Addressing 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

Abstract

The basic building block of on-chip nanophotonic interconnects is the microring resonator [14], and these resonators change their resonant wavelengths due to variations in temperature - a problem that can be addressed using a technique called "trimming", which involves correcting the drift via heating and/or current injection. Thus far system researchers have modeled trimming as a per ring fixed cost. In this work we show that at the system level using a fixed cost model is inappropriate - our simulations demonstrate that the cost of heating has a non-linear relationship with the number of rings, and also that current injection can lead to thermal runaway. We show that a very narrow Temperature Control Window (TCW) must be maintained in order for the network to work as desired. However, by exploiting the group drift property of co-located rings, it is possible to create a sliding window scheme which can increase the TCW. We also show that partially athermal rings can alleviate but not eliminate the problem.

1 Introduction

Computer architects have long been intrigued by the potential of optics, but a variety of barriers (difficulty of storing photons, cost of converting from photons to electrons and back, fabrication technology incompatibilities, etc.) have stood in the way of their use in conventional designs. Recently, however, this has begun to change. In particular, the ability to fabricate microring resonators and optical waveguides on silicon [25] has enabled the use of optics for both on and off-chip communication [14] [17] [16]. Using an external laser and a comb generator, researchers are now able to create dozens of wavelengths and steer them around on-chip using these resonators, creating optical interconnection networks [27] [3] [29].

This capability will be particularly important in future designs as the number of processors on a chip continues to climb. These multiprocessors on a chip will require high bandwidth communication networks, and electrical networks are not likely to scale up well, primarily for latency and power consumption reasons. For example, the 16-tile MIT RAW and Intel Polaris interconnection networks consume 36% and 28% of total chip power, respectively [26]. As the size of an on-chip network grows, this problem will only become worse. Fortunately, optical networks [15] have less signal crosstalk, lower power loss, and higher switching speeds than electrical networks, so they are ideal candidates for use in future large scale chip-level multiprocessors.

These facts have lead researchers from HP [27], Cornell [12] [4], Northwestern [19], Columbia [24] [23] [22] and MIT [3] [10] to propose different topologies and arbitration/flow control schemes for such networks. The details of each design vary, but what they all have in common is that photons are generated by an off-chip laser and routed around the chip through silicon waveguides using photonic resonators (anywhere from a few hundred to hundreds of thousands). However, because the refractive index of Si changes with temperature, microring resonators are very susceptible to temperature-induced changes to their resonance wavelength. And in a multicore processor there is a potential for significant variations in temperature, so this per-ring drift in resonance must be addressed.

One approach to dealing with this problem is to use a technique called *trimming* - correcting the drift in the resonant wavelength using heating (which shifts the resonant frequency towards the red) and current injection (which causes a shift towards the blue). Trimming gives the architect the ability to precisely and adaptively control the wavelength at which a ring resonates - however, it will also increase the overall power consumption of the network. Researchers have thus far modeled trimming as a fixed additional cost and treated it as little more than a nuisance in terms of additional power consumption overhead, but the practical issues underlying the system level implementation of trimming, especially in the context of networks that contain hundreds of thousands of rings, have not been studied.

As we began this work, we were faced with a number of

questions. For example, is a network of hundreds of thousands of microrings by itself thermally stable? What is the relationship between die size, number of rings, ring density and trimming? Is trimming at the system level even possible, or is thermal runaway the inevitable result? Can the problem be solved by using partially *athermalized* rings, which employ PMMA (polymethyl methacrylate) as an upper cladding [30] to reduce the thermal sensitivity of rings?

In order to attempt to find answers to these questions, we developed a tool called Mintaka to model the thermal and power issues in nanoscale photonics networks. Mintaka is based on Booksim [5] and Orion [11] [28] (popular on-chip network simulators) and Hot-Spot [9] [8] - a widely used tool to study the temperature effects on the microarchitecture of a processor. Using this tool we were able to show that nanophotonic networks themselves are thermally stable, i.e., the micro resonators that comprise the network do not experience temperature fluctuations in spite of the fact that a significant amount of laser power is being pumped into the network. In addition, we found that preventing thermal runaway requires extraordinarily tight control of the system level ambient temperature (less than one degree Kelvin). This is a serious problem, since these nanophotonic networks reside in a very hostile thermal environment, because of the large number of processor cores, memory controllers, and caches (all potentially operating at several GHz) – controlling the temperature of such a system to within one degree Kelvin will be extremely challenging.

Fortunately, we also found that thermal drift in rings is correlated (meaning a set of co-located rings used to implement a communication interface will all drift the same way). We call this property *group drift*, and show that by using additional resonators a sliding window scheme can be created, allowing a set of rings to be trimmed as a group instead of one ring at a time. This group trimming ability is particularly important because it means the designer does not have to route control signals to each ring individually in order to direct the current injection and heating - doing so at the system level would be an implementation nightmare, if there are hundreds of thousands of rings. There is nice trade-off between the size of the Temperature Control Window (TCW) and the number of additional rings necessary, which allows the architect to choose the number depending on the application of the system and where the system resides (the cooling infrastructure available).

Our investigation also revealed that when using heating the die size is more important than the actual number of rings, while for current injection the ring density and number of rings are both influential factors. Finally, we examined the potential of partially athermalized rings and found that they do not *eliminate* the thermal runaway problem, although they alleviate it somewhat.

In Section 2 we describe the necessary background and motivation for the proposed research and in Section 3 we describe work related to this paper from research literature. This is followed by the description of the Mintaka tool in Section 4, the details of the experimental setup in Section 5, and results and discussion in the subsequent sections.

2 Background

Microring resonators are designed to resonate when presented with specific individual wavelengths and remain quiescent at all other times. This ability is exploited to realize the basic building blocks of photonic interconnects - such as add/drop filters and switches. Figure 1(a) shows a passive microring that is biased during fabrication to extract only λ_1 from the incoming waveguide and steer it down a perpendicular waveguide. Modulating specific wavelengths requires an active microring resonator, which uses the presence or absence of an electrical current to extract or ignore a specific wavelength. Figure 1(b) illustrates an active microring resonator modulating wavelength λ_1. If the electrical current is present, λ_1 is extracted and sent down waveguide II – if there is no current applied, λ_1 will continue down waveguide I unaffected.

The wavelengths individual microrings respond to are set during fabrication. However, variations in fabrication may require that certain microrings have their resonance frequency adjusted up or down slightly. Passive or post fabrication techniques such as using UV light to modify microrings have been proposed [20] [13], which have the distinct advantage of not requiring any additional power once the rings have been moved into their correct spectral position. However, it is not clear how practical this approach is at the system level, given the number of rings involved.

The refractive index (n) of silicon changes due to changes in ambient temperature (ΔT), which can be modeled as - $\Delta n \approx 1.84 \times 10^{-6} \times \Delta T$. As a result microring resonators are very sensitive to temperature and drift spectrally approximately 0.09nm/°C. The resonance frequency of a microring can be changed by heating it[1], which causes a shift towards the red end of the spectrum, or by electrically injecting current (to shift the resonance towards the blue) [1]. This dynamic modification of resonance frequency is referred to as "trimming", and is illustrated in Figure 2.

3 Related Work

Within the research community there has been a growing interest in harnessing the benefits of optical interconnects for addressing the shortcomings of electrical interconnects [15, 16]. The idea of using microring resonators

[1] This effect can be implemented in devices, for example, with thin-film platinum surface heaters near waveguide sections [7]

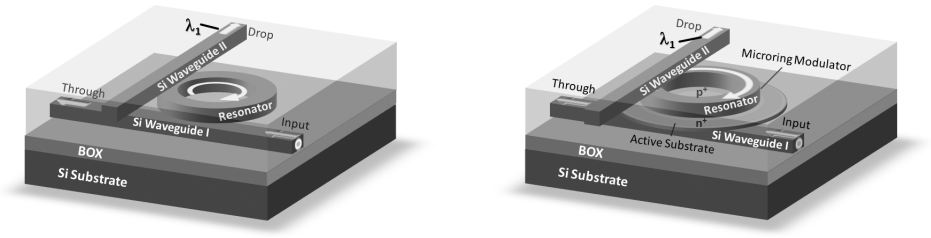

(a) Filter PASSIVE Microring (b) Modulator ACTIVE Microring

Figure 1: Microring Resonators

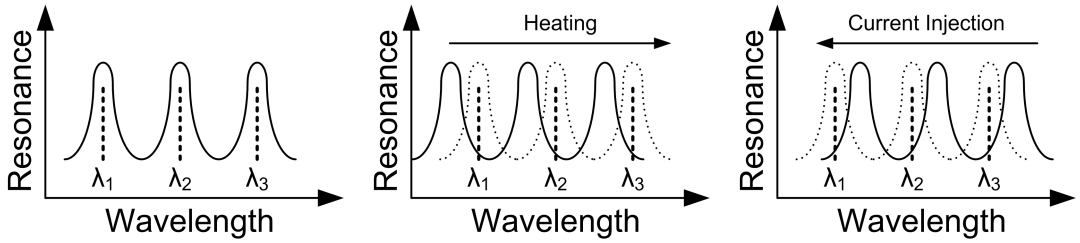

(a) On Resonance (b) Red Shift Needed (c) Blue Shift Needed

Figure 2: Microring Resonance vs. Wavelength for On Resonance (a), Heating (b), and Current Injection (c)

for modulation in on-chip optical networks has been around for some time, although as stated in [14], "... the disadvantage of using resonators for modulation is the high temperature sensitivity of the device." However, as discussed before, this can be compensated by trimming. In current literature, researchers typically estimate the required microring trimming power by multiplying the estimated average trimming power per microring by the number of microrings [1] [18] [10]. A global estimate for microring heating was provided in [19], and we assume that a similar approach was used to derive the estimate

In [27] HP researchers describe a 64x64 WDM based crossbar (called Corona) for a 256-core CMP and Ahn in [1] estimates that a total of ∼26W is necessary for trimming of the Corona network (which is ∼54% of the estimated∼48W total network power.) Cornell researchers described a bus-based scheme to connect clusters of processors in [12], and more recently propose a hybrid opto-electronic on-chip network called Phastlane [4] that uses a low complexity nanophotonic crossbar supported by an electrical network for buffering and arbitration. Neither [12] or [4] explicitly discuss the required power for trimming of the nanophotonic networks.

MIT and Berkeley researchers [10] describe a multistage Clos network using a mixture of electronic routers that are connected by WDM based photonic links. A fixed thermal power was assumed in [10] to tune the microrings over a 20°K range. The researchers also use a different set of constraints - they assume the microrings are 10μm in diameter, and place the rings on the same die as the cores. These two constraints lead to their conclusion that optical crossbar designs are impractical. However, microrings can be as small as 3μm and still function correctly, and if performance is important it is possible to implement the communication network on another level of a 3D design as illustrated in [27]. The Clos design reported in [10] uses only thermal trimming, which means the rings must be designed to operate at a temperature higher than that which the cores could ever raise them (since if the temperature gets above that, there is no mechanism for shifting the rings back towards the blue.)

The authors in [24] [23] propose a photonic 2D torus network that employs an electrical network for arbitration and flow control, but no estimate for trimming power of the nanophotonic network is explicitly discussed. Firefly [19] is another hybrid opto-electronic network proposal that uses an electrical network for intra-cluster communication and a nanophotonic crossbar for inter-cluster communication. A global estimate of 3.6W is assumed for microring heating in the Firefly network. FlexiShare network crossbar [18] uses a token stream for arbitration and credit sharing. A 1μW per ring per °K with a 20°K tuning range was assumed for trimming power in FlexiShare.

These approaches to estimating trimming power are reasonable given the absence of a full integrated power/thermal simulation. And it could be argued that theoretically any microring based network will not require any trimming power if operated under ideal conditions. Temperature fluc-

tuations in the environment external to the chip will occur in the real world, and it is vital to understand how temperature fluctuations affect the amount of power necessary to support trimming.

4 Mintaka

In order to evaluate the power and thermal characteristics of on-chip optical networks we developed a simulator called Mintaka[2]. The photonic power estimates were developed using a link loss approach similar to that done in [1], with relevant data gathered from the literature and extrapolated from laboratory test results. Hot-Spot 5.0 was chosen to perform the thermal analysis since it can be compiled as a library, allowing it to be integrated into Mintaka [9] [8]. The core of Mintaka – like Hot-Spot – has been designed to compile as a library for future integration into other network performance simulators.

The electrical components used in Mintaka were constructed in a manner similar to that used in ORION 1.0, although electrical technology data such as transistor sizing and wire capacitances were taken from CACTI 6.5 (for technology parameters from 90nm to 32nm), and ITRS 2009 [21] (for beyond 32nm). The energy required per transition for each sub-component was calculated, and the number of transitions per sub-component was maintained as an integer in order to mitigate potential floating point round off errors that could have occurred during long simulation runs if a floating point only implementation had been applied. The equations for calculating transistor sizing, capacitance, and other low level electrical characteristics were also taken directly from CACTI 6.5. Static power loss was also accounted for in Mintaka, since CACTI 6.5 calculates static power values as a function of temperature; the temperature of each floor-plan unit was passed into the simulated network in order to ascertain the static power loss.

The link loss calculation starts at a photodetector and works backwards towards the source, adding the attenuation losses along the way; this is done for the worst case path for all networks. Attenuation sources include the photodetector, waveguide, waveguide intersection, waveguide bends, grating coupling, on-resonance rings, and off-resonance rings. Once the total attenuation has been calculated the power output necessary to switch the photodetector at the desired rate is determined, and given these two values the required laser power per wavelength is a simple calculation of $P_{PD}10^{\frac{A}{10}}$ (where P_{PD} is the required power at the photodetector and A is the attenuation of the path.)

Once the minimum laser power has been calculated for the worst-case path, multiplying by the total number of wavelengths provides the total amount of laser power required by the network. Once this value has been determined, the energy lost (absorbed) by each optical component can then be calculated, starting at the source and working downstream towards the photodetector. The energy loss of each sub-component is maintained as a floating point number and the number of times the loss occurs is maintained as an integer, for reasons described earlier.

Mintaka was validated by comparing its link loss calculations to those published for Corona [1], when the same parameters were input. Mintaka calculated an attenuation loss of 13dB for Corona, which matched the published values – a 3dB photodetector loss, plus the 11dB transmission loss, minus the 1dB coupling loss (the coupling loss is not included since Mintaka calculates the on-chip required laser power only). We also validated Mintaka by hand-calculating the link losses of some other optical network configurations and comparing those results to the ones generated by Mintaka. In addition, the total power of the floor plan units was compared to the laser power plus the estimated electrical switching and static power loss for all the simulations. None of the simulations deviated from the expected power consumption, given the required laser power and the input traffic pattern.

In Mintaka, the floor-plan layout for each network is integrated into the electrical/optical power/sizing calculations. The floor-plan layout is necessary since both the optical and electrical power requirements depend upon the distance which the signals must travel, and the minimum size of some sub-components is dependent upon the power requirements.

The power consumed in the network, both electrical and optical, was maintained for each floor-plan component. This floor-plan, floor-plan power, and floor-plan temperature data was passed to the Hot-Spot library (using appropriate thermal constants for SOI) to calculate the new floor-plan temperatures. The Hot-Spot steady state solver was used to determine the updated temperatures for the floor-plan components. The updated floor-plan temperatures were used by Mintaka to provide more accurate power consumption numbers for the next iteration, since some components (such as static power loss and trimming power) are a function of temperature. The iterative process continued until Mintaka/Hot-Spot converged on a steady state solution, or a thermal runaway was detected.

5 Experimental Setup

The base architecture we used in our simulations was a 64 node crossbar network with a 64-bit data path between nodes, built using 22nm technology. The nodes were assumed to operate at 5GHz and capable of generating and consuming one 128-bit flit per cycle. We chose a crossbar as

[2]Mintaka -faintest of the three belt stars in ORION, which hints at the faint influence of the ORION simulator.

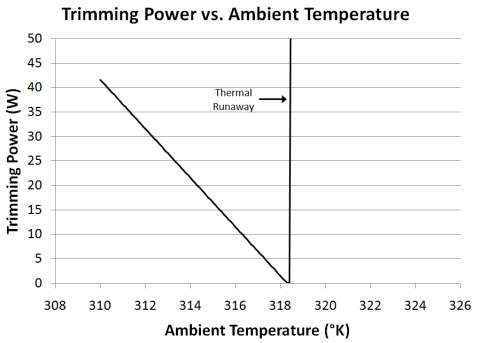

Figure 3: Trimming Power (W) vs. Ambient Temperature (°K) for Network

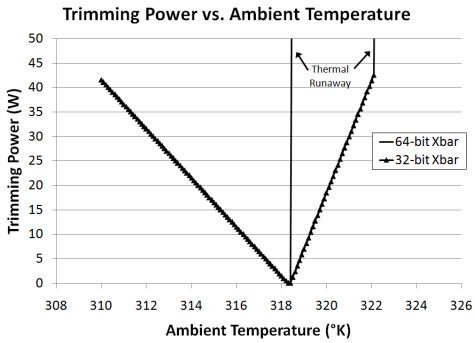

Figure 4: Trimming Power (W) vs. Ambient Temperature °K for 64-bit and 32-bit Network

our base model because it uses a similar number of microring resonators as other proposed on-chip optical networks, and in addition a very detailed floor-plan was available. In our case, the on-chip network consists of ∼524K microring resonators and occupies an entire level of a 3D stacked processor design, with an area of 484mm^2. The workload used is synthetic random traffic pattern, since the goal of this work is to determine the power/thermal sensitivity of large on-chip optical networks and not to analyze the performance.

6 Stability Issues with Trimming

The first step in our investigation was to determine if the network itself was thermally stable (does the network generate too much internal heat even in the absence of trimming?). The network settled at less than half a °C above the ambient. The initial simulation also yielded the "ideal" operating temperature for the network, 45.38°C. This value was used to estimate the required trimming power for the subsequent experiments - as the temperature of the microrings drops below the "ideal" temperature heating is required, and as the temperature climbs above the "ideal" current injection is applied.

The first trimming experiment consisted of a sweep of the ambient temperatures from 310°K to 325°K. The microring thermal sensitivity was assumed to be 0.09nm/°C, and the channel separation was assumed to be 0.16nm with a 2% tolerance for resonance.

The required trimming power was initially assumed to be 130μW/nm [1] for current injection (blue shift) and 240μW/nm [6] for heating (red shift). A red shift is not achieved until the microring reaches the appropriate temperature – however, using a fixed value of power per nm to accomplish red shift does not make sense, so that approach was quickly abandoned in favor of a closed loop solver that determined the actual power required to maintain a minimum required temperature. Figure2(b) demonstrates a resonance shift towards red for heating. A resonance blue shift can be directly achieved by current injection (see Figure2(c)), but Resonance deteriorates under high current, meaning that blue shift using current injection will be capable of less than 1nm of shift. It was assumed that a maximum of 1nm of blue shift could be achieved through current injection.

The results of the trimming power sensitivity simulations can be seen in Figure 3. The network required approximately 5.1W of heating for every degree the ambient temperature dropped below the design target of 318.15°K (45°C). In addition, the network became thermally unstable within a one degree increase above the optimal ambient temperature - this is because the current injection becomes a positive feedback system (current injection heats the rings, heat in the rings causes red shift, requiring more current injection, etc.). Even though the network may settle at a steady state when ambient temperature rises, the required shift was beyond 1nm.

The thermal runaway observed in the baseline model led us to look at a 32-bit version of the network operating at twice the frequency. The 32-bit network contains ∼270K microrings, and it has been assumed that reducing the number of microrings will reduce the required trimming power. The results of the trimming power sensitivity simulations using the 32-bit network are shown in Figure 4. As one might expect the power required for current injection is lower for the 32-bit network, and the network is slightly more stable than the 64-bit version (becoming thermally unstable within four degrees above the optimal ambient temperature.) What seems surprising is that the required power for heating is the almost identical for both networks. In situations where heating is required for trimming, the amount of trimming power required appears to have a non-linear relationship with microring count. It was this observation that led us to investigate varying the die area.

Figure 5 shows the required trimming power for the 64-bit network with varying die areas. The die areas shown

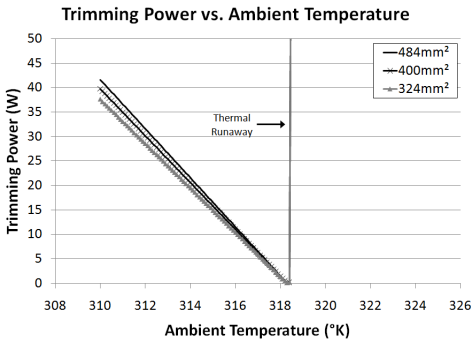

Figure 5: Trimming Power (W) vs. Ambient Temperature (°K) for 64-bit Network with 484mm^2, 400mm^2, and 324mm^2 Die Area

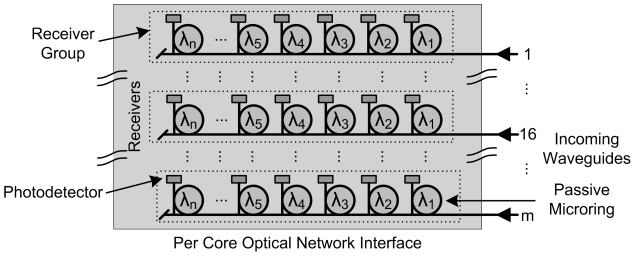

Figure 6: Receive Section of a Node Illustrating Microring Grouping

are 484mm^2, 400mm^2, and 324mm^2 (22mm, 20mm, and 18mm squares). The trimming power required for heating is clearly related to (although not directly proportional to) the die area. This might at first seem somewhat counter-intuitive (one might expect there to be some correlation between microring count and trimming power), but upon further analysis it is clear that the power required to maintain a given temperature is dominated by the area that has to be heated (the total die area) and not the number of microrings. Using simple thermodynamic analysis, the required trimming power for heating should equal the rate at which heat can be removed from the die - and as the die area is reduced, so is the rate at which heat is removed from the die.

The results of the trimming power necessary for current injection is not visible in Figure 5 because of the thermal runaway. Further analysis of the simulation results shows that the trimming power for current injection appears to have a direct correlation to average microring density. As the die area shrinks there is an increase in the required current injection trimming power, because the heating that occurs during injection is spread among the rings. This can be seen in Figure 4, since the 32-bit network has a dramatically lower microring density than the 64-bit equivalent.

These results show that reducing the number of microrings or the microring density can reduce the trimming power necessary for current injection, but only reducing the die area is effective at lowering the required trimming power for heating. We define the Temperature Control Window (TCW) as the range of temperatures within which the network must be kept in order to remain within a given trimming budget and prevent thermal runaway. The TCW for the 64-bit network is less than 1.1°C and 4.1°C for a trimming budget of 5W and 20W, respectively.

If heating is the only form of trimming implemented to address thermal drift then power will always be used to heat the microrings unless the system is being operated in an environment on the edge of the operational range. Our simulation results imply that a 20°K TCW as suggested in [18] and [10] would require a maximum trimming power of 103W (51.6W average) for the 484mm^2 die area of the simulated network. Even for the 400mm^2 die area assumed in [10] the maximum power for heating over the 20°K TCW is 98.9W (49.4W average). According to [2] the recommended temperature range of data centers is 18-27°C. This is a 9°C temperature range for the room ambient temperature - the temperature range the chip will experience will most likely be larger. Thus, a 20°K TCW is justifiable given the air conditioning recommendations of ASHRAE [2].

7 Increasing the Thermal Window

The granularity of the floor-plan units used in the previous results was set to a transmitter or receiver *group* (i.e. 64-bits), and it was assumed a constant temperature was maintained across the entire group of colocated rings used to implement a transmitter or receiver. Figure 6 illustrates the receive section of a node and shows how microrings were grouped together into single floor-plan units.

The assumption that rings drift thermally as group was made because finding a thermal solution in the simulator for approximately half a million floor plan units is not currently practical. In order to ascertain the accuracy of this assumption, the simulator was modified to keep track of each individual microring in a specific transmitter or receiver group. Simulations were run over sweeps of ambient temperatures ranges with trimming disabled (having trimming on would defeat the purpose of the test), and transmitter and receiver groups to be studied were chosen strategically (i.e. corners, centers, and edges).

The largest intra group temperature delta observed was 5.13e-4°C, which corresponds to less than 0.03% channel separation when assuming a wavelength separation of 0.16nm and 0.09nm/°C thermal drift. This result is very encouraging for a number of reasons - it means that our previous results are valid because our basic assumption was not erroneous, it means that the amount of circuitry necessary to support trimming can be greatly reduced since microrings can be trimmed as a group instead of needing to be trimmed

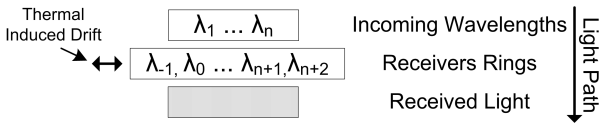

Figure 7: Node to Node Drift Resilience

individually, and perhaps most importantly it means that we can extend the TCW by adding resonators at each end of the incoming wavelength spectra.

7.1 Sliding Ring Window

The results presented so far indicate that the Temperature Control Window (TCW) will be impractically small for realistic trimming power budgets, and that a reasonable TCW will require potentially unreasonably large trimming budgets. In order to overcome this problem, we propose incorporating additional rings on either end of the spectral range in order to maintain the same usable data path width. The additional microrings will create a Sliding Ring Window (SRW), exploiting the fact that the entire group will slide the same amount spectrally in either direction. This concept is shown in Figure 7, and works as follows: Current injection is used to maintain the spectral position of the entire group of rings (see Figure8(b)) until the rings become so hot that they will naturally resonate at the next (red shifted) frequency (see Figure8(c)). At this point the current injection is turned off and the entire group begins resonating one wavelength over. As the rings cool current injection can be reapplied to correct the spectral position of the microrings to the previous (blue shifted) channel. The existence of the additional rings prevents the current injection positive feedback system from thermally running away by creating a lower power trimming state at a higher temperature. We will designate the number of additional rings in the name, so that SRW-1 indicates 1 extra resonator, SRW-2 indicates there are two extra rings, etc.

Guiding the correct electrical signals to/from the correct microrings of the SRW will require integration with the trimming circuitry. However, the SRW control is expected to be of minimal additional circuit complexity since any trimming system implemented must maintain microrings in the correct spectral position. The control of the SRW is managed locally (a global feedback channel is unnecessary) since the temperature of the microring group is the feedback.

In order to evaluate the effectiveness of the SRW, simulations were run with SRW-2 – increasing the network total microring count to ∼540K. Whenever possible the steady state thermal solver was used, although some of the transitional temperatures could not be solved using the steady state thermal solver since the microrings (by design) cycled back and forth between lower temperature/higher power and higher temperature/lower power states. The transitional temperatures were determined by simulating using the time step solver until the total trimming power converged.

Figure 9 shows the impact of SRW-2 on the amount of trimming power required. The sawtooth (or wave-like) pattern is due to the fact that the microrings require maximum current injection trimming before becoming hot enough to be on resonance at the next channel. The peak current injection trimming power is a function of the channel separation of 0.16nm – the peak trimming power could be reduced if channels were able to be more densely packed, although this would also shrink the TCW. The TCW for the 64-bit network using SRW-2 is ∼5.6°C for a trimming budget of 10W (a dramatic increase from the ∼2.1°C TCW for the 64-bit network without SRW), and the TCW is greater than 7.5°C for the 20W trimming budget discussed in the previous section.

7.2 Increasing TCW with SRW

The SRW mechanism can be expanded to incorporate more than just two additional microrings. Each additional microring per group will result in an additional peak as seen in Figure 9 and will increase the TCW. The separation of the peaks seen in Figure 9 is ∼1.8°C, which corresponds to the channel separation divided by thermal sensitivity (0.16nm / 0.09nm/°C). Achieving the 20°K TCW discussed previously would require roughly seven more microrings (creating SRW-9), raising the total microring count to ∼595K for the 64-bit network.

Expanding the TCW by using additional microrings comes at a cost of increased area, trimming power, and laser power. The increased area needed by the microrings is obvious, but compared to the current microring count and die area of the proposed 64-bit network the additional area for SRW microrings is not a concern. The required laser power increases with each additional microring since the number of off-resonance microrings which light must travel through is increased, but this increase is also not a great concern since the additional attenuation of off-resonance microrings is relatively small (1.5e-3dB). The additional trimming power is likely to be the greatest concern, since each microring added will also need to be trimmed. These additional microrings will cause an increase in the peak current injection power seen in the SRW sawtooths, though the increase in peak trimming power is relative to the number of additional rings. The proposed SRW-9 would have an approximately 10% greater peak power than that of Figure 9, yielding a roughly 9.7W peak.

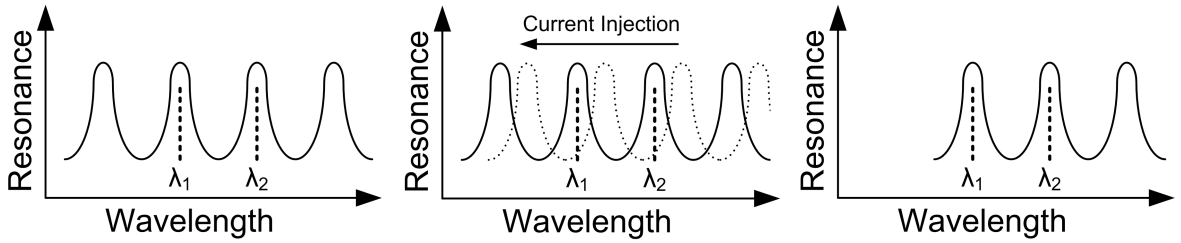

(a) On Resonance (b) Blue Shift Needed (c) Remove Current Slide to Next Ring

Figure 8: Sliding Ring Window Microring Resonance vs. Wavelength for On Resonance (a), Current Injection (b), and Current Removal (c)

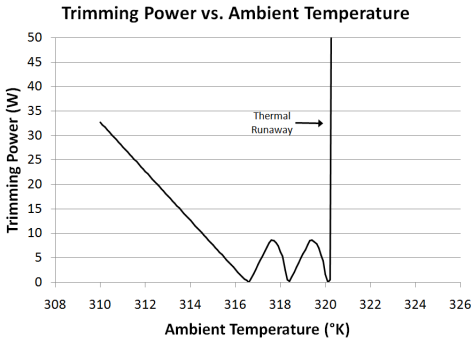

Figure 9: Trimming Power (W) vs. Ambient Temperature (°K) for 64-bit Network Using SRW-2

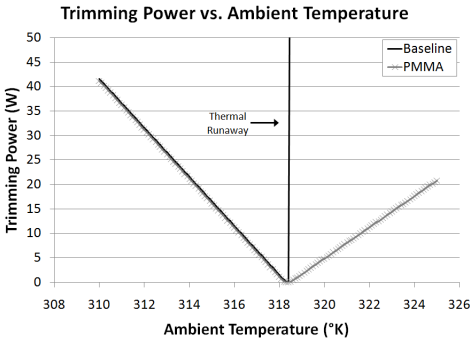

Figure 10: Trimming Power (W) vs. Ambient Temperature (°K) for Baseline and PMMA 64-bit Network

8 Impact of Partially Athermalized Microrings

Another promising approach to increasing the TCW is to use rings which have been clad with polymethyl methacrylate [30]. These partially athermalized rings are substantially less sensitive to temperature variations - unclad rings change approximately 0.09nm/°C, while for PMMA clad rings the change is closer to 0.027nm/°C. In order to analyze their impact on the TCW, the original simulation was rerun using PMMA clad rings instead of unclad ones. Figure 10 shows the results for both the baseline and PMMA-clad 64-bit network. What may be somewhat surprising is that both networks require the same amount of additional heating power for every degree below the ambient temperature. While the PMMA cladding reduces the thermal sensitivity of the microring resonators, it does not change the power required to maintain a minimal temperature (although the minimal temperature that must be maintained is slightly lower for the PMMA-clad microrings than the baseline, which can be seen in the offset of the PMMA-clad line in the figure). On the other hand, when the rings are too hot and current injection is required, the PMMA-clad microrings substantially outperform their unclad counterparts. As the figure shows, only ∼3W of current injection is require for every degree the ambient temperature climbs above the optimal. Thus the TCW for PMMA clad network is less than 3°C and 10.6°C for a trimming budget of 5W and 20W, respectively. However, this still implies that a trimming budget of nearly 39W would be required to meet the 20°K TCW.

Since cladding and SRW are orthogonal techniques, we decided to see how well they would work in conjunction. The simulations using SRW-2 were rerun assuming PMMA-clad rings, and the results are presented in Figure 11. In this figure one can see that the PMMA-clad network using SRW-2 provides a TCW of 19°C with a trimming budget of less than 10.5W. Using athermalized rings, the 20°K target TCW can be obtained with a trimming budget below 12.2W.

Another interesting point to note in the figure is that the peaks in the baseline network are lower than those of the PMMA-clad network. One would expect both the baseline and the PMMA-clad networks to require the same peak potential trimming power - however, the baseline network has a much higher thermal sensitivity, and therefore it oscillates between lower temperature/higher power and higher temperature/lower power more often than does the PMMA upper clad network. Thus, the baseline peak is actually

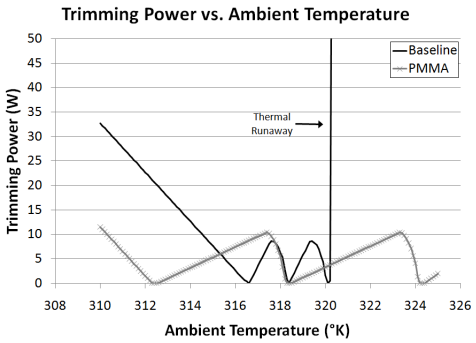

Figure 11: Trimming Power (W) vs. Ambient Temperature (°K) for Baseline and PMMA 64-bit Network with SRW-2

smoothed out by its thermal sensitivity.

9 Architectural Design Guidelines

Based on what we have learned so far, we can make some observations that architects must keep in mind when designing large nanophotonic systems:

1. The power required to maintain temperature using heating has a non-linear relationship with microring count and thermal sensitivity, and is more affected by the die area, ambient temperature, and rate at which heat can be transferred from the die to the ambient.

2. Trimming using current injection is highly sensitive to microring count, density, and the thermal sensitivity of the rings, and thus thermal runaway can happen very easily - networks with ∼524K microrings and a 484mm^2 die area experience runaway within a change of a single degree. This is significant, because optical network topologies have been proposed that employ higher microring counts and densities than those analyzed in this paper. The use of current injection in these networks will only be feasible if they employ techniques like PMMA cladding or the Sliding Ring Window.

3. Any microring based nanophotonic network cooling system will need to be carefully designed, since it impacts the efficiency of the trimming system. The design of a nanophotonic network that trims using only heating will obviously benefit from a weaker cooling system. A cooling system that only removes the amount of heat generated by the laser and modulation power would be ideal, since the microrings could be brought up to temperature and then maintained by the heat generated by the laser and modulation. Conversely, any system that utilizes current injection as a form of trimming will benefit from an efficient cooling system - otherwise, the network will become thermally unstable.

10 Conclusion and Future Work

Up to now in on-chip photonics interconnect literature, the power required to perform trimming has been treated primarily as a fixed cost per ring. In this paper we have shown that this is not the case from either direction - the energy required to shift the resonance to the red via heating has a non-linear relationship with the number of rings, and shifting the resonance to the blue using current injection can quickly lead to thermal runaway. This is key information for the architect to have, because a more accurate and realistic estimation of the energy required to trim a photonic network is of critical importance when deciding which network topology to use.

In order to avoid the thermal runaway problem without having to rely exclusively on the cooling system, we have examined different ways to increase the Temperature Control Window (TCW) - in particular by adding extra rings to create a Sliding Ring Window (SRW), and by using PMMA-clad resonators. We have shown the effectiveness of each technique, and also shown that the a combination of the two works even better.

The work presented in this paper assumed a consistent heat sinking/cooling system and did not incorporate the power/thermal fluctuations from the core layer. Accurately modeling potential hotspots from the core may uncover even more potential trimming issues. Very large microring resonator counts were also assumed in this work, it is possible that the trimming power required for lower density/microring count networks may yield very different results. Communication power dissipated in the microrings (photonic, modulation, etc.) affects the required power for trimming, therefore a fair comparison of networks trimming power is difficult when the required communication power differs.

The trade-offs of channel separation with SRW should be further investigated, our simulations assumed channels with minimum possible separation, there is the potential to widen the TCW at a cost of higher peak trimming power. Widening the channel separation to match the microrings free spectral range also has the potential for SRW scheme without the need of additional microrings. The use of SRW causes the trimming power to be a non-continuous function - therefore, trimming will be most efficient in periodic temperature bands. An investigation of system level techniques to maintain the network within those temperature bands should also be completed.

References

[1] J. Ahn, M. Fiorentino, et al. Devices and architectures for photonic chip-scale integration. *Applied Physics A: Materials Science & Processing*, 95:989–997, June 2009.

[2] ASHRAE. 2008 ashrae environmental guidelines for datacom equipment. Technical report, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., Atlanta, GA, USA, 2008.

[3] C. Batten, A. Joshi, et al. Building manycore processor-to-dram networks with monolithic silicon photonics. In *HOTI '08: Proceedings of the 2008 16th IEEE Symposium on High Performance Interconnects*, pages 21–30, Washington, DC, USA, 2008. IEEE Computer Society.

[4] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi. Phastlane: a rapid transit optical routing network. *SIGARCH Comput. Archit. News*, 37(3):441–450, 2009.

[5] W. Dally and B. Towles. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, San Francisco, 2004.

[6] M. Fiorentino. personal communication about trimming work being done by Andrei Faraon at HP., 2010.

[7] B. J. Frey, D. B. Leviton, and T. J. Madison. Temperature-dependent refractive index of silicon and germanium. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6273 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2006.

[8] W. Huang, S. Ghosh, K. Sankaranarayanan, K. Skadron, and M. R. Stan. Hotspot: Thermal modeling for cmos vlsi systems. *IEEE Transactions on Component Packaging and Manufacturing Technology*, 2005.

[9] W. Huang, K. Sankaranarayanan, et al. Accurate, pre-rtl temperature-aware design using a parameterized, geometric thermal model. *IEEE Transactions on Computers*, 57(9):1277–1288, 2008.

[10] A. Joshi, C. Batten, 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*, pages 124–133, Washington, DC, USA, 2009. IEEE Computer Society.

[11] A. Kahng, B. Li, et al. Orion 2.0: A fast and accurate noc power and area model for early-stage design space exploration. In *DATE*, pages 423–428, April 2009.

[12] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonesi. Leveraging optical technology in future bus-based chip multiprocessors. In *MICRO 39: Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 492–503, Washington, DC, USA, 2006. IEEE Computer Society.

[13] Y. Kokubun, N. Kobayashi, and T. Sato. Uv trimming of polarization-independent microring resonator by internal stress and temperature control. *Opt. Express*, 18(2):906–916, 2010.

[14] M. Lipson. Guiding, modulating, and emitting light on silicon-challenges and opportunities. *Lightwave Technology, Journal of*, 23(12):4222–4238, Dec. 2005.

[15] D. Miller. Rationale and challenges for optical interconnects to electronic chips. *Proceedings of the IEEE*, 88(6):728–749, Jun 2000.

[16] D. Miller. Device requirements for optical interconnects to silicon chips. *Proceedings of the IEEE*, 97(7):1166 –1185, july 2009.

[17] D. Miller, A. Bhatnagar, et al. Opportunities for optics in integrated circuits applications. In *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, pages 86–87 Vol. 1, Feb. 2005.

[18] Y. Pan, J. Kim, and G. Memik. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *High Performance Computer Architecture, 2010. HPCA 2010. IEEE 16th International Symposium on*, Jan. 2010.

[19] Y. Pan, P. Kumar, et al. Firefly: illuminating future network-on-chip with nanophotonics. *SIGARCH Comput. Archit. News*, 37(3):429–440, 2009.

[20] J. Schrauwen, D. V. Thourhout, and R. Baets. Trimming of silicon ring resonator by electron beam induced compaction and strain. *Opt. Express*, 16(6):3738–3743, 2008.

[21] Semiconductor Industry Association. International technology roadmap for semiconductors 2009. 2009.

[22] A. Shacham and K. Bergman. Building ultralow-latency interconnection networks using photonic integration. *IEEE Micro*, 27(4):6–20, 2007.

[23] A. Shacham, K. Bergman, and L. P. Carloni. The case for low-power photonic networks on chip. In *DAC '07: Proceedings of the 44th annual Design Automation Conference*, pages 132–135, New York, NY, USA, 2007. ACM.

[24] A. Shacham, K. Bergman, and L. P. Carloni. On the design of a photonic network-on-chip. In *NOCS '07: Proceedings of the First International Symposium on Networks-on-Chip*, pages 53–64, Washington, DC, USA, 2007. IEEE Computer Society.

[25] R. A. Soref and B. R. Bennett. Electrooptical effects in silicon. *IEEE Journal of Quantum Electronics*, 23:123–129, Jan. 1987.

[26] S. Vangal, J. Howard, et al. An 80-tile sub-100-w teraflops processor in 65-nm cmos. *Solid-State Circuits, IEEE Journal of*, 43(1):29–41, Jan. 2008.

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

[28] H.-S. Wang, X. Zhu, et al. Orion: A power-performance simulator for interconnection networks. *Microarchitecture, IEEE/ACM International Symposium on*, 0:294, 2002.

[29] Q. Xu, B. Schmidt, et al. Micrometre-scale silicon electro-optic modulator. *Nature*, 435(7040):325–327, 2005.

[30] L. Zhou, K. Okamoto, and S. Yoo. Athermalizing and trimming of slotted silicon microring resonators with uv-sensitive pmma upper-cladding. *Photonics Technology Letters, IEEE*, 21(17):1175–1177, Sept.1, 2009.