

Low-power Integration of On-chip Nanophotonic Interconnect for High Performance Opto-electrical IC

Duo Ding and David Z. Pan

Electrical and Computer Engineering Dept., the Univ. of Texas at Austin, Austin, TX, U.S.A;

ABSTRACT

In this manuscript we study the potentials of nanophotonics on-chip integration and propose a set of automation methodologies to construct low power on-chip interconnect with flexible geometry shapes. We show that with such techniques, a systematic design aid environment can be developed to generate optimized integration configurations meanwhile honoring complex sets of photonic device constraints. Due to their unique characteristics, not only do these techniques benefit the optimization of on-chip photonic networks, but also they can be efficiently applied to build low-power high-throughput application specific ICs with opto-electrical interconnection.

Keywords: On-chip Nanophotonics Interconnect, Low-power Integration, Wavelength Division Multiplexing

1. INTRODUCTION

As semiconductor technology road map extends, the development of future high performance low power silicon systems is facing many critical challenges. Amongst them, on-chip interconnect is starting to play more and more important roles due to: (1) growing interconnect versus gate delay ratio; (2) potentially longer global interconnects due to higher levels of on-die functional integration; (3) timing closure becomes harder for complex design; (4) economical and power-efficient interconnect design becomes more and more difficult.

To address the interconnect challenges, various technologies are proposed as potential solutions.¹⁻³ Among them, nanophotonics devices and interconnect demonstrate unique potentials for constructing high speed and low power on-chip communication links.³⁻⁹ As recent advances in nanophotonics fabrication makes individual devices smaller in footprint and better in performance, photonics interconnect is expected to qualify as a feasible on-chip interconnect solution with high integration density, consisting of nano-scale modulators, detectors, couplers, switches, waveguides and WDM (Wavelength Division Multiplexing) components, etc. With these building blocks, device modeling/roadmapping⁷ and on-chip integration¹⁰⁻¹⁶ of nanophotonics have been studied and exploited in many ways that changed the traditional IC design methodologies.

Nanophotonics has been recently employed in the research and development of on-chip networks and architectures^{11,14,17} to generate dynamic data traffic routing with high-throughput Time Division Multiplexing and WDM on-chip nano-photonic links. These works show promising potential for nanophotonics to address inter-core and memory bandwidth limitations. Lately, studies have been carried out to control the temperature dependency of nanophotonic devices at both fabrication^{18,19} and network level^{20,21} to assist the optimization of thermal reliability and power efficiency.

At circuit (physical) design level, studies for automatic and efficient nanophotonics interconnect planning have been limited, especially under complex device constraints that are usually hard to satisfy via custom design. An early design automation work¹⁰ was presented where straight-line single-channel optical waveguides were deployed into a system-on-package under timing driven metrics. However, physical device characterizations were not considered for modulators, waveguide, photo-detectors, and important issues such as optical link configuration, loss figure, thermal reliability and signal integrity were not properly explored. Physical-layer effects (loss, power) are modeled and applied to photonics Network-on-Chip performance evaluation,¹⁵ yet for a complex design, it could still be difficult to construct photonics architectures with optimal performance meanwhile satisfying all device constraints. At circuit implementation level, our early works presented a parameterized nanophotonics interconnect library^{12,13} together with a synthesis framework for low power on-chip optical routing through

Further author information:

Duo Ding: E-mail: ding@cerc.utexas.edu, David Z. Pan: E-mail: dpan@ece.utexas.edu

architectural/physical co-design. However, thermal reliability and signal multiplexing mechanisms were not exploited.

In this paper, we study the characteristics and requirements of on-chip nanophotonic interconnect construction for low power integration. We develop a set of design automation environments that allow us to fully exploit the photonic interconnect design space automatically. These techniques will not only benefit photonic architecture/network designs, but also provide circuit level implementation for ICs intended for specific applications. The rest of the paper is organized as follows. In Section 2 we give a brief overview of the photonic devices modeled in the scope of this paper and highlight a few critical device level requirements in building reliable photonic on-chip systems. In Section 3 we explore the design space of low power photonic on-chip integration using non-rectilinear interconnect geometries. In Section 4 we propose an interconnect implementation engine using WDM signal carriers under various physical constraints such as signal integrity, throughput capacity and thermal reliability. Section 5 concludes the paper.

2. NANOPHOTONICS DEVICE MODELS AND DESIGN REQUIREMENTS

Quantified design space exploration and CAD optimization require a properly selected and well parameterized set of nano-photonics elements/devices to build high performance on-chip optical links that are power efficient and thermal reliable. Therefore, we extend¹³ with WDM related modules and thermal models to configure/analyze different on-chip optical links, with respect to critical considerations in the physical design stages, such as power, loss, timing, temperature variation and thermal reliability.




Figure 1. Data link comparison between optical and Cu wires

2.1 Device Characterization

A hybrid link consists of the *optical link* and a segment of Cu interconnect. An *optical link* is configured with a combination of certain modulator, OWG (on-chip optical waveguide), photo-detector and corresponding driver/amplifier circuits, as shown in comparison to a Cu wire in Fig. 1. In particular, we use optical modulators to convert electrical signals into optical (E-to-O) domain (onto OWG channels), and photo-detectors to convert the light pulses into electrical (O-to-E) domain under detection constraints. Couplers are also employed to enable optical waveguide cross-couplings for planar routing. These required building blocks of on-chip optical links are characterized with respect to device operating speed, optical/electrical power consumption, on-chip loss and footprint, etc. Various types of existing nanophotonic devices allow us to configure different high performance optical links in terms of power and/or speed. Based on current photonics fabrication technology, optical signalling has great speed advantage over low-K Cu interconnect (11ps versus 37ps per mm on Metal5/6). The devices employed in this paper are summarized in Table 1.

From Table 1, we can draw several observations. First, multiple devices as modeled enable us to configure different on-chip optical links featuring low power, high performance or a trade-off in-between. Second, signal propagation speed is estimated to be 3.4X (11ps v.s 37ps) faster on currently fabricated OWG than a global Cu interconnect under optimal repeater insertion in 22nm technology. However, further studies must take into considerations the modulation and detection delay for E-to-O and O-to-E data conversions. Also, each sink's photo-detection optical power threshold must be satisfied for successful O-to-E conversion, which is assumed to be 100μW for the devices in Table 1. Last but not least, photo-detection speed should be lower bounded by modulation speed on an optical link to avoid data corruptions during O-to-E conversion. The goal of this paper is to construct an optimized circuit implementation to satisfy all the physical layer requirements.

Table 1. Device and interconnect model details

	footprint	speed	on-chip loss	E-power
mod	30X40um	14Gb/s ²²	2dB	0.7mW
	footprint	speed	O-det power	E-power
detector	20X20um	40Gb/s	0.1mW	1.3mW
	delay	optical loss	thickness	width
WDM	11ps/mm ⁷	1.5dB/cm	230nm	450nm
	delay	thickness	width	repeater
Cu	37ps/mm	1um	0.4um	1.4mm

Cu interconnect on 22nm technology Metal5/6 with $\rho=2.2\mu\Omega\cdot\text{cm}$, $R_{sheet}=0.022\Omega$, $C=2\text{pF/cm}$. MOSFET models for optimal gate sizing/repeater insertion are from Metal Gate/High-K/strained-Si PTM.²³

2.2 Thermal Reliability Modeling for WDM

Current on-chip WDM techniques mainly fall into the following categories: AWG (array waveguide) based, ring resonator based and thin film filter based, among which ring resonator cavity based add-drop filter techniques are most widely employed in architecture designs^{14,17,21} due to its compact footprint (potential ultra density) and demonstrated high quality factor (Q). Unfortunately, all these nanophotonic devices are prone to thermal variations, especially ring resonator structures.

In particular, on-chip temperature fluctuation causes the central operating frequency (wavelength) of a photonic device to drift. If such a drift results in an off-set that falls outside the range of operating bandwidth (BW), the device will degrade or even malfunction. Especially for high energy efficiency on-chip UDWDM devices with ring resonator structure, the quality factor Q ²⁴ (defined as the energy stored in the cavity versus the energy dissipated per unit cycle) is very high and BW is very narrow, rendering the devices highly sensitive to ambience thermal variations. The relationships between thermal reliability, device operating BW , quality factor Q and energy efficiency are defined in Eqn.(1)-(3).

$$Q = \frac{\lambda_0}{\Delta\lambda_{FWHM}} = \frac{\sqrt{r_1 r_2 a} L \pi n_g}{(1 - r_1 r_2 a) \lambda_0} \quad (1)$$

$$n_g(\lambda) = n_e(\lambda) - \lambda \frac{dn_e(\lambda)}{d\lambda} \quad (2)$$

$$BW = \Delta f = \frac{f_{resonant}}{Q} \quad (3)$$

In the equations above, r_1 , r_2 , a , L are ring geometry related parameters, λ_0 is the central working(resonant) wavelength of the ring modulator or detector. n_e is a temperature dependent term, denoting the refractive index of the ring material (e.g., silicon). From the above discussions we can observe that within a relatively small range, one can trade-off Q value for thermal reliability of a certain ring resonator device, without causing aliasing issues in-between of separate channels on a WDM waveguide. However, such a trade-off comes at a power loss penalty that needs to be minimized for power efficient designs.

Based on Eqn.(1)-(3), we investigate and establish the thermal reliability models for WDM related devices that are mainly based on cavity based components (e.g., ring resonators and ring couplers). The thermal reliability models are obtained through exhaustive temperature dependent refractive index modeling/simulation, working bandwidth characterization, power consumption/dissipation simulation and numerical methods such as Finite-difference Time-domain (FDTD) device simulations on powerful computing platforms.²⁵

2.3 Critical Requirements for On-chip Integration

We investigate the on-chip integration potential of various types of optical links in terms of power, performance(timing) and thermal considerations, based on the characterized devices and the circuit models.

Considering the delay overhead introduced by E-to-O and O-to-E data conversions, we define *critical length* L_{crit} as the dimension of an on-chip link above which nanophotonics yield shorter signal delay than Cu. Therefore we have Eqn. (4):

$$T_{mod} + T_{det} + \tau_o \cdot L \leq \tau_e \cdot L \quad (4)$$

where T_{mod} is the E-to-O modulation delay/bit and T_{det} is the O-to-E photo-detection delay/bit; τ_o is signal delay per *mm* on OWG, τ_e is the delay per *mm* on Cu interconnect, L is the length of the link. Solving Eqn. (4) gives us the range of L , whose lower boundary defines L_{crit} value in *mm*.

Due to lack of buffering in the optical domain, optical link configuration requires the speed of modulator be upper-bounded by the speed of photo-detection to avoid data corruptions during O-to-E conversions, thus Eqn. (5) must hold:

$$T_{det} \leq T_{mod} \quad (5)$$

Also for O-to-E conversion, laser power at sink must be equal to or higher than the photo-detection threshold for a logic “1” to be detected successfully, thus Eqn. (6) must hold:

$$P_{O-sink} \geq P_{O-det} \quad (6)$$

where P_{O-sink} is the laser power at sink and P_{O-det} is the detector’s minimal detection threshold, as listed in Table 1

Considering on-chip temperature variations and the ring resonator thermal reliability modeling, we define *temp_th* as the temperature variation threshold value above which a ring resonator based WDM device malfunction. *temp_th* corresponds to a scenario in which even trading off the quality factor Q would not compensate the temperature variation, owing to aliased transmission frequencies between different channels of a WDM trunk and related ring modulators/detectors. Therefore, the following Eqn.(7)(8)(9) must hold:

$$Max_Temp_Var(trunk_i) \leq temp_th \quad (7)$$

$$Temp_Var(Ring_i^{mod}) \leq temp_th \quad (8)$$

$$Temp_Var(Ring_j^{det}) \leq temp_th \quad (9)$$

In Eqn.(7), all the nodes along the path of a WDM trunk must satisfy the *temp_th* condition; while in Eqn.(8)(9), both modulation and detection ring resonators on link $node_i \rightarrow node_j$ must also meet *temp_th*. Therefore with Eqn.(7)(8)(9), the whole optical link’s thermal reliability constraint is properly setup.

3. NON-RECTILINEAR PHOTONIC INTERCONNECT IMPLEMENTATION

As aforementioned, optical routing has unique characteristics when compared with traditional copper routing. Manhattan (X/Y) routing may not be favored on optical layer in a lot of situations because of the large amount of loss caused by sharp turnings. In this section, we present *O-Router* to explore the non-rectilinear geometry space for optical waveguide placement and circuit implementation.

3.1 O-Router: Waveguide Placement/Routing with Flexible Geometries

O-Router performs gridless optical routing with waveguide couplings/crossings on a single layer. As a result, routing geometry becomes very flexible, with different geometries and penalties according to their respective optical interconnect loss. In order to further explore optical routing geometry, we define the following 3 types of losses (with dB unit) on an optical interconnect path in equations 10- 14:

$$L_{loss} = \alpha \cdot length_{path} \quad (10)$$

$$B_{loss} = \beta \cdot \theta \cdot r^{-\eta} \quad (11)$$

$$C_{loss} = \gamma \cdot Num_{couplers} \quad (12)$$

$$P_{loss} = L_{loss} + B_{loss} \quad (13)$$

$$Total_{loss} = P_{loss} + C_{loss} \quad (14)$$

where L_{loss} is straight line waveguide loss, it is proportional to the length of optical interconnect, with a coefficient α ; B_{loss} is the bending loss, since waveguide cross-section width is negligible compared to the bending radius in *O-Router*, we assume B_{loss} to be proportional to the degree of the optical interconnect (silicon waveguide) arc angle θ , and inversely proportional to the radius r of the interconnect, with an index η ; C_{loss} is the coupling loss, proportional to the number of couplers (crossings) on the interconnect, with a coefficient γ . All related coefficients are determined by our numerical library.

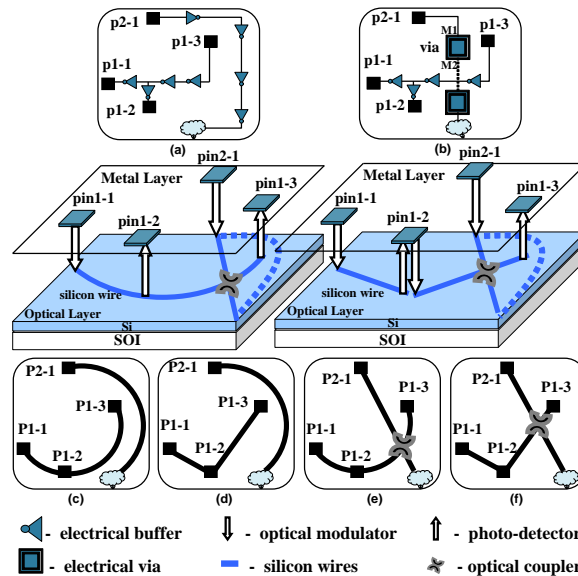


Figure 2. Motivational example for electrical routing v.s optical routing

3.2 A Case Study Example

We first briefly explore the different trade-offs for optical routing. As shown in Figure. 2, there are 2 nets to be routed on a chip, noted as pin_i-j , meaning it is the j th member of net i ; Fig. 2(a) and (b) shows two alternatives for conventional routing on electrical layer with buffers and/or metal via inserted to alleviate the timing penalty caused by the long wires across the chip. Buffers are inserted since RC delay increases quadratically with electrical wire length. Yet buffer insertion is not all-powerful technique. Generally speaking, cross-chip timing critical nets are tough to fix thus impose great difficulty to VLSI design timing closure. As technology further scales down and system integration level rockets, issues with electrical interconnect will get more severe.

Fig. 2(c)-(f) show 4 possible routing geometries for the 2 nets on optical layer, according to our optical routing. Routing geometry (c) requires a total of 2 optical modulators: 1 inserted at P1-1, 1 inserted at P2-1, while for (d), 1 extra modulator will be inserted at P1-2, in order to drive P1-3, since sharp turning at P1-2 is either too lossy or too costly to fix other than using an extra modulator. In (e) and (f), optical coupler is introduced for coupling optical signal across 2 wires, with certain amount of loss. In these 2 cases, couplers can be employed either because doing so results in less amount of loss than taking detours as in (c) and (d), or because taking detours results in more coupling loss with other nets on chip, etc.

We can learn that geometries (c)(e) result in least among of modulating power among (c)-(f), yet optical interconnect bending loss: B_{loss} is also introduced, as well as the coupling loss: C_{loss} (in (e)) so that the constraint for successful detection at P1-3 may be violated due to too much loss on interconnect. To optimally pick the best routing geometry from the (c)-(f) 4 cases is the motivation of *O-Router*.

O-Router targets at finding optimal optical routing geometry to minimize total modulating power, subject to various constraints imposed by the device characterizations.

Given the pin locations of certain circuit netlist for optical routing, *O-Router* seeks optimal routing solution with Integer Linear Programming to minimize total modulating power, meanwhile satisfying various detection constraints according to established library parameters.

3.3 Automatic Implementation based on Integer Linear Programming

To reduce custom design workload, we propose a mathematical technique and develop an automatic engine to implement the on-chip photonic interconnects. First we enumerate all routing geometries for the 2-pin, 3-pin and 4-pin nets, shown in Fig. 3(a) concave shape 4-pin net is shown as an example). Each X_{ij} is an integer variable, where $i \in net_space$, $j \in sol_space(net\ i)$. When $X_{ij} = 1$, the corresponding routing geometry from Fig. 3 will




Figure 3. A partial list of optical interconnect implementation geometries for 2, 3 and 4 pin nets

be adopted, as part of the final routing solution space. Number of modulators in each X_{ij} is also recorded; our numeric library returns the actual modulating power based on this number and the ij index.

The ILP formulation is as follows in Equation 15- 24, with all terms and variables explained in Table 2. The objective function is the total power required to drive all the on-chip optical modulators for our optical interconnect framework. The ILP solver will minimize the objective function, subject to constraints imposed from Eq. 16 to Eq. 24. In Eq. 15, the first term $MPow_{X_{ij}}$ is total modulating power consumption for routing geometry X_{ij} using 1X modulators, while the second term $(MPow_{penalty} - P_0) \cdot M_{ij} \cdot N_{ij}$ is for penalizing the usage of 10X driving power ModulatorX: if M_{ij} is 1(hard constraint violation), then ModulatorX will be used to replace all Modulator1/s in geometry X_{ij} to meet the constraint (P_0 is the laser power consumption of Modulator1).

$$\min\left\{ \sum_{i \in \text{net_space}()}^{j \in \text{sol_space}(i)} [MPow_{X_{ij}} \cdot X_{ij} + (MPow_{penalty} - P_0) \cdot M_{ij} \cdot N_{ij}] \right\} \quad (15)$$

s.t

$$\forall i, m \in \text{net_space}, i \neq m, j \in \text{sol_space}(i), n \in \text{sol_space}(m) :$$

$$P_{loss_{X_{ij}}} \cdot X_{ij} + \text{net}_{loss_{X_{ij}}} \leq \text{loss_th}_{X_{ij}} + \text{pow} \cdot N_{ij} \cdot M_{ij} \quad (16)$$

$$P_{loss_{X_{ij}}} = L_{loss_{X_{ij}}} + B_{loss_{X_{ij}}} \quad (17)$$

$$\text{net}_{loss_{X_{ij}}} = \sum_{m \in \text{net_space}}^{n \in \text{sol_space}(m)} C_{loss_{X_{ij_mn}}} \cdot X_{ij_mn} \quad (18)$$

$$C_{loss_{X_{ij_mn}}} = \gamma_{ij_mn} \cdot \text{cross_num} < X_{ij}, X_{mn} > \quad (19)$$

$$X_{ij} + X_{mn} \leq 1 + X_{ij_mn} \quad (20)$$

$$(1 - X_{ij}) + (1 - X_{mn}) \leq 2 - 2X_{ij_mn} \quad (21)$$

$$\sum_{j \in \text{sol_space}(i)} X_{ij} = 1, \quad \text{where } X_{ij} = 0 \text{ or } 1 \quad (22)$$

Table 2. Descriptions for ILP involved terms and variables.

Name	Description	Name	Description
$\text{net_space}()$	set of nets for an optical netlist	$\text{sol_space}(i)$	set of possible routing geometries for net i
$MPow_{X_{ij}}$	total modulator power consumption of routing geometry X_{ij}	$MPow_{penalty}$	power consumption penalty for using each ModulatorX. Set to 10 times of P_0
P_0	power consumption of Modulator1	$P_{loss_{X_{ij}}}$	propagation loss power on silicon wires of X_{ij}
N_{ij}	least number of optical modulators used for geometry X_{ij}	$C_{loss_{X_{ij}}}$	coupling loss power between routing geometry X_{ij} and X_{mn}
X_{ij}	integer variable. $X_{ij} = 1$ means to accept the j th routing geometry of net i	M_{ij}	integer variable. $M_{ij} = 1$ means to insert modulatorX into j th routing geometry of net i
X_{ij_mn}	integer variable. numerically equals to $X_{ij} \cdot X_{mn}$	γ_{ij_mn}	coupling loss coefficient dependent on geometry X_{ij} and X_{mn}
$\text{loss_th}_{X_{ij}}$	loss threshold for O-E conversion for X_{ij}	pow	extra driving power penalty of ModulatorX

Table 3. Performance comparisons between *O-Router* and Minimum Spanning Tree algorithm.

	photo-detection threshold : 55%				photo-detection threshold: 75%			
	ibm01	ibm02	ibm03	ibm04	ibm01	ibm02	ibm03	ibm04
Net number	5	20	50	137	5	20	50	137
Pin number	15	50	155	391	15	50	155	391
Pin/net ratio	3	2.5	3.1	2.85	3	2.5	3.1	2.85
MST-routing (normalized power)	3.5	6	35.66	305.13	3.5	12.75	39	306.25
<i>O-Router</i> (normalized power)	1	2.88	10.75	57.75	2.13	5.38	16.5	100.25
Improvement	71.40%	52.00%	69.90%	81.10%	39.10%	57.80%	57.70%	67.30%

$$X_{ij_mn} = 0 \text{ or } 1 \quad (23)$$

$$M_{ij} = 0 \text{ or } 1 \quad (24)$$

Constraint Eq. 16 is set for each routing geometry X_{ij} , such that its total loss (propagation loss P_{loss} and coupling loss C_{loss}) is bounded by an upper bound of loss threshold: $loss_th_{X_{ij}}$, once the upper bound of loss is exceeded, it means the photo-detection requirements in routing geometry X_{ij} are violated. If among all feasible X_{ij} , some of such constraint is inevitably violated, then ModulatorX will be inserted into the corresponding geometry X_{ij} and replace existing 1X modulators. Constraint Eq. 19 maps the crossing number to coupling loss.

For the calculation of optical interconnect coupling number, we introduced the cross-term integer variables: X_{ij_mn} . Numerically, it is the product of term X_{ij} and X_{mn} . Since variable multiplications are not supported by ILP solver, we add the constraint pair Eq. 20- Eq. 21. Integer constraints Eq. 20 and Eq. 21 bound the X_{ij_mn} term so that it always equals the product of its two corresponding routing geometries. Equality constraint Eq. 22 makes sure that the ILP solver eventually picks only 1 geometry out of each net for the final optimal solution.

3.4 Simulation Results and Analysis

Simulations are carried out according to the aforementioned steps and original electrical benchmarks come from ISPD98/08 routing benchmarks. ibm01-04 are the final 4 optical netlists benchmarks, listed as in Table 3. Due to considerations of silicon wire spacing/low coupling noise communication, the sizes of the optical netlists are kept from small to medium, and the optical layer pin density is kept from low to medium. As a baseline for *O-Router*, Minimum Spanning Tree (MST) routing algorithm is implemented on ibm01-04. Both *O-Router* framework and MST algorithm are repeated on ibm01-04 for 2 different photo-detection threshold values: 55% and 75%. Such percentages signify the photo-detection power threshold for received signals at the end of optical interconnect. Therefore, 75% threshold photo-detectors impose stricter detection requirements on *O-Router* framework.

In Table 3, the simulated power consumptions are normalized by the amount of power reported by *O-Router* on ibm01, under photo-detection threshold of 55%. For 55% threshold, *O-Router* achieves above 50% of power reduction compared to MST baseline, with a max of 81.1% on ibm04. For the 75% threshold, *O-Router* reports slightly less power reductions due to higher detection requirements; still an average of above 50% reduction, with a max of 67.3% of power reduction on ibm04.

4. GLOW: LOW-POWER THERMAL-AWARE INTERCONNECT IMPLEMENTATION WITH WDM

In this section, we employ nanophotonic on-chip WDM interconnect (Fig. 4) to achieve high density/capacity implementation in IC backend design global routing stage. We propose *GLOW*, a new hybrid global router for power-efficient thermal-reliable physical synthesis featuring WDM waveguide placement, optical channel allocations and optical-electrical data converter planning.

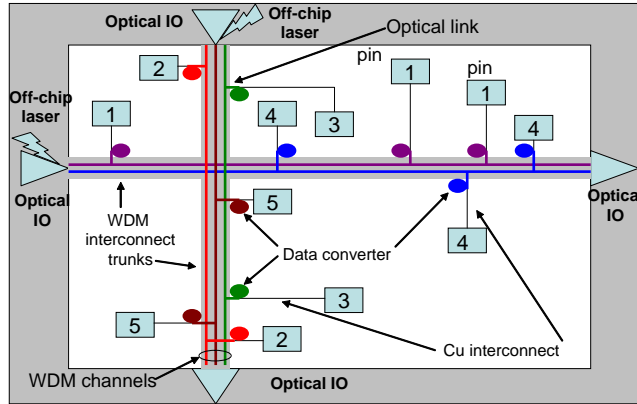


Figure 4. High routing capacity interconnect using signal WDM

4.1 A Case Study Example

With on-chip WDM providing great signal multiplexing capacity, we motivate a global router to take the advantages of WDM channels under various physical design constraints such as thermal reliability and timing. A simple scenario is illustrated in Fig. 5. Given a net (A,B,C,D) to be routed with node A as the driver, B,C,D as sinks, we aim to find a global routing solution in optical-electrical domain to satisfy various integration requirements, such as: thermal reliability and functionality, minimal driving power, signal integrity and data conversion quality, timing considerations and WDM channel utilization rate, etc.

DEFINITION I. **WDM link**: A piece of on-chip interconnect that solely or partially employs WDM. It consists of laser source, nanophotonic waveguide (OWG) and modulation/detection devices.

DEFINITION II. **WDM trunk**: The body of the OWG in a WDM link is also referred to as a WDM trunk.

DEFINITION III. **WDM channel**: The working wavelength of a WDM link. Each channel is assigned a unique wavelength λ which signifies the carrier frequency of the optically modulated signals.




Figure 5. Example of thermal-aware hybrid routing with WDM

In Fig. 5, thermal issue refers to the scenario for which on-chip temperature variation causes extra power loss, signal degradation or even malfunction to the nanophotonics devices, such as modulator, photo-detector and WDM waveguide. Without careful considerations and planning, an opto-electrical link could fail to operate in reality due to fallacies introduced by high thermal variations. We can simply set the thermal killer regions as blockages in the routing stage, yet still moderate temperature variation affects the rest of the chip that suffer from different degrees of thermal induced power loss. Under power objective, this part of power loss must be minimized, although it would not cause circuit malfunction. Other sources of power loss comes from waveguide crossings, propagation loss, etc.

During routing in the opto-electrical domain, timing condition must be met such that a hybrid data link does not generate longer signal delay than an otherwise routing path in the electrical domain. Under such conditions in Fig. 5, link A→B is routed with Cu interconnect while links A→C, A→D are partially merged with WDM

trunks, meanwhile link A→D takes trunk1 due to the thermal blockage between sink D and trunk2. Data links from different nets must be assigned different wavelengths (i.e., channels) when sharing the same trunk.

For high WDM channel utilization rate, sharing onto a single WDM trunk is encouraged unless timing and/or thermal conditions are violated. In this case, path A→C would tend to merge with link A→D onto trunk1, however is prohibited by the long delay from trunk1 to sink C.

4.2 Overall CAD Flow

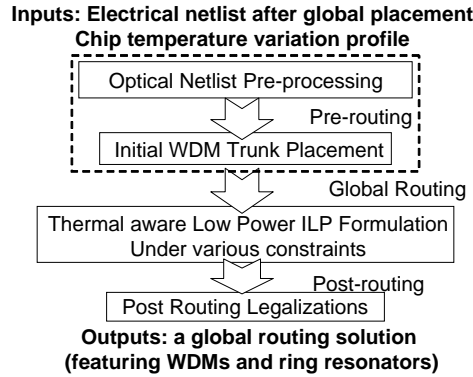


Figure 6. An overview of our proposed CAD flow

In Fig. 6, we illustrate a top level flow diagram of our proposed method, starting from a given input netlist and on-chip temperature variation profile. Such a CAD flow is consisted of 3 major stages: a **Pre-routing** stage that prepares the optical netlist and WDM trunk placement; a **Global Routing** stage that serves as the core formulation of the WDM channel assignment problem based on various physical design constraints; and a **Post-routing** stage that further examines the legalization issues in both the optical and electrical domains. In the following subsections, we describe each function block in Fig. 6 in detail.

4.2.1 Netlist Pre-processing

Netlist pre-processing step prepares the optical netlist with an initial consideration of the *timing condition* which guarantees that the circuit timing does not degrade after employing nanophotonics (since each data conversion takes significant time). This step is mainly proposed to derive optical netlist test cases from existing electrical benchmarks such as ISPD global routing netlists. This step is very critical since it selects proper pins (nets or partial nets) from the electrically placed netlist to synthesize in the **Global Routing** stage. The selection is designed such that the minimal manhattan distance of all driver-sink pairs mapped onto the optical domain is lower bounded by the *critical length* L_{crit} . This step serves to yield *non-negative timing gain* in the optical domain than in the electrical domain. This aligns well with *critical length* definition and discussions in Section 2. The main technique involved is described as follows,

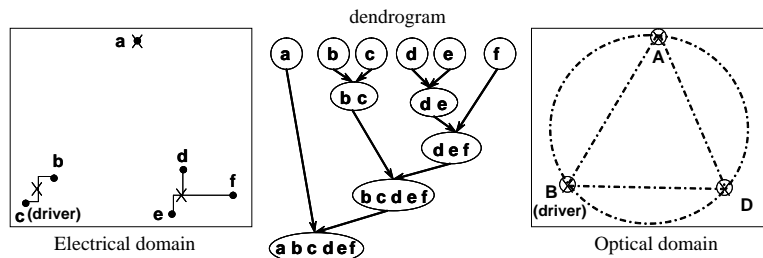


Figure 7. A brief illustration of netlist pre-processing

Pin Clustering: To cluster the electrically placed input netlist based on manhattan distance using hierarchical clustering method. In this case, we first construct the *dendrogram* (illustrated in Fig. 7) and then pick out the




Figure 8. Our WDM based global routing scenario

clusters satisfying the L_{crit} dimension with a *depth first search* on the *dendrogram*. The result of this procedure is a set of clusters whose respective geometric medians are mapped to the optical domain as pseudo-pins. These pseudo-pins form the *Optical Netlist*, while the rest of pins within each cluster remain on the electrical domain and are electrically interconnected to their geometric median. Therefore, only 1 O-to-E or E-to-O conversion is needed per cluster. This procedure is briefly illustrated in Fig. 7, where $a-f$ are pins of certain net in the electrical netlist and ABD are pseudo pins (a partial net) mapped onto the optical plane to represent clusters with edges larger than L_{crit} in the *dendrogram*. B is the driver pin in optical domain since driver pin c lies in the bc cluster in electrical domain.

4.2.2 Initial WDM Trunk Placement

Initial WDM trunk placement depend on the median of geometry distributions of optical nets in the *Optical Netlist* and is carried out in a partitioned manner across the whole chip area according to Eq. (25) as a general guideline, until the total number of WDM channels is sufficient to hold the total number of optical nets/links.

$$Place_{trunk^k} = med\{med[net_i]\}^{i \in Partition^k} \quad (25)$$

The partition based initial placement executes in the following steps:

- Continues for both horizontal and vertical directions
- Avoids over-heated regions marked as thermal blockages
- Partition ends when the number of WDM channels are sufficient for the total links in the optical netlist.
- Extra WDM trunks may need be added in **Post-routing**

4.2.3 Thermal-aware Low Power Routing

First, we define *timing condition* as the condition that guarantees smaller signalling delay on the opto-electrical link than on Cu interconnect. This is a critical consideration since each additional O-E/E-O data conversion brings significant delay. The *thermal condition* is defined to make sure the local temperature variation does not fall out of the working range of the ring modulators. In case of a violated *thermal condition*: (1) Q value will be adjusted to trade-off power efficiency for thermal reliability; (2) if (1) can not be done without causing aliases between separate WDM channels, that particular region is set as a thermal blockage. For the core routing problem of the **Global Routing** stage, we propose 2 approaches, namely *CAT* and *GLOW*.

In Fig. 8, we illustrate the routing problem after **Pre-routing**, with laser sources from off-chip whose driving power to each WDM waveguide trunk differs according to the total number of channels assigned/ utilized after the routing stage. To constrain the solution space for the global routing stage, we assume to take the shortest distance route when a pin is to connect to certain WDM trunk, i.e, data conversion (mod/det) only happens on WDM trunks. Based on Fig. 8, we will briefly discuss both approaches.

CAT: A greedy heuristic approach for WDM Channel Assignment under Thermal considerations. The basic motivation for *CAT* is to assign optical nets/links to WDM trunks in a sequential manner, meanwhile to combine timing and thermal-awareness constraints locally for each WDM trunk. In particular, *CAT* picks all the local nets/link satisfying the timing condition and assign the least power consuming links to fill the available channels to certain WDM waveguide, then move onto the next waveguide. If at the end of the process, there are still unassigned nets, then the *Initial WDM Placement* stage will be appended with extra WDM resources to route the remaining nets.

CAT's advantage mainly include run time and the simplicity of the implementation, however, key power related factors are neglected such as WDM trunk crossings and the co-relation between thermal reliability and Q value related resonant power loss. Also, there is guarantee for a solution with a global minima.

GLOW: An ILP based global routing approach for low power driven thermal-reliable WDM synthesis. *GLOW* is a low power driven global router with various physical design constraints. With careful selection of IV (integer variables) and BV (binary variables), we not only formulate the key power related terms, but also the cross-related variables and constraints that are otherwise very hard to capture. We will discuss its formulations in Section 4.4.

4.3 CAT Routing Algorithm

CAT is designed and implemented as a greedy heuristic approach for thermal-aware WDM channel assignment under timing constraints. It is performed in 3 major steps: first, *Initial WDM Trunk Placement*; second, *Timing and Thermal Condition Calculation*; third, *Greedy Channel Assignment*. In this paper, *CAT* is used as a baseline example for performance evaluations and analysis.

Initial WDM Trunk Placement: *CAT* uses the same trunk initial placement result as in *GLOW*.

Timing and Thermal Condition Calculation: In this step, all the WDM trunks are traversed in low-power priority order sequentially. For each trunk, timing/thermal conditions for all optical links are calculated and updated using models from our numeric library.

Greedy Channel Assignment: For the channel assignment, we use a greedy heuristic method which executes in 3 phases: **Phase1:** Form set $S(link_i)$ for WDM *trunk_i* with the optical links that guarantee smaller signalling delay than in the electrical domain. $S(link_i)$ is a set of link candidates to be assigned to WDM *trunk_i*. **Phase2:** Sort the links in $S(link_i)$ with *Thermal Condition* metric in ascending order. **Phase3:** Assign links from $S(link_i)$ to *trunk_i* in ascending order, until the total number of optical nets assigned reaches $Cmax$.

4.4 GLOW Routing Algorithm

4.4.1 ILP Formulation

To formulate the optical global routing problem, we introduce parameters and binary/integer variables as shown in Table 4. In particular, we emphasize the following terms:

- n, m : total number of WDM trunks in the row and column directions after initial placement, respectively.
- W_i : binary variables denoting the assignment status of WDM trunk i . If W_i is 0, trunk i is not assigned any optical nets in the final routing solution, therefore will not be turned on (no input laser power from its optical IO port); if W_i is 1, trunk i is assigned certain nets, but may still has available channels.
- W_{ij} : binary variables numerically equal to the product of W_i and W_j , where $i \in [0, n-1]$, $j \in [n, n+m-1]$. If W_{ij} is 0, trunk i and trunk j are not physically crossed, vice versa.
- $S_{link_k}^{trunk_i}$: binary variables, with 0 meaning link k is assigned onto WDM trunk i .
- $Sum_{net_i}^{trunk_j}$: integer variables, representing the total number of optical nets assigned onto trunk j in the final solution.
- $\lambda_{net_i}^{trunk_j}$: binary variables, with 0 meaning net i is assigned onto WDM trunk j in the final routing solution; vice versa.

Table 4. Variables/parameters in ILP formulation

Name	Description	Name	Description
P_{total}	total laser power consumed	P_{loss}	total on-chip laser power loss
$P_{dynamic}$	total on-chip laser power for optical signaling	P_0	base power consumption for a WDM trunk
P_{cross}	total power loss due to trunk crossings	P_{trunk_thm}	total power loss due to trunk thermal effects
P_{ring_thm}	total power loss due to ring thermal effects	P_{path}	total power loss due to photon propagation
P_{λ_i}	laser power on channel λ_i for optical signalling	P_{thp}^{ij}	laser power loss when trunk i , trunk j cross
$P_{trunk_thm}^i$	thermal related power loss on trunk i	$P_{ring}^{link_i}$	laser power loss on the rings of link i
W_i	BV: allocation status of trunk i	W_{ij}	BV: crossing status of trunk i and trunk j
$S_{link_j}^{trunk_i}$	BV: assignment status of link i onto trunk j	$Sum_{net_i}^{trunk_j}$	IV: # of links in net i assigned to trunk j
$\lambda_{net_i}^{trunk_j}$	BV: assignment status of net i onto trunk j	$T_{var}^{link_i}$	temperature variation on the rings of link i
C_{max}	channel capacity of each WDM trunk	PIN_{max}	max pin # in certain net of the optical netlist
$temp_th$	temperature variation tolerance threshold	τ_e	delay per unit length on Cu interconnect
τ_o	delay per unit length on optical links	τ_{conv}	delay overhead by data conversions
WL_o^i	Cu wire length on link i	WL_o^i	optical wire length on link i
$HPWL^{link_i}$	half parameter wire length of link i	L_{crit}	critical length of on-chip nanophotonic interconnect

• C_{max} : channel capacity of each WDM trunk. It is total available channel number that serves at an upper bound limit condition when assigning optical nets.

• PIN_{max} : max pin number of certain net in the optical netlist. For our proposed formulation, PIN_{max} can take any number.

Please see Table 4 for the complete list of variables and parameters.

We propose the following objective function for *GLOW*'s thermal-aware low power routing featuring on-chip photonics WDM:

$$\text{Minimize}\{P_{total}\} \text{ w.r.t } W_i, W_{ij}, S_{link_i}^{trunk_j}, \lambda_{net_i}^{trunk_j} \quad (26)$$

such that:

$$P_{total} = P_{loss} + P_{dynamic} \quad (27)$$

$$P_{loss} = P_{cross} + P_{trunk_thm} + P_{ring_thm} + P_{path} \quad (28)$$

$$P_{cross} = \sum_{\substack{j \in [n, n+m-1] \\ i \in [0, n-1] \\ i \in \text{all trunks}}} W_{ij} * P_{thm}^{ij} \quad (29)$$

$$P_{trunk_thm} = \sum_i W_i * P_{trunk_thm}^i \quad (30)$$

$$P_{ring_thm} = \sum_i \sum_{j \in \text{all links}} S_{link_j}^{trunk_i} * P_{ring}^{link_j} \quad (31)$$

$$P_{dynamic} = \sum_i \sum_{j \in \text{all nets}} \lambda_{net_j}^{trunk_i} P_{\lambda_i} + \sum_i W_i P_0 \quad (32)$$

Eq. (26) above gives the objective function of *GLOW* as the total power P_{total} required to drive the circuit. As shown in Eq. (27), P_{total} is divided into 2 parts: the total optical power loss on chip P_{loss} , which is the amount of power the drivers need to compensate for the guarantee of *detection conditions* on photo-detectors; and $P_{dynamic}$, the signal switching power on WDM channel carriers.

P_{loss} is divided into 4 terms: waveguide crossing power, thermal related WDM trunk power, thermal related ring resonator power and the power to compensate propagation loss of on-chip waveguide.

$P_{dynamic}$ consists of 2 terms: P_0 is the base power consumption for each WDM trunk, it is a constant power cost when turning on a N-channel WMD trunk; the 2nd term is the switching power on all WDM channels, which is linearly proportional to the number of channels utilized. Apparently, WDM trunk multiplexing/sharing rate is to be maximized in order to avoid unnecessary P_0 's.

All power related terms are modeled according to our previous discussions in Section 2 and Section 4.2, please also see Table 4 for further explanations of each term.

4.4.2 Physical Design Constraints

Following the discussions in Section 2, we present the detailed mathematical formulations of various routing constraints for *GLOW*:

- Timing constraint: for each optical link, the routing solution must not result in longer signal delay than HPWL estimated delay in the electrical domain:

$$S_{link_i}^{trunk_j} [\tau_e * WL_e^i + \tau_o * WL_o^i + \tau_{conv}] \leq \tau_e * HPWL^{link_i} \quad (33)$$

- Selection constraint: to make sure each link i is only assigned to one WDM trunk. For each link i , we have:

$$\sum_j^{j \in \text{all trunks}} S_{link_i}^{trunk_j} = 1 \quad (34)$$

- Channel capacity constraint: a WDM trunk does not exceed its capacity limit. For each WDM trunk j :

$$\sum_i^{i \in \text{all nets}} \lambda_{net_i}^{trunk_j} \leq Cmax \quad (35)$$

- Detection constraint: the final optical power at each sink on each link must be large enough to be detected.
- Thermal constraint: for each link (pair of pins from source to sink), local temperature variation must not result in performance degradation or malfunction. For each link i and trunk j :

$$S_{link_i}^{trunk_j} * T_{var}^{link_i} \leq temp_{th} \quad (36)$$

- Binary/Integer variable constraints: since W_{ij} and $\lambda_{net_i}^{trunk_j}$ are introduced to eliminate non-linear terms, the following constraints must be enforced:

$$2W_{ij} \leq W_i + W_j \leq 1 + W_{ij} \quad (37)$$

where $i \in [0, n - 1]$, $j \in [n, n + m - 1]$

$$\frac{(2 \sum_k^{k \in net_i} S_{link_k}^{trunk_j} - 1)}{2PINmax} \leq \lambda_{net_i}^{trunk_j} \leq 2 \sum_k^{k \in net_i} S_{link_k}^{trunk_j} \quad (38)$$

$$\frac{(2 \sum_{i=1}^{all\ nets} \lambda_{net_i}^{trunk_j} - 1)}{2Cmax} \leq W_j \leq 2 \sum_{i=1}^{all\ nets} \lambda_{net_i}^{trunk_j} \quad (39)$$

Here Equation(38) and (39) are enforced for two-fold reasons: (1) we are able to calculate the number of optical nets assigned to certain WDM trunk via optical link related variables; (2) to introduce non-linear relation between $\lambda_{net_i}^{trunk_j}$ and $S_{link_k}^{trunk_j}$ under ILP formulation. For this part an intermediate term $Sum_{net_i}^{trunk_j}$ is introduced by Equation(40) as follows,

$$Sum_{net_i}^{trunk_j} = \sum_k^{k \in net_i} S_{link_k}^{trunk_j} \quad (40)$$

Equation(38)(39)(40) together make sure that if $Sum_{net_i}^{trunk_j} = 0$, then $\lambda_{net_i}^{trunk_j} = 0$; if $Sum_{net_i}^{trunk_j} > 0$, then $\lambda_{net_i}^{trunk_j} = 1$.

Table 5. Simulation result comparisons between our proposed *CAT* and *GLOW*

Method	CAT						GLOW					
	CK1	CK2	CK3	CK4	CK5	CK6	CK1	CK2	CK3	CK4	CK5	CK6
Optical Netlist												
Net #	35	70	137	240	437	996	35	70	137	240	437	996
Pin #	95	187	391	658	1357	2698	95	187	391	658	1357	2698
Sink #	60	117	254	418	920	1702	60	117	254	418	920	1702
Trunk # ^a	4	11	12	25	46	138	5	16	22	40	87	193
Channel # ^b	36	72	138	286	570	1314	35	79	152	295	602	1408
Avg. Channel/trunk	9.0	6.55	11.5	11.44	12.39	9.52	7.0	4.94	6.9	7.38	6.92	7.29
Total trunk-length	4.8	13.2	14.4	30	55.2	165.6	6.0	19.2	26.4	48.0	104.4	231.6
Total power ^c	1.45	4.68	6.81	13.8	27.26	65.52	1.00	2.48	5.27	7.25	16.63	32.86
Power reduction%	-	-	-	-	-	-	31.0%	47.0%	22.6%	47.5%	39.0%	49.8%

^a Each WDM trunk has a maximum of 32 available channels in the initial placement stage. Unassigned trunks will be turned off in the global routing stage.

^b Unassigned WDM channels will be turned off (no laser input from off-chip) in the global routing stage.

^c Total power consumption is normalized to the power consumed on CK1 by *GLOW*.

4.5 Simulation and Testing

CAT and *GLOW* are assessed by various testing cases derived from ISPD netlists. We describe the benchmark preparation and discuss the simulation results as follows,

Benchmarks and Simulation Setups: In Table 5 we list 6 optical benchmarks: CK1-6, with net number ranging from 35 upto 996. These test cases are derived from ISPD global routing contest benchmarks (with over 100K nets) by: (1) up-scaling the chip dimension into centimeter scale; (2) employing our proposed *Optical Netlist Pre-processing* techniques to generate optical netlists. Considering the limited integration volume of current on-chip WDM nanophotonics, the sizes of these testing netlists are suitable.

For the hierarchical clustering procedure, L_{crit} is set to 3.7mm for centimeter-scale chips. We assume all the inserted ring resonators are legalized and initially thermally tuned. The on-chip thermal variation profiles are randomly generated based on measured data of real processor chips. The tolerance threshold $temp_{th}$ of the maximal range of temperature variation is set to between 15 to 20 degrees, as hard constraints in our problem formulation. Corresponding wavelength off-set sensitivity of the WDM interconnect is set to 0.12nm/degree C. For the WDM trunk initial placement, we use 32-channel WDM trunks to start with, then run the proposed global routing algorithms on 3.0GHz Linux workstations with 8GB memories.

Result and Analysis: In Table 5, we show simulation results of *CAT* and *GLOW*, with total power consumption normalized to the power value that *GLOW* gives on CK1. Compared with *CAT*, *GLOW* demonstrates significant 22%-49% of total power reductions on CK1-6, respectively.

Reasons of such improvement are mainly two-fold: first, *CAT* only searches for local optimal solutions and assign optical nets/links to WDM trunks in a sequential/local manner, while *GLOW* aims at a global optimal solution with mathematical programming techniques; second, *CAT* is not aware of the waveguide crossing power, nor does it consider the thermal related ring resonator power-reliability trade-off in a global manner; while on the other hand, the ILP formulation of *GLOW* makes it possible to model all the key power contributors.

Also in Table 5 we show the WDM channel/trunk allocation status of *CAT* and *GLOW* on CK1-6. We see that compared with *GLOW*, *CAT* assigns fewer number of WDM trunks, resulting in a slightly higher number of average WDM channels per trunk and shorter total length of on-chip WDM waveguide. *GLOW*, however, works by assigning WDM trunks/channels across the chip aiming at the global solution of power consumption minimization under given thermal reliability requirements. This helps *GLOW* to bring down the total power at the cost of some extra OWG wirelength. This is acceptable since the fabrication cost of straight OWGs are relatively low meanwhile the silicon layer provides rich resources for monolithic integration of the required nanophotonics components.

In some few cases when there are no feasible solutions exist, the ILP formulation will not return valid WDM channel/trunk allocation strategy and the WDM trunk initial placement must be adjusted (by adding more trunks). In this paper, such adjustments are carried out in a progressive and heuristic manner until feasible integer solutions are found. With accelerated ILP, *GLOW* manages to locate the optimal solutions for all 6 optical netlists within 0.8 CPU hours. Such run-time is well acceptable as the optical routing problem size is fairly limited, i.e., only the top global nets/pins are mapped into the optical domain while the rest nets are placed and routed in the electrical domain.

5. CONCLUSION

This paper explored the design space of low power on-chip nanophotonics interconnect implementations using flexible interconnect geometries and automated techniques. In particular, we have studied and showed that the proposed set of techniques allow us fully exploit the non-rectilinear waveguide placement and the signal multiplexing mechanisms meanwhile satisfying various complex sets of physical device constraints such as thermal conditions and signal integrity, etc. They can be employed to efficiently build low-power and economical on-chip optical interconnect for high performance application specific IC. We believe a lot of future research can be done to co-optimize the CAD and nanophotonics technologies using automated design techniques.

ACKNOWLEDGMENTS

This work is supported in part by Texas Advanced Research Program.

REFERENCES

- [1] M.-C. Frank Chang et al. RF Interconnects for Communications On-Chip. In *Proc. Int. Symp. on Physical Design*, 2008.
- [2] Navin Srivastava et al. Performance Analysis of Carbon Nanotube Interconnects for VLSI Applications. In *ICCAD*, 2005.
- [3] David A. B. Miller. Device Requirement for Optical Interconnects to Silicon Chips. In *IEEE Special Issue on Silicon Photonics*, 2009.
- [4] Yongqiang Jiang et al. 80-micron Interaction Length Silicon Photonic Crystal Waveguide Modulator. In *Applied Physics Letters*, 2005.
- [5] Yurii Vlasov. Silicon Photonics for Next Generation Computing Systems. In *European Conference on Optical Communications*, 2008.
- [6] Ian. O'Connor. Optical Solutions for System-Level Interconnect. In *Proc. System Level Interconnect Prediction*, 2004.
- [7] Kyung-Hoae Koo et al. Compact Performance Models and Comparisons for Gigascale On-Chip Global Interconnect Technologies. In *IEEE Trans. on Electron Devices*, 2009.
- [8] Ian A. Young et al. Optical I/O Technology for Tera-Scale Computing. In *IEEE J. Solid-State Circuits*, 2010.
- [9] M. R. Watts. Silicon Photonics in High Performance Computing. In *Proc. Photonics in Switching (PS) Conf.*, 2010.
- [10] Jacob R. Minz et al. Optical Routing for 3D System-on-Package. In *Proc. Design, Automation and Test in Europe*, 2006.
- [11] Assaf Shacham et al. Photonic Networks-on-Chip for Future Generation Chip Multiprocessors. In *IEEE Trans. on Computers*, 2008.
- [12] Duo Ding et al. O-Router: An Optical Routing Framework for Low Power On-Chip Silicon Nano-photonics Integration. In *DAC*, 2009.
- [13] Duo Ding et al. OIL: A Nanophotonic Optical Interconnect Library for a New Photonic Networks-on-Chip Architecture. In *SLIP*, 2009.
- [14] Yan Pan et al. Firefly: Illuminating Future Network-on-Chip with Nanophotonics. In *Proc. Int. Symp. on Computer Architecture*, 2009.
- [15] Johnnie Chan et al. Physical-Layer Modeling and System-Level Design of Chip-Scale Photonic Interconnection Networks. In *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 2011.
- [16] Pranay Koka et al. Silicon-Photonic Network Architectures for Scalable, Power-Efficient Multi-Chip Systems. In *Proc. Int. Symp. on Computer Architecture*, 2010.
- [17] Ajay Joshi et al. Silicon-Photonic Clos Networks for Global On-Chip Communication. In *Int. Symp. on Networks-on-Chip*, 2009.
- [18] Jong-Moo Lee et al. Controlling Temperature Dependence of Silicon Waveguide Using Slot Structure. In *Optics Express*, 2008.
- [19] Biswajeet Guha et al. CMOS-Compatible Athermal Silicon Microring Resonators. In *Optics Express*, 2010.
- [20] Moustafa Mohamed et al. Power-Efficient Variation-Aware Photonic On-Chip Network Management. In *Int. Symp. on Low Power Electronics Design*, 2010.
- [21] Zheng Li et al. IRIS: A Hybrid Nanophotonic Network Design for High-Performance and Low-Power On-Chip Communication. In *J. on Emerging Technologies in Computing Systems*, 2011.
- [22] Po Dong et al. Low Vpp, Ultralow-energy, Compact, High-speed Silicon Electro-Optic Modulator. In *OPTICS EXPRESS*, 2009.
- [23] Predictive Technology Model, <http://ptm.asu.edu>.
- [24] Payam Rabiei et al. Polymer Micro-Ring Filters and Modulators. In *J. of Lightwave Technology*, 2002.
- [25] Rsoft is a CAD Suite for Photonics Device Simulation.