Where Does Light Beat Copper? How Manufacturing Variation and Heater Strategy Set the Crossover Length of On-Chip Optical Links

Jone Ng

***Abstract*—**Photonic interconnects can be a more efficient alternative to copper for communication, and the efficiency is determined by the distance and tuning power for ring resonators. I analyzed the effects of manufacturing variations and thermal tuning adjustments on the distance where light becomes more energy-efficient than copper. For the baseline assumption of 2 nm of adjustment per ring, the crossover distance is 4.7 mm, but with fixed wavelength assignments, the variation changes the median crossover to 6.9 mm. Using matched rings with channel reassignment decreases this to 4.5 mm. With copper values from SKY130, the crossover length varies between 0.66 and 18 mm. The heater strategy can change the crossover distance by 1.55x in all four copper configurations. To get these results, I used models of electrical/optical links, models of waveguides, and Monte Carlo sampling on 10,000 virtual chips.

***Index Terms*—**photonics, microring resonators, thermal tuning, manufacturing variations, copper interconnects, crossover distance

I. INTRODUCTION

On-chip communication usually uses copper wires, which consume energy when they charge and discharge. Photonics offers another option by using light instead of electrical signals, but optical links also use energy for heaters, lasers and converting signals to keep them working [1], [2]. The crossover length is how far data needs to travel before light beats copper for energy efficiency.

In previous research, Chen et al. [3] and Samanta et al. [4] compared the energy costs of electrical and optical links without looking into how manufacturing variation shifts the crossover through heater power. Krishnamoorthy et al. [5], Su et al. [6], and Mirza et al. [7] studied how imperfections in the manufacturing can shift the microring resonances, and Zheng et al. [8] looked at ways to cut the energy needed to correct them. DSENT [9] and work by Kurian et al. [10], Ramini et al. [11], and Li et al. [12], looked at optical tuning costs at the network level. However, these studies did not invetigate how the heating used to correct manufacturing variations can change the crossover length.

I used manufacturing variation data from previous experiments to simulate 10,000 chips and estimate how the crossover lengths differed. I also compared three heater-control strategies across two microring designs, and tested whether the link between heater power and crossover length is consistent across four SKY130-based configurations for copper.

II. LINK MODELS AND VARIATION

A. Copper-link energy

I used two copper models to compare: the reference wire has 100 Ω/mm, 200 fF/mm, and 0.8V, and the SKY130 metal 4 values come from the process design kit as described in Table I [15]. The repeater gate time for the wire came from my ngspice simulation using the SKY130 inverter, giving me the gate time of 12.32 ps in Table I. The electrical baseline had repeaters sized and

spaced according to an Elmore-delay model and established methods [13], [14]. Its energy use per bit is around

$$E_{\mathrm{Cu}}(L) = \alpha\, V_{\mathrm{DD}}^2\, L\left(c_{\mathrm{wire}} + \frac{C_{\mathrm{rep}}}{s}\right) \quad (1)$$

where L is length, α is the charging activity, VDD is the supply voltage, cwire is the capacitance per unit length, Crep is the switched repeater capacitance, and s is the distance between repeaters.

B. Photonic-link energy

As the distance increases, more laser power is needed to keep the signal up as light gets lost in the waveguide. The link also has a modulator, a waveguide, a filter ring, and a photodetector. With N wavelength channels each transmitting at rate B, the energy per delivered bit is about

$$E_{\mathrm{opt}}(L) = \frac{P_{\mathrm{sens}}\, 10^{\mathcal{L}(L)/10}}{\eta_{\mathrm{laser}}\, B} + E_{\mathrm{mod}} + E_{\mathrm{rx}} + \frac{P_{\mathrm{heat}}}{NB} \quad (2)$$

Here, Psens is the receiver optical sensitivity, $\mathcal{L}$(L) is the end-to-end loss in dB, ηlaser is the laser wall-plug efficiency, and Pheat is the total heater power across the N channels, assuming that all channels are active. Lower use increases the overhead per bit. Halving the per-channel bit rate from 10 to 5 Gb/s increases the reference crossover from 4.73 to 5.82 mm. That crossover L* satisfies ECu(L*) = Eopt(L*).

TABLE I
PRINCIPAL MODEL PARAMETERS

Input	Value	Basis
SKY130 met4 resistance	156.7 Ω/mm	PDK
Wire capacitance, isolated/packed	75.9/260.1 fF/mm	PDK + est.
Repeater gate time	12.32 ps	ngspice
Supply voltage	1.8/0.8 V	Scenarios
Charging activity *α*	0.25	Chosen
Bit rate / channels	10 Gb/s / 8	Chosen
Receiver energy	191 fJ/bit	[4]
Receiver sensitivity	−17 dBm	[4]
Laser efficiency / waveguide loss	30% / 1 dB/cm	[4]
Coupler / modulator loss	3/1 dB	[4]
Wafer width / thickness 1σ	5.52/1.35 nm	[6]
Temperature range	20 K	[2]
Local resonance spread 1σ	0.27 nm	[7]

Input	Value	Basis
Heater power / shift	100 μW/nm	[2]
Ring radius	5/7.11 μm	Chosen

C. Manufacturing variation and validation

Each of 10,000 virtual chips get a shared wafer-scale geometry and a temperature shift over a 20 K range, combined with ring-to-ring variation in resonance. Eight channels are spaced apart by about 1.6 nm. A mode solver calculates the geometry and temperature sensitivity for a silicon waveguide. I compared three heater strategies: fixed channel assignment, cyclic barrel shift, and minimum cost one-to-one reassignment. I chose two ring sizes: the traditional 5-μm ring and 7.11-μm rings because its free spectral range of 12.8 nm matches the span of the 8 channels spaced 1.6 nm apart. The heaters only shifted their resonances to work towards longer wavelengths of light.

When using matched rings for reassignment, the average wavelength adjustment needed was 1.19 nm, which is close to the value below 1.2 nm that was reported from experiments [5]. Checking the electrical model against ngspice, I found that the Python wire solver matched the result. The delay approximation differed from the simulation of the circuit by 0.9-5.9%. The calculated ring sensitivities differed by 1 and 16% from the published measurements [6].

III. RESULTS

A. Reference energy crossover

The reference optical link uses 307.5 fJ/bit at zero length: 191 fJ/bit for the receiver, 50fJ/bit for the modulation, 40 fJ/bit for the heaters, and 26.5 fJ/bit for the operation of the laser. The reference electrical wire uses about 65.6 fJ/bit/mm. These values give an energy crossover of 4.73 mm (Fig. 1). The 40 fJ/bit contribution from the heater is the standard assumption of two rings per channel, with each requiring 2 nm of shifting at 100 μW/nm and operated at 10 Gb/s. Across all tested parameter ranges, switching activity from 0.1 to 0.5 changes the crossover from 12.0 to 2.4 mm, and changing the receiver energy from 20-200 fJ/bit changes it from 2.1 to 4.9 mm. Waveguide loss from 0.5 to 6 dB/cm only changes it from 4.7 to 5.1 mm.

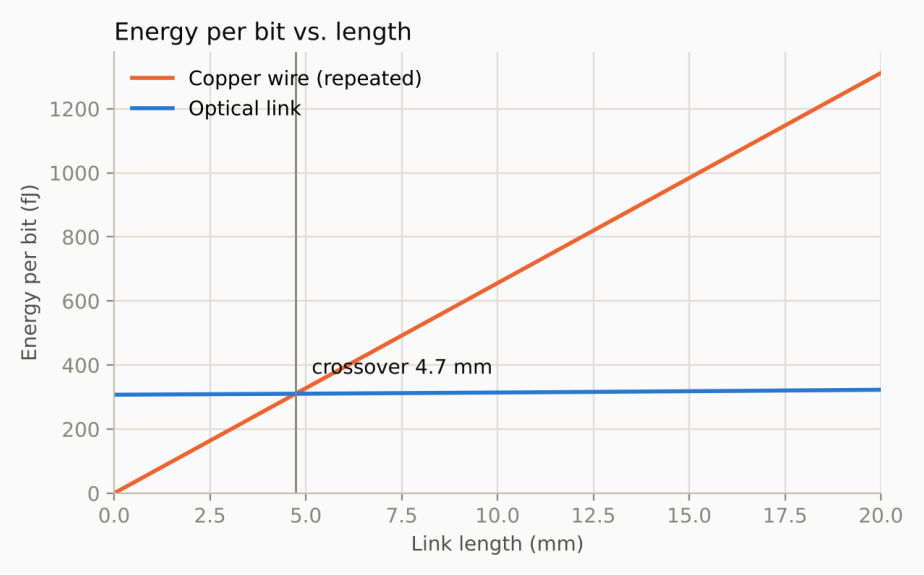

Fig. 1. The copper wire's energy per bit vs. the optical link's; the lines intersect at 4.73 mm.

B. Tuning strategies under variation

The most improvement came from using 7.11-μm rings with reassignment instead of 5-μm rings with fixed assignment. A 7.7x reduction was observed when the average adjustment went from 9.13 to 1.19 nm. The crossover distance went from 6.93 to 4.49 mm, a 1.55x difference. Between chips, the ranges containing 90% of crossover distances were vastly different, with 5.03 compared to 0.40 mm.

TABLE II

HEATER TUNING/CROSSOVER LENGTH BY RING DESIGN AND CHANNEL ASSIGNMENT

Ring design	Channel assignment	Mean heating per ring (nm)	Median crossover (mm)	90% of chips (mm)
5 μm (18.2 nm free spectral range)	Fixed	9.13	6.93	4.41–9.44
5 μm	Barrel shift	2.94	4.96	4.41--5.86
5 μm	Full reassignment	2.94	4.96	4.41–5.86
7.11 μm (12.8 nm free spectral range)	Fixed	6.42	6.10	4.41–7.78
7.11 μm	Barrel shift	1.19	4.49	4.28–4.68
7.11 μm	Full reassignment	1.19	4.49	4.28–4.68

Medians of 10,000 chips with reference wire. Barrel shift and full reassignment have identical results.

As seen in Table II, allowing reassignment was the most effective, with the 6.93 to 4.96 mm decrease for 5-μm rings and the 6.10 to 4.49 mm decrease for the 7.11-μm ring. Increasing the ring radius only dropped it slightly from 6.93 to 6.10 mm. Surprisingly, the simpler barrel shift method also worked just as well as reassignment; this is because most of the variations affect all of the rings similarly.

Fig. 2 shows the percentage of chips where light uses less energy than copper at each length. One problem with fixed assignments is that heaters can only move rings towards longer wavelengths. This means that even if a wavelength is slightly higher than its channel, it may need almost a full 18.2 nm of adjustment. This explains the split where 43% cross before 6 mm but another 47% do not cross until after 7.5 mm. Matched rings avoid this problem by allowing rings to switch to better-fitting channels when the wavelength is too high rather than needing a large adjustment. This is shown through the crossover being between 4.48 and 4.49 mm even if the width variation increases. Heater efficiency also has a great effect; when the power needed for a 1 nm shift increases from 50 to 225 μW/nm, the crossover changes from 4.30 to 4.93 mm, compared to a much greater change of 5.36 to 9.73 mm for fixed channels.

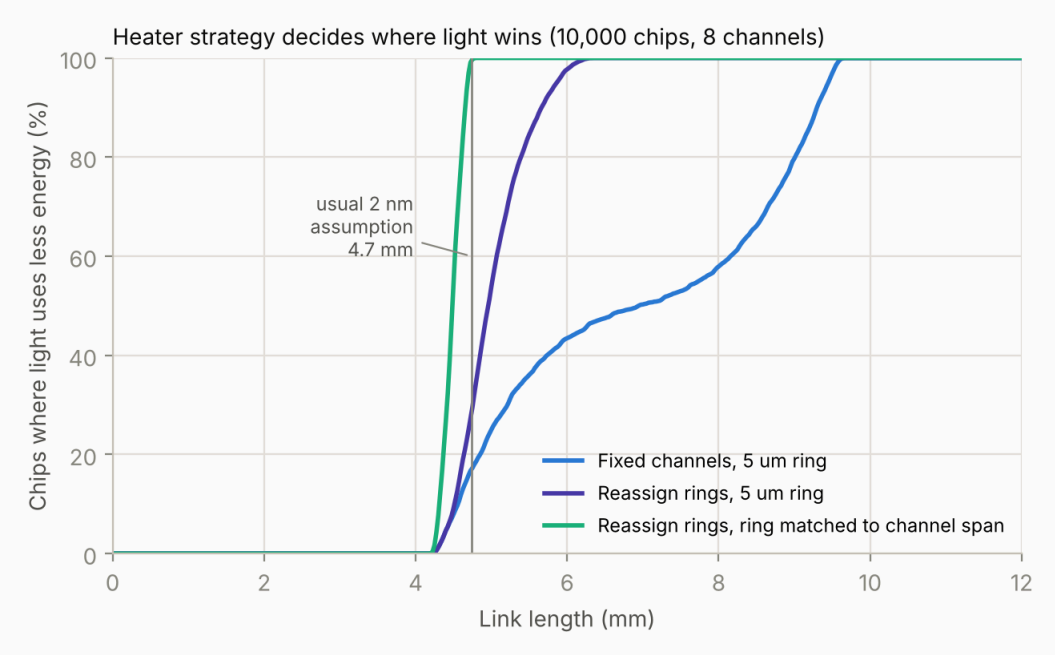

Fig. 2. Share of 10,000 simulated chips on which the optical link uses less energy than copper, versus link length. Vertical line: the conventional 2-nm-per-ring assumption (4.7 mm).

C. SKY130 sensitivity

For SKY130 met4, the modeled copper energy ranges from 26 to 443 fJ/bit/mm depending on the voltage and nearby conductors (Table III). The combined range of optical crossover distances is approximately 0.66 to 18 mm. Here, the matched rings with reassignment are the most efficient in each configuration, while fixed channels require approximately 1.55x the

crossover length. The absolute values depend much more on the electrical technology and coupling assumptions.

TABLE III
SKY130 COPPER ENERGY AND OPTICAL CROSSOVER LENGTH

Copper case	Copper (fJ/bit/mm)	Matched (mm)	Fixed (mm)	2-nm (mm)
1.8 V, packed	443	0.66	1.02	0.70
1.8 V, isolated	129	2.26	3.50	2.39
0.8 V, packed	88	3.35	5.18	3.54
0.8 V, isolated	26	11.73	18.17	12.38

Medians of 10,000 chips. Matched refers to the 7.11-μm rings with reassignment. Fixed refers to the fixed channels with 5-μm rings. 2-nm refers to the conventional fixed tuning assumption.

IV. DISCUSSION AND CONCLUSION

Despite the crossover distances varying from 0.66 to 18 mm, the fixed-channel crossover length was consistently around 1.55x the crossover of matched rings. At length zero, the optical energy without heaters is 267.5 fJ/bit; heating adds 182.6 fJ/bit using fixed channels or only 23.8fJ/bit with matched rings and reassignment. By setting (1) equal to (2) and assuming the extra energy from waveguide loss is small compared to the copper energy, the crossover is around

$$L^{*} \approx \frac{E_{\mathrm{opt}}(0)}{e_{\mathrm{Cu}}} \quad (3)$$

where eCu is the copper energy per bit per unit length and Eopt(0) is the optical energy per bit at zero length. Dividing the two crossovers cancels eCu, so the ratio only depends on the fixed optical energies. The equation predicts (267.5+182.6)/(267.5+23.8) = 1.545, which is close to the 1.55x pattern that I observed. This ratio even holds for the longest crossover in Table III, where 18.17/11.73=1.549. This value only applies to this specific optical link and is not a universal constant.

The model did not account for the heater driver energy, which could have made fixed-channel operation less effective. My SKY130 copper values came from a 130 nm process, while the photonic aspects were based on newer technologies, and the main results used a 10 Gb/s connection. The coupling of the packed wire was a geometric estimate. Finally, this study was simulation-only, and it examined a single link instead of an entire network.

In conclusion, manufacturing variation affects the crossover through the heating strategy used. Also, the results suggest that the 2 nm assumption might not accurately represent all variations in manufacturing, as matched rings achieved a 4.49 mm crossover and fixed channels had a crossover of 6.93, differing from the 4.73 mm crossover from the conventional assumption. Using just one value can underestimate or overestimate the crossover distance.

CODE AND ACKNOWLEDGMENT

Code/model assumptions: https://github.com/jonemng-hub/light-vs-electrons

I installed and configured the simulation tools, ran each simulation, generated figures and tables, checked reported values against the code outputs, and checked references against their sources. AI was used to help write and debug the simulation code, search the literature, and structure/review the manuscript.

REFERENCES

[1] M. Haurylau, G. Chen, H. Chen, J. Zhang, N. A. Nelson, D. H. Albonesi, E. G. Friedman, and P. M. Fauchet, "On-chip optical interconnect roadmap: Challenges and critical directions," *IEEE J. Sel. Topics Quantum Electron.*, vol. 12, no. 6, pp. 1699–1705, 2006, doi: 10.1109/JSTQE.2006.880615.

[2] S. Manipatruni, M. Lipson, and I. A. Young, "Device considerations for nanophotonic CMOS global interconnects," arXiv:1207.6819, 2012. [Online]. Available: https://arxiv.org/abs/1207.6819

[3] G. Chen, H. Chen, M. Haurylau, N. A. Nelson, D. H. Albonesi, P. M. Fauchet, and E. G. Friedman, "Predictions of CMOS compatible on-chip optical interconnect," *Integration, the VLSI Journal*, vol. 40, pp. 434–446, 2007, doi: 10.1016/j.vlsi.2006.10.001.

[4] A. Samanta, S.-H. Lee, C.-Y. Cheng, S. Palermo, and S. J. B. Yoo, "3D electronic-photonic heterogenous interconnect platforms enabling energy-efficient scalable architectures for future HPC systems," arXiv:2510.03943, 2025. [Online]. Available: https://arxiv.org/abs/2510.03943

[5] A. V. Krishnamoorthy, X. Zheng, G. Li, J. Yao, T. Pinguet, A. Mekis, H. Thacker, I. Shubin, Y. Luo, K. Raj, and J. E. Cunningham, "Exploiting CMOS manufacturing to reduce tuning requirements for resonant optical devices," *IEEE Photon. J.*, vol. 3, no. 3, pp. 567–579, 2011, doi: 10.1109/JPHOT.2011.2140367.

[6] Z. Su, E. S. Hosseini, E. Timurdogan, J. Sun, G. Leake, D. D. Coolbaugh, and M. R. Watts, "Reduced wafer-scale frequency variation in adiabatic microring resonators," in *Proc. Optical Fiber Communication Conf. (OFC)*, San Francisco, CA, USA, 2014, doi: 10.1364/OFC.2014.Th2A.55.

[7] A. Mirza, F. Sunny, P. Walsh, K. Hassan, S. Pasricha, and M. Nikdast, "Silicon photonic microring resonators: A comprehensive design-space exploration and optimization under fabrication-process variations," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 41, no. 10, pp. 3359–3372, 2022, doi: 10.1109/TCAD.2021.3132555.

[8] Y. Zheng, P. Lisherness, M. Gao, J. Bovington, S. Yang, and K.-T. Cheng, "Power-efficient calibration and reconfiguration for on-chip optical communication," in *Proc. Design, Automation & Test in Europe Conf. (DATE)*, 2012, doi: 10.1109/DATE.2012.6176711.

[9] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic, "DSENT—a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling," in *Proc. IEEE/ACM Int. Symp. Networks-on-Chip (NOCS)*, 2012, doi: 10.1109/NOCS.2012.31.

[10] G. Kurian, C. Sun, C.-H. O. Chen, J. E. Miller, J. Michel, L. Wei, D. A. Antoniadis, L.-S. Peh, L. C. Kimerling, V. Stojanovic, and A. Agarwal, "Cross-layer energy and performance

evaluation of a nanophotonic manycore processor system using real application workloads," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. (IPDPS)*, 2012, doi: 10.1109/IPDPS.2012.103.

[11] L. Ramini, P. Grani, H. Tatenguem Fankem, A. Ghiribaldi, S. Bartolini, and D. Bertozzi, "Assessing the energy break-even point between an optical NoC architecture and an aggressive electronic baseline," in *Proc. Design, Automation & Test in Europe Conf. (DATE)*, 2014, pp. 1–6, doi: 10.7873/DATE.2014.321.

[12] Z. Li et al., "Reliability modeling and management of nanophotonic on-chip networks," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 20, no. 1, pp. 98–111, 2012, doi: 10.1109/TVLSI.2010.2089072.

[13] W. C. Elmore, "The transient response of damped linear networks with particular regard to wideband amplifiers," *J. Appl. Phys.*, vol. 19, no. 1, pp. 55–63, 1948, doi: 10.1063/1.1697872.

[14] K. Banerjee and A. Mehrotra, "A power-optimal repeater insertion methodology for global interconnects in nanometer designs," *IEEE Trans. Electron Devices*, vol. 49, no. 11, pp. 2001–2007, 2002, doi: 10.1109/TED.2002.804706.

[15] SkyWater Technology and Google, "SkyWater SKY130 open source process design kit," https://github.com/google/skywater-pdk, 2020, accessed Oct. 2026.