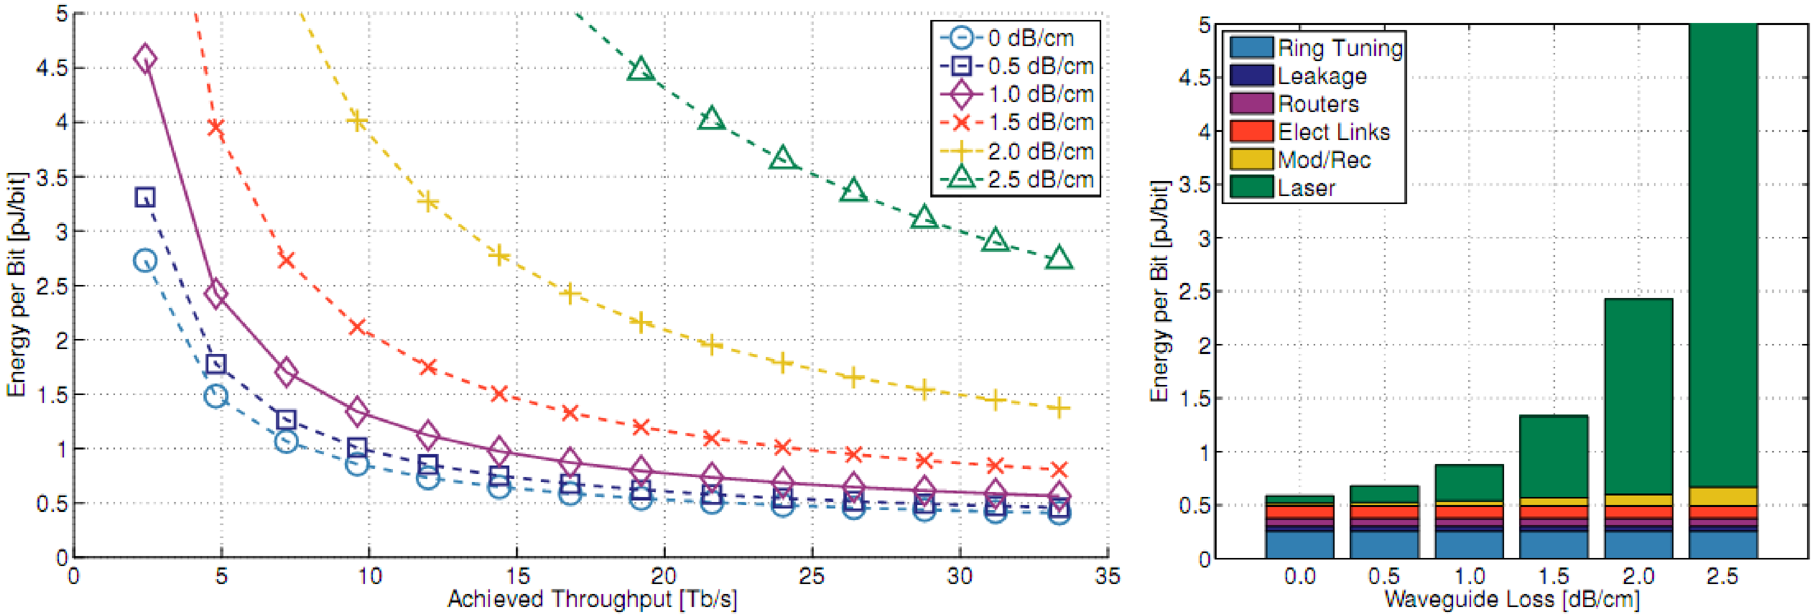


(a) Sensitivity to waveguide loss. Energy per bit vs throughput (left) and energy per bit breakdown at 16 Tb/s throughput (right)

- Ring heating efficiency

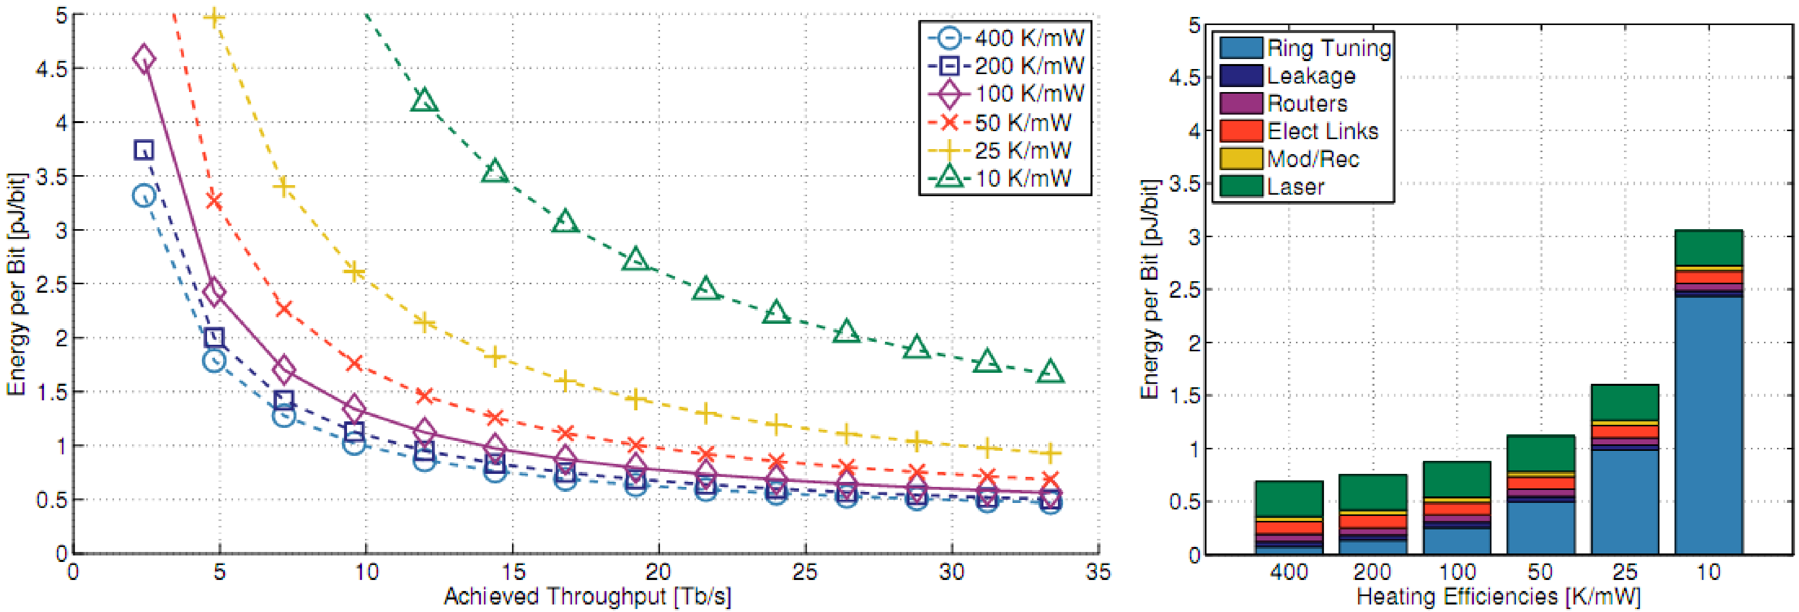


(b) Sensitivity to heating efficiency. Energy per bit vs throughput (left) and energy per bit breakdown at 16 Tb/s throughput (right)

# Tool Validation (45nm SOI)

| Model | | Reference Point | DSENT | Orion2.0 + | Orion2.0 Mod* |
|---|---|---|---|---|---|
| Ring Modulator Driver (fJ/b) | | 50 (11 Gb/s) | 60.87 | N/A | N/A |
| Receiver (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 ($mm^2$) | Encounter – 0.070 | 0.062 | 0.129 | 0.067 |

+ Default Orion 2.0 technology parameters for 45nm

*Correctly specified 45nm tech params

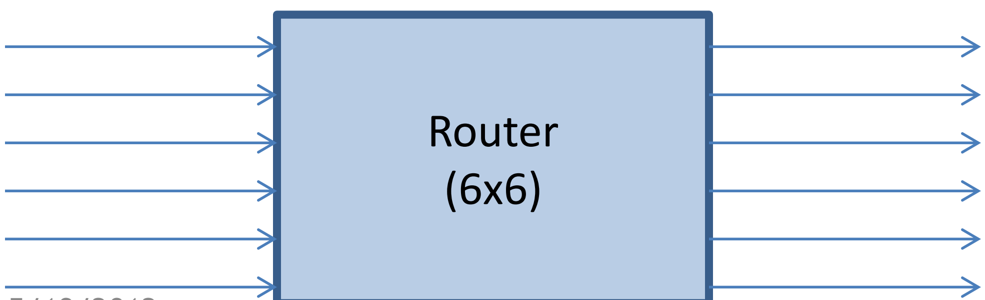


- 6 Input ports, 6 output ports
- 64-bit flit width
- 8 VCs/Port, 16 Buffer FIFO
- 1 GHz clock
- 0.16 flit injection rate

# DSENT Framework

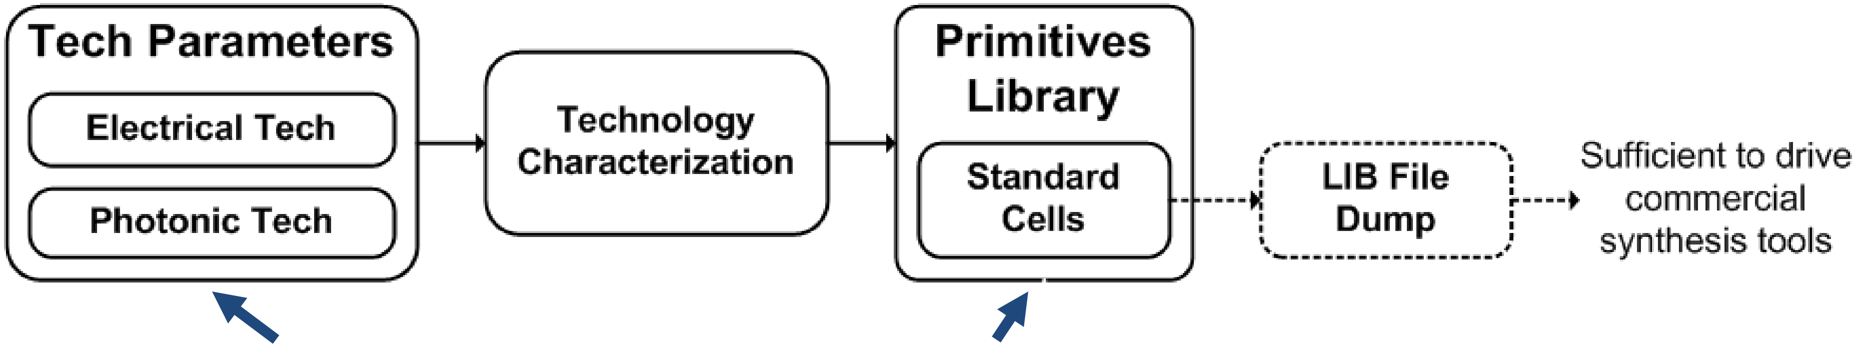


| Technology | Value |
|---|---|
| Supply Voltage | 1.0 V |
| Gate Capacitance / width | 1.0 fF/um |
| Effective on current / width | 650 uA/um |
| Off-current / width | 100 nA/um |
| DIBL | 150 mV/V |
| Sub-threshold Swing | 100 mV/dec |
| Photodetector Responsivity | 1.0 mA/mW |
| … | … |

| Primitive Cells |
|---|
| NAND2 |
| INVERTER |
| BUFFER |
| … |
| Receiver |
| Modulator |
| … |

- Use only basic technology parameters
- Build a usable set of primitives for modeling

- Required technology input mostly limited to what is attainable through ITRS projections and other roadmaps

# DSENT Framework

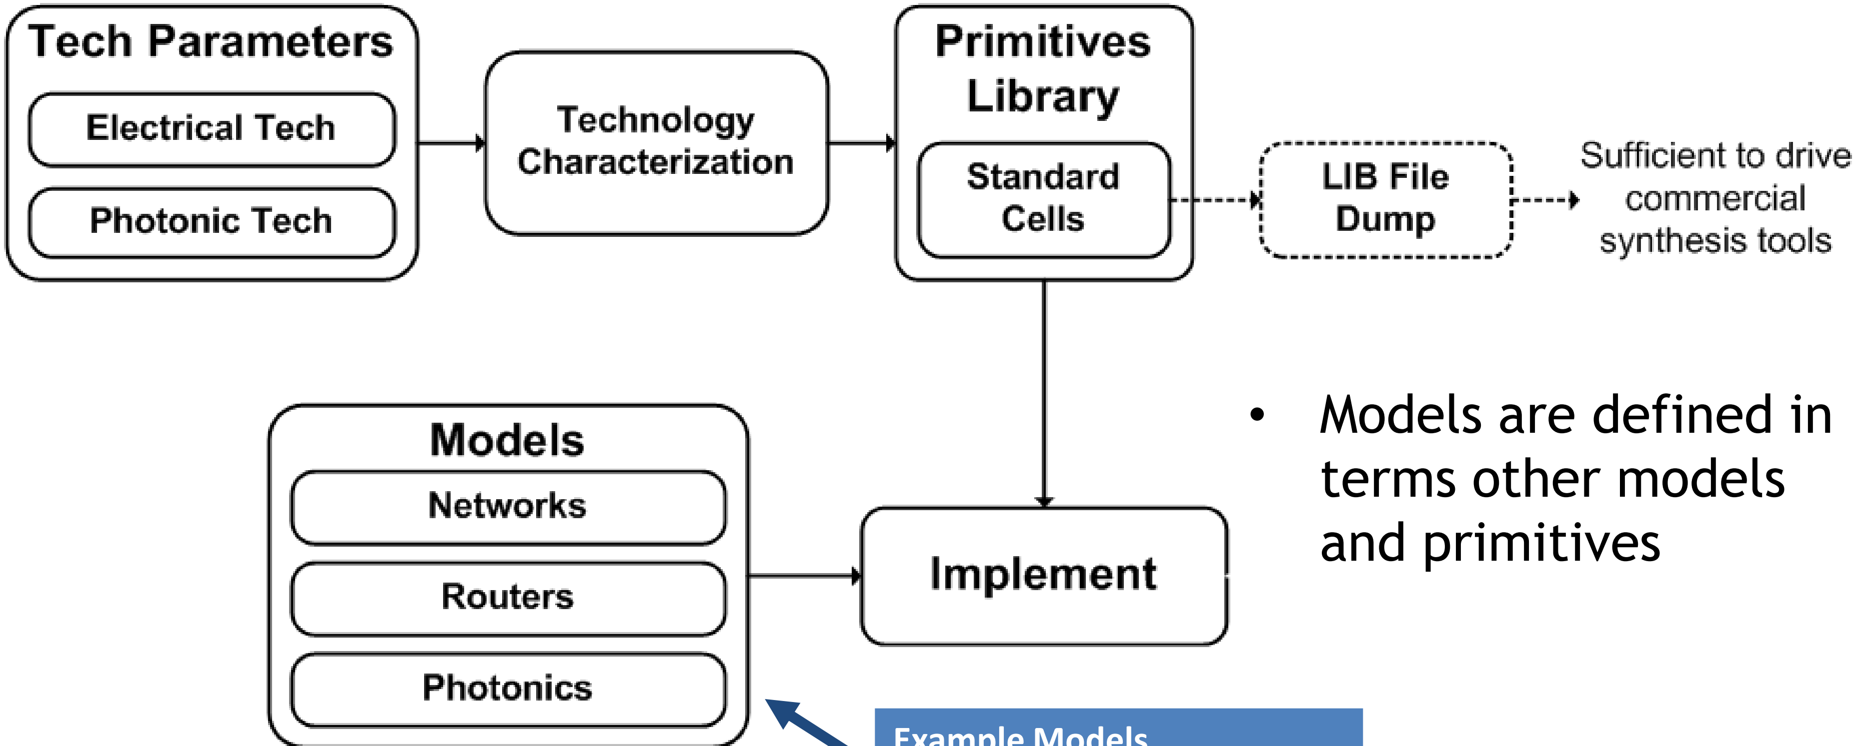


- Models are defined in terms other models and primitives

| Example Models |
|---|
| Mesh Network |
| Clos Network |
| Routers |
| Optical links (SWSR, SWMR) |
| Serializer/Deserializer |
| … |

# DSENT Framework

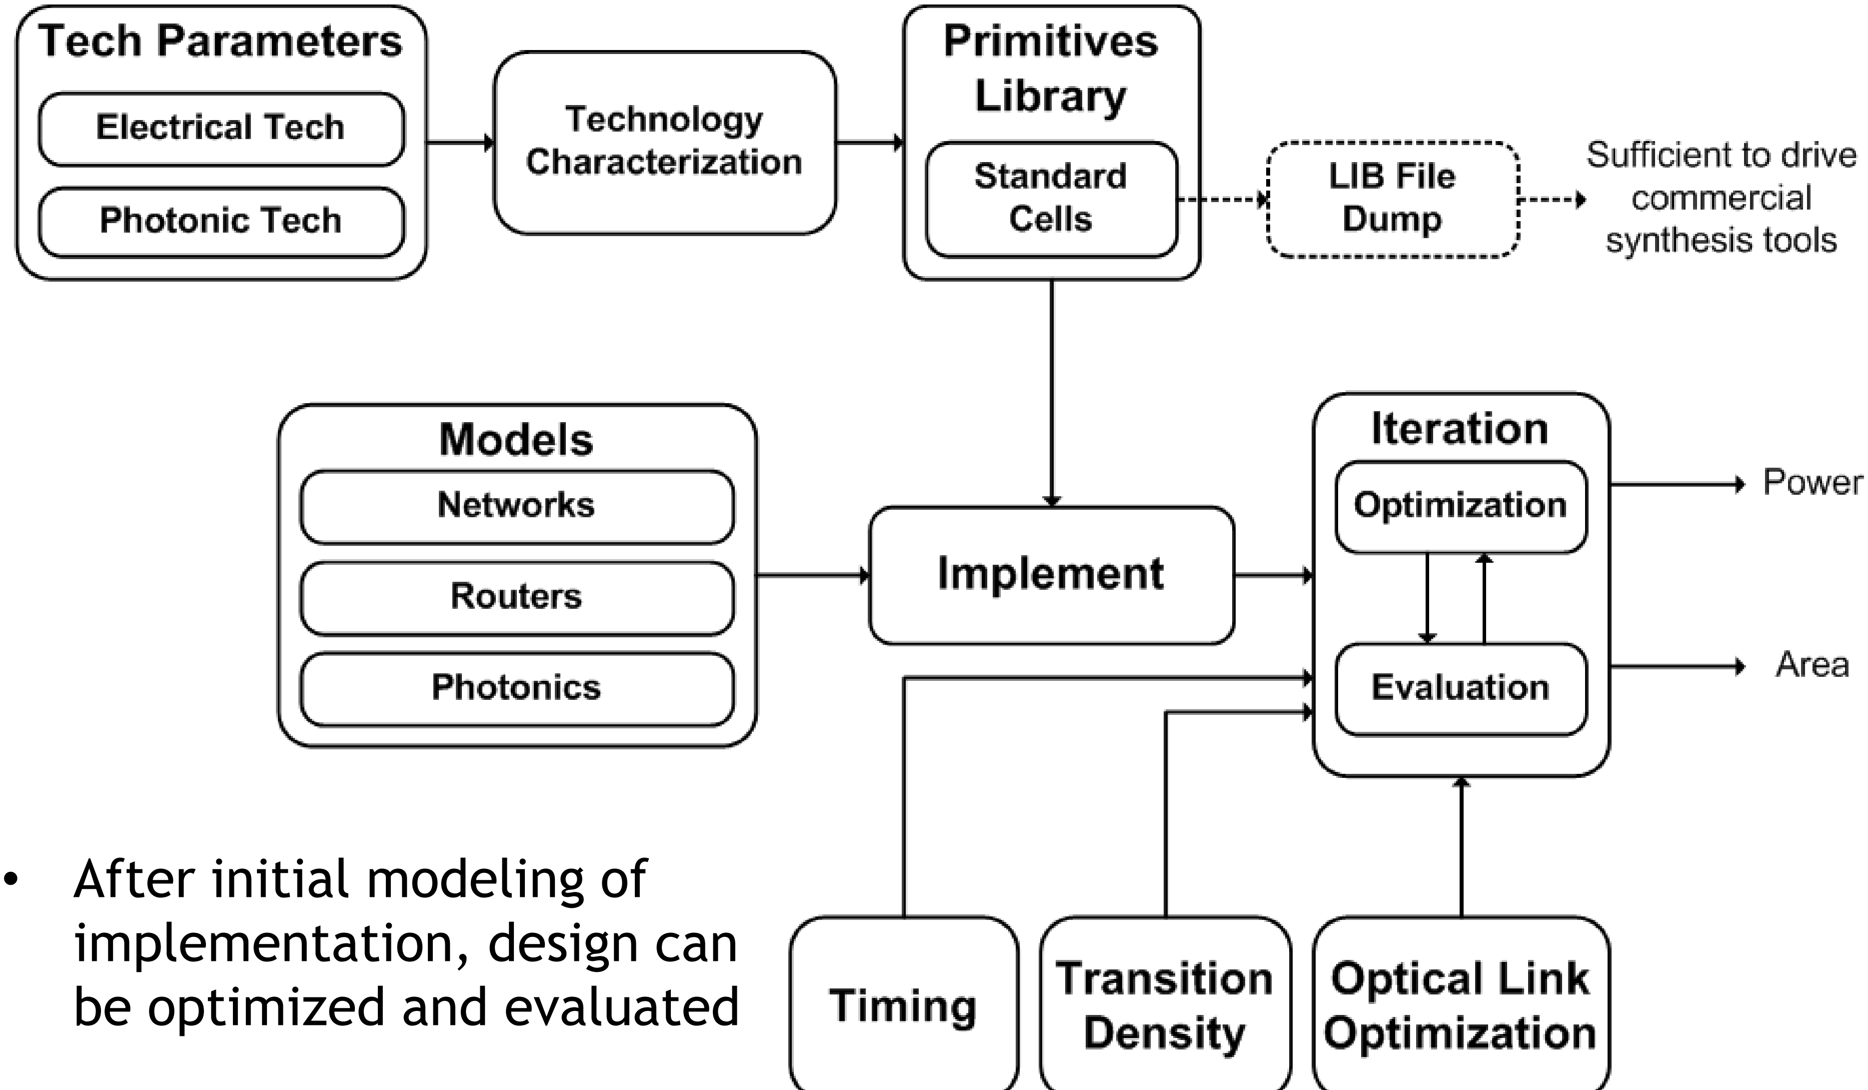


- After initial modeling of implementation, design can be optimized and evaluated

[Georgas, CICC 2011]

# Misc

Error in Cacti 6.5

| Validation Targets | Published $P_{leakage}$ | 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]