An Investigation into System-Level Trimming Issues in On-Chip Nanophotonic Networks

Christopher Nitta, Matthew Farrens, and Venkatesh Akella
University of California, Davis
Davis, CA 95616
Email: cjnitta@ucdavis.edu, mkfarrens@ucdavis.edu, akella@ucdavis.edu

Variations in temperature cause the microrings used in nanophotonic interconnects [1] to change their resonant wavelengths - a problem which can be addressed by heating (which moves the resonant frequency towards the red) and/or current injection (which moves it towards the blue). Correcting the drift in resonant frequency in this manner is called "trimming", and estimates indicate it will comprise a significant portion of the total photonic power - for example, in [2] Ahn estimates that ∼26W is necessary for trimming in the Corona network [3], which is ∼54% of the estimated ∼48W total network power.

Thus far researchers have not examined the practical implications of using trimming in an on-chip nanophotonic networks comprised of hundreds of thousands of rings. In particular, the energy required to perform trimming has been modeled as a per ring fixed cost [2], [4]. In order to evaluate the system-wide power and thermal characteristics of on-chip optical networks, we developed a simulator called Mintaka. Using this simulator we show that at the system level, using a fixed cost model is inappropriate - our simulations show that the power required to maintain temperature using heating is not dependent on microring count and thermal sensitivity, and is instead dictated by the die area, ambient temperature, and rate at which heat can be transferred from the die to the ambient. We show, for example, that a nanophotonic network using roughly half as many microring resonators (∼270K) as our baseline (∼524K) saves a maximum of 2.7mW on heating (equivalent to ∼0.01% of total system heating power).

Some researchers have proposed networks [4] that use only thermal trimming. This can be problematic, because the fixed cost model ignores the fact that the microrings must be designed to initially resonate at a temperature higher than the core will ever achieve (since there is no mechanism for shifting the rings back towards the blue). It is not clear how the extra heat that is being fed to the rings will affect the operation of the cores.

Other network models assume the use of current injection to move the resonance frequency towards the blue, which can be used to compensate for excess heating. In this work we show that current injection can very quickly lead to thermal runaway (because injecting current causes the rings to heat up, requiring more and more current that must be injected to overcome the heating), and that a very narrow Temperature Control Window (TCW) of about 1 degree Kelvin must be maintained in order for the rings to work as desired (which from a practical standpoint will be extremely difficult to accomplish). However, by incorporating additional microrings on either end of the spectral range and exploiting the group drift property of co-located rings [5], it is possible to create a Sliding Ring Window

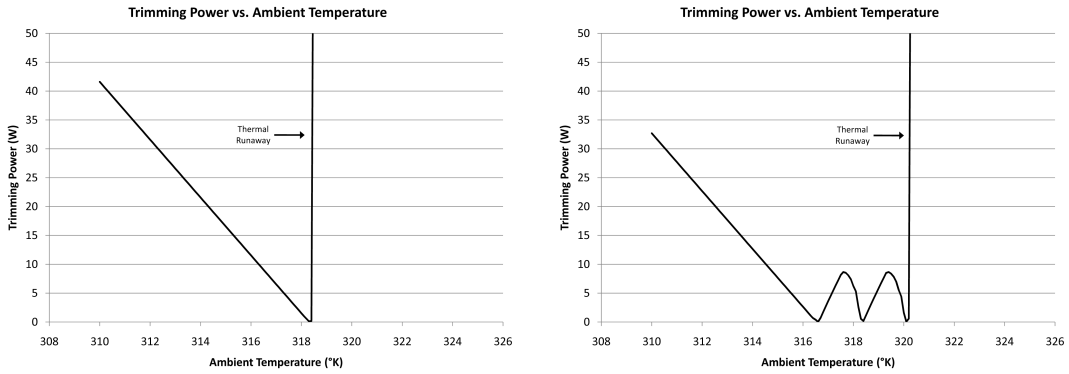

(a) Standard Network

(b) SRW Network

Fig. 1: Trimming Power (W) vs. Ambient Temperature (°K) for Standard Network (a) and Network Using SRW (b)

(SRW) scheme which can significantly increase the TCW, and thus increase the feasibility of using trimming in real systems.

The current injection positive feedback is avoided with the SRW by providing a higher temperature/lower trimming power state (instead of higher temperature always requiring higher trimming power). Figure 1(a) shows the trimming power required as a function of ambient temperature. The current injection positive feedback can be seen in the near vertical line labeled as thermal runaway. Figure 1(b) shows the trimming power for the same network using SRW - the two wave-like patterns represent the maximum current injection trimming before the rings become hot enough to resonate at the next channel.

We also investigate the use of partially athermalized rings (based on techniques such as PMMA cladding), and show that they can alleviate but not eliminate the problem. Athermalization techniques only minimally reduce trimming power when heating is used, but can dramatically reduce trimming power when current injection is applied. While the athermalization techniques reduce the thermal sensitivity of the microring resonators, they do not change the power required to maintain a minimal temperature (although the minimal temperature that must be maintained is slightly lower for the partially athermalized microrings than for the baseline). On the other hand, when the rings are too hot and current injection is required, the partially athermalized microrings substantially outperform their more thermally sensitive counterparts. Using partially athermalized rings in combination with SRW works even better - a 20°K target TCW can be obtained with a trimming budget below 12.2W (much lower than the 51.6W average when heating is used alone).

As we discovered in this work, it is not clear that large-scale networks with over a hundred thousand rings (as proposed in [3] and [6]) can be continuously trimmed in a stable manner; therefore, malfunctioning microrings may need to be *tolerated* instead of corrected (as is done in areas such as hard-disks, flash-memory, wireless communication, and small geometry DRAMs). In order to develop resilience techniques, we need information on the types of faults that are likely to occur in the optical domain. Unfortunately, since the fabrication of nanophotonic components is still in the nascent stage, there is very little in the literature on either the nature of defects or how to model them. Therefore, we derive a set of simple models for faults due to malfunctioning microring resonators, based on the underlying physics of the optical devices. We classify faults into one of two categories: *interfering* or *non-interfering*. Through the use of our fault model we discovered that the bit errors occurring from *non-interfering* faults can be made unidirectional if the correct transmitter/receiver structures are chosen, and we also observed that *interfering* faults result in a unique behavior, one that does not appear to have a direct analogue in the electrical domain.

Our work has shown that the energy required to shift the resonance to the red via heating is not directly dependent of the number of rings, and shifting the resonance to the blue using current injection can quickly lead to thermal runaway. Fortunately, there are techniques that can be used to deal with this problem. Our work on the fault model shows that the choice of transmitter/receiver structures, in combination with intelligently selected encoding schemes, can make the photonic link resilient to any number of *non-interfering* faults. This is key information for the architect to have, because a more realistic estimation of the energy required to trim a photonic network and the manifestation of potential faults is of critical importance when deciding which network topology to use.

REFERENCES

[1] M. Lipson, "Guiding, modulating, and emitting light on silicon-challenges and opportunities," *Lightwave Technology, Journal of*, vol. 23, no. 12, pp. 4222–4238, Dec. 2005.

[2] J. Ahn *et al.*, "Devices and architectures for photonic chip-scale integration," *Applied Physics A: Materials Science & Processing*, vol. 95, pp. 989–997, Jun. 2009.

[3] D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," in *ISCA '08: Proceedings of the 35th International Symposium on Computer Architecture*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 153–164.

[4] A. Joshi *et al.*, "Silicon-photonic clos networks for global on-chip communication," in *NOCS '09: Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*. Washington, DC, USA: IEEE Computer Society, 2009, pp. 124–133.

[5] C. Nitta *et al.*, "Addressing system-level trimming issues in on-chip nanophotonic networks," in *High Performance Computer Architecture, 2011. HPCA 2011. IEEE 17th International Symposium on (to appear)*, Feb. 2011.

[6] Y. Pan *et al.*, "Firefly: illuminating future network-on-chip with nanophotonics," *SIGARCH Comput. Archit. News*, vol. 37, no. 3, pp. 429–440, 2009.