

GLOW: A Global Router for Low-Power Thermal-Reliable Interconnect Synthesis using Photonic Wavelength Multiplexing

Duo Ding, Bei Yu and David Z. Pan
ECE Dept. Univ. of Texas at Austin, Austin, TX 78712
{ding, bei, dpan}@cerc.utexas.edu

ABSTRACT

In this paper, we examine the integration potential and explore the design space of low power thermal reliable on-chip interconnect synthesis featuring nanophotonics Wavelength Division Multiplexing (WDM). With the recent advancements, it is foreseen that nanophotonics holds the promise to be employed for future on-chip data signalling due to its unique power efficiency, signal delay and huge multiplexing potential. However, there are major challenges to address before feasible on-chip integration could be reached. In this paper, we present *GLOW*, a hybrid global router to provide low power opto-electronic interconnect synthesis under the considerations of thermal reliability and various physical design constraints such as optical power, delay and signal quality. *GLOW* is evaluated with testing cases derived from ISPD07-08 global routing benchmarks. Compared with a greedy approach, *GLOW* demonstrates around 23%-50% of total optical power reduction, revealing great potential of on-chip WDM interconnect synthesis.

1. INTRODUCTION

As semiconductor technology roadmap extends into deeper sub-micron domain, the development of future high performance low power systems faces many key challenges. Among them, VLSI interconnect plays more and more critical roles due to: (1) growing ratio of interconnect versus gate delay; (2) higher operating frequency and design complexity; (3) challenging interconnect design for low power systems.

To address the interconnect challenges for future computing systems, various alternative techniques have been proposed as potential solutions (e.g., [1–3]). Among them, nanophotonics devices and interconnect attract active researches (e.g., [3–8]) due to their unique potential to make high speed low power on-chip interconnect.

The many recent advances in nanophotonics devices have demonstrated great on-chip integration potential in nano-scale optical modulators, photo-detectors, couplers, switches, waveguides and WDM (Wavelength Division Multiplexing) devices. Meanwhile researches on photonics device modeling (e.g., [7]) and on-chip integration (e.g., [9–14]) have also introduced new opportunities and challenges to the traditional architectural and physical design methodologies. Specifically, on-chip networks and special architecture designs have been proposed (e.g., [10, 13, 15]) to enable high throughput network communication with on-chip nanophotonic WDM links. Lately, active studies have been carried out to compensate the temperature dependence of nanophotonic devices at both fabrication level (e.g., [16, 17]) and on-chip

network level (e.g., [18, 19]) to assist the design and optimization of thermal-reliable and power-efficient optical-electrical systems-on-chips.

At physical design level, however, studies for efficient on-chip photonics interconnect synthesis have been limited. An early work by [9] employed straight line single channel optical waveguides to perform system-on-package optical routing under some timing consideration. However, physical characteristics were not considered for photonic devices. Important issues such as optical link configuration, loss, thermal reliability and signal integrity were not properly studied. In [14], physical-layer effects (loss, power) are applied to photonics Network-on-Chip performance evaluation. Yet without a systematic CAD environment, it is difficult to design photonics architectures with optimal performance meanwhile taking full advantage of the power budget. In [11, 12], a photonic interconnect library was presented with a physical synthesis framework for low optical power routing. But thermal reliability and WDM mechanisms were not included.

In this paper, we employ nanophotonic on-chip WDM interconnect (Fig. 1) to achieve high density/capacity in the global routing stage. Based on device characterization and modeling, 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. The rest of the paper is organized as follows: in Section 2, we motivate the WDM based optical routing problem under the critical consideration of thermal reliability and summarize the main contributions of this paper. In Section 3 we extend the Optical Interconnect Library (OIL) in [12] by introducing thermal and power related models of nanophotonics devices. In Section 4 we present an overview of our proposed CAD flow, followed by Section 5 and 6, in which we explain the detailed formulation and algorithms for *GLOW* together with an alternative greedy approach *CAT*. Section 7 presents the results, followed by conclusion in Section 8.

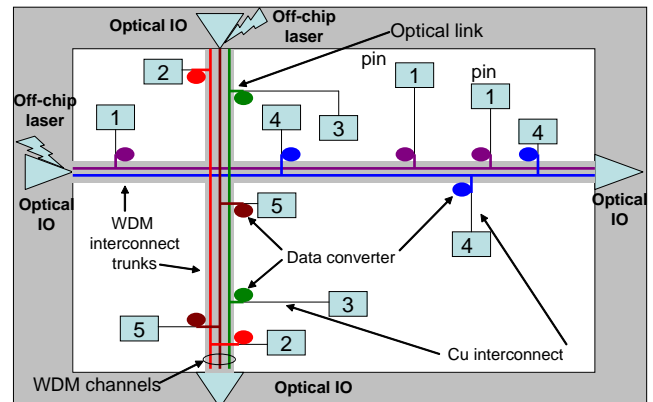


Figure 1: High routing capacity interconnect using WDM

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

Asia and South Pacific Design Automation Conference (ASPDAC), Jan. 30 – Feb. 2, Sydney, Australia, 2012




Figure 2: Thermal-aware WDM routing example

2. MOTIVATION AND CONTRIBUTIONS

With on-chip WDM providing great signal multiplexing capacity, we motivate a global router to take advantage of WDM channels under various physical design constraints such as thermal reliability and timing. A simple scenario is illustrated in Fig. 2. Given a net (A,B,C,D) to be routed with node A as the driver pin, we aim to find a global routing solution in optical-electrical domain to satisfy:

- Thermal reliability and functionality
- Minimal optical driving power required
- Signal integrity and data conversion quality
- Timing considerations and WDM channel utilization
- Legalization based on opto-electrical domain design rules

DEFINITION I. WDM link: A special hybrid on-chip interconnect consisting of electrical wires and optical wavelength division multiplexing devices for large signal capacity.

DEFINITION II. WDM trunk: The nanophotonic WDM waveguide on a WDM link is also referred to as a *trunk*.

DEFINITION III. WDM channel: The carrier of a modulated photonic signal on a WDM trunk. Each *channel* is assigned a unique central wavelength λ . An optical signal can be transmitted on a channel if its wavelength is between $\lambda - \frac{1}{2}BW$ and $\lambda + \frac{1}{2}BW$, where BW is the channel bandwidth.

In Fig. 2, thermal issue refers to the scenario for which on-chip temperature variation causes extra power loss, signal degradation or even malfunction to the nanophotonic devices, such as modulator, photo-detector and WDM waveguide. Without careful planning, an opto-electrical link could fail due to big temperature change. During global routing, regions with excessive thermal variation can be simply set as blockages to avoid. Even for regions with acceptable thermal variation, there are still trade-offs to seek between power efficiency and thermal reliability, e.g., over-optimizing optical power efficiency could result in thermal failure, whereas over-margining the thermal reliability can cause more optical power loss, especially for ring-structure resonators.

Moreover, during routing in the opto-electrical domain, there are extra timing constraints to consider such that the optical-electrical interconnects provide no worse critical path delay than the routing solutions in the electrical domain. For example in Fig. 2, 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, but it is prohibited by the long delay from trunk1 to sink C.

Last but not the least, the final routing solution needs to deliver signals strong enough to be picked up by the photo-detector, meanwhile be legalized according to design rules in both optical and electrical domains.

We summarize key contributions of this paper as follows,

- We propose a systematic CAD framework for on-chip WDM synthesis to co-optimize power and thermal reliability
- We develop a new thermal reliability models for nanophotonics devices considering WDM
- We formulate the optimal global routing problem with Integer Linear Programming technique
- We evaluate the CAD framework with various testcases derived from ISPD global routing benchmarks

3. NANOPHOTONICS DEVICE MODELS

We extend [12] with WDM modules to analyze on-chip optical link configurations, taking into account of power, loss, timing, temperature variation and thermal reliability.

3.1 Device Characterization

Based on current photonics fabrication technology, optical signalling has great advantage over low-K Cu interconnect (11ps versus 37ps per mm on Metal5/6) for global nets. 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 pure Cu interconnects:

$$T_{mod} + T_{det} + \tau_o \cdot L \leq \tau_e \cdot L \quad (1)$$

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. (1) gives us the range of L , whose lower boundary defines L_{crit} value in mm. The devices employed in this paper are summarized in Table 1.

3.2 Thermal Reliability 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 [13, 15, 19] due to its compact footprint (potential ultra density) and demonstrated high quality factor (Q). However, these devices can be very sensitive to ambience temperature change.

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 WDM devices with ring resonator structure, the quality factor Q [22] (defined as the energy stored in the cavity versus the energy dissipated per unit cycle) is very high and BW is very narrow, thus their temperature sensitivity could lead to signal failure. The relationships between thermal rela-

Table 1: Device and interconnect model details

	footprint	speed	on-chip loss	E-power
mod	30X40um	14Gb/s [20]	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 [21].




Figure 3: An overview of our proposed CAD flow

bility, operating BW , quality factor Q and energy efficiency are defined in Eqn.(2)-(4).

$$Q = \frac{\lambda_0}{\Delta\lambda_{FWHM}} = \frac{\sqrt{r_1 r_2} a L \pi n_g}{(1 - r_1 r_2 a) \lambda_0} \quad (2)$$

$$n_g(\lambda) = n_e(\lambda) - \lambda \frac{dn_e(\lambda)}{d\lambda} \quad (3)$$

$$BW = \Delta f = \frac{f_{resonant}}{Q} \quad (4)$$

where 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). Therefore within a relatively small range, we can trade-off Q value for thermal reliability without causing aliasing between WDM channels of trunk. Such a trade-off comes at a power loss penalty that needs to be minimized.

Based on Eqn.(2)-(4), we 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 using [23].

4. OVERALL CAD FLOW

In this section, we present the overall flow of our systematic CAD framework for low power thermal-aware on-chip WDM integration, using the models from Section 3.

In Fig. 3, we illustrate a top level flow diagram of our proposed method, starting from a given input netlist and on-chip temperature variation profile. The flow consists 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. We detail the flow in Fig. 3 as follows.

4.1 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 ISPD07-08 global routing netlists. This step is very critical since it selects proper pins (nets or partial nets) from the




Figure 4: A brief illustration of netlist pre-processing

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 3. The main technique involved is described as follows,

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. 4) and then pick out the 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. 4, 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 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. (5) as a general guideline, until the total number of WDM channels is sufficient to hold the total number of optical nets/links in the netlist.

$$Place_{trunk} = med\{med[net_i]\}^{i \in Partition^k} \quad (5)$$

The partition based initial placement executes in 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 number of links in the optical netlist.
- Extra WDM trunks may need be added in **Post-routing**

4.3 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.

Algorithm 1 *CAT*: Channel Assignment for Thermal reliability

```

Input: (1) Initial WDM trunk placement
Input: (2) Temperature variation profile
Input: (3) Optical netlist
Generate  $L(link)$  as a set of all unassigned links/nets
Generate  $U(link)$  as a set of remaining available link resources
for each WDM  $trunk_i$  do
  for each  $link_j$  in the optical netlist do
    Calculate timing constraint on  $(link_j, trunk_i)$ 
    Calculate thermal variation constraint on  $(link_j, trunk_i)$ 
  end for
  Form set  $S(link_i)$  by links that satisfy Timing Condition
  Sort  $S(link_i)$  based on ascending Thermal Variation
  Select  $link_k \in S(link_i)$  in ascending index order
  while  $link_k$  unassigned AND  $trunk_i$  has available channels do
    Assign  $link_k$  to WDM  $trunk_i$ 
    Update set  $L(link)$ 
  end while
end for
if set  $U(link) = \emptyset$  AND  $L(link) \neq \emptyset$  then
  Revise initial placement with more WDM resources
  Execute CAT on unassigned links
end if
return WDM channel assignment AND optical power

```

Fig. 5 illustrates the routing problem after **Pre-routing** stage, with off-chip laser sources whose driving power to each WDM waveguide trunk varies depending on the total number of channels assigned. To constrain the solution space for the global routing stage, we take the shortest distance route when a pin is to connect to certain WDM trunk, i.e., data converters (mod/det) are placed along WDM trunks.

For the core formulation of the **Global Routing** stage, we propose 2 algorithms: (1) global routing with low power WDM (*GLOW*) as a major contribution of this paper; (2) channel assignment under thermal consideration (*CAT*) as a comparison baseline. Details are in Section 5 and 6.

4.4 Post Routing Legalization

This stage is mainly to resolve the cases when multiple rings are contending the same geometry location, causing design rule violations in the optical domain. For this paper, we use simple perturbation based re-routing/adjustment techniques, leaving other perspectives to detailed routing stages.

5. CAT ALGORITHM

CAT is designed as a greedy approach for Channel Assignment under Thermal considerations. The basic motivation 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




Figure 5: Our WDM based global routing scenario

fill the available channels to certain WDM waveguide, then move onto the next waveguide. If eventually there are still remaining nets unassigned, then the *Initial Placement* stage will be repeated to include more WDM trunks/channels.

CAT is performed in 3 major steps: first, *Initial WDM Trunk Placement*; second, *Timing and Thermal Condition Calculation*; third, *Greedy Channel Assignment*.

Initial WDM Trunk Placement: The same as Section 4.2, which is used for both *CAT* and *GLOW*.

Timing and Thermal Condition Calculation: In this step, all the WDM trunks are traversed in certain order sequentially. For each trunk, timing/thermal conditions for all optical links are calculated and updated.

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 C_{max} . For more details of *CAT*, please refer to Algorithm 1.

6. GLOW ROUTING ALGORITHM

6.1 ILP Formulation

To formulate the optimal optical routing problem, we represent the assignment status of channels on WDM trunks with integer binary variables (occupied or available) and use the cross-term variables to model the optical power loss introduced by WDM signal crossings, which is otherwise very difficult to accurately characterize before the channel assignment takes place. Table 2 details all the variables defined.

- 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 unassigned 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 $W_i \cdot W_j$, where $i \in [0, n-1]$, $j \in [n, n+m-1]$. 0 meaning trunk i and trunk j are not physically crossed; vise 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, 0 meaning net i is assigned onto WDM trunk j in the global routing; vise versa.

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

$$\text{Minimize}\{P_{total}\} \text{ w.r.t } W_i, W_{ij}, S_{link_i}^{trunk_j}, \lambda_{net_i}^{trunk_j} \quad (6)$$

where

$$P_{total} = P_{loss} + P_{dynamic} \quad (7)$$

$$P_{loss} = P_{cross} + P_{trunk_thm} + P_{ring_thm} + P_{path} \quad (8)$$

$$P_{cross} = \sum_{i \in [0, n-1]} \sum_{j \in [n, n+m-1]} W_{ij} * P_{thm}^{ij} \quad (9)$$

$$P_{trunk_thm} = \sum_i^{i \in \text{all trunks}} W_i * P_{trunk_thm}^i \quad (10)$$

Table 2: Variables/parameters in ILP formulation

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
$P_{\lambda i}$	laser power on channel λi for optical signalling
P_{thm}^{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
$S_{link_j}^{trunk_j}$	BV: assignment status of link i onto trunk j
$Sum_{net_j}^{trunk_j}$	IV: # of links in net i assigned to trunk j
$\lambda_{net_j}^{trunk_j}$	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
$temp^{threshold}$	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_e^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

$$P_{ring_thm} = \sum_i \sum_{j \in \text{all links}} S_{link_j}^{trunk_j} * P_{ring}^{link_j} \quad (11)$$

$$P_{dynamic} = \sum_i \sum_{j \in \text{all nets}} \lambda_{net_j}^{trunk_j} P_{\lambda i} + \sum_i W_i P_0 \quad (12)$$

Eq. (6) above gives the objective function of *GLOW* as the total power P_{total} required to drive the circuit. As shown in Eq. (7), 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 Section 3 and Section 4. Table 2 further details each term.

6.2 Physical Design Constraints

We present the detailed mathematical formulations of various routing constraints for *GLOW* as follows:

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

$$S_{link_i}^{trunk_j} [\tau_e * WL_e^i + \tau_o * WL_o^i + \tau_{conv}] \leq \tau_e * HPWL^{link_i} \quad (13)$$

- Selection constraint: to make sure each link i is only

Algorithm 2 *GLOW*: Global Routing for Low Power WDM

Input: (1) Initial WDM trunk placement
Input: (2) Temperature variation profile
Input: (3) Optical netlist
for each WDM trunk i **do**
 for each optical link j **do**
 Calculate P_{cross} , $P_{WDM_thermal}$, $P_{ring_thermal}$
 Calculate $P_{dynamic}$
 Update *Timing Constraint*, *Thermal Constraint*
 end for
end for
Invoke ILP solver
return WDM channel assignment AND optical/laser power

assigned to one WDM trunk. For each link i :

$$\sum_{j \in \text{all trunks}} S_{link_i}^{trunk_j} = 1 \quad (14)$$

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

$$\sum_i \lambda_{net_i}^{trunk_j} \leq C_{max} \quad (15)$$

- 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 be upper bounded by $temp^{threshold}$ to avoid performance degradation or malfunction. For each link i and trunk j :

$$S_{link_i}^{trunk_j} * T_{var}^{link_i} \leq temp^{threshold} \quad (16)$$

- 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:

$$2W_{ij} \leq W_i + W_j \leq 1 + W_{ij} \quad (17)$$

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

$$\frac{(2 \sum_{k \in net_i} S_{link_k}^{trunk_j} - 1)}{2PIN_{max}} \leq \lambda_{net_i}^{trunk_j} \leq 2 \sum_{k \in net_i} S_{link_k}^{trunk_j} \quad (18)$$

$$\frac{(2 \sum_{i=1}^{\text{all nets}} \lambda_{net_i}^{trunk_j} - 1)}{2C_{max}} \leq W_j \leq 2 \sum_{i=1}^{\text{all nets}} \lambda_{net_i}^{trunk_j} \quad (19)$$

Here Equation(18) and (19) 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(20):

$$Sum_{net_i}^{trunk_j} = \sum_k S_{link_k}^{trunk_j} \quad (20)$$

Equation(18)(19)(20) 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 $\lambda_{net_i}^{trunk_j} = 1$. They serve as a binary comparator meanwhile still preserve the linear programming formulation. Algorithm 2 summarizes the pseudo-codes of the main steps of *GLOW*.

7. SIMULATION AND TESTING

Benchmarks and Simulation Setups: In Table 3 we list 6 optical benchmarks: CK1-6, with net number ranging

Table 3: Simulation result comparisons between our proposed *CAT* and *GLOW*

Method	CAT						GLOW					
Optical Netlist	CK1	CK2	CK3	CK4	CK5	CK6	CK1	CK2	CK3	CK4	CK5	CK6
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. Chan/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 reduc.%	-	-	-	-	-	-	31.0%	47.0%	22.6%	47.5%	39.0%	49.8%

^a Each WDM trunk has up to 32 available channels at initial placement. Unassigned trunks will be turned off after routing.

^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*.

from 35 to 996. These test cases are derived from IPSD07-08 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 3, 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 around 23%-50% 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 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 globally; on the other hand, the ILP formulation of *GLOW* allows us to model all the key factors of optical power.

Also in Table 3 we show the WDM channel/trunk allocation 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 costs of straight OWGs are relatively low meanwhile there are resources on the silicon layer for the nanophotonics integration.

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 the 6 optical netlists in about 0.2 to 0.9 hours. Such run-time is acceptable as the optical routing problem size is fairly small, i.e., only the top global nets/pins are mapped into the optical domain while the rest nets are routed in the electrical domain.

8. CONCLUSION

In this paper we explored the synthesis of on-chip nanophotonic WDM interconnects and presented *GLOW*, a low power optical router under thermal-reliability considerations. It is evaluated on various testing cases showing significant optical power reduction compared with a baseline greedy method. We believe a lot of future research can be done to co-optimize the CAD and the nanophotonics technologies in the physical design area.

9. 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 *Proc. of the 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] Jacob R. Minz et al. Optical Routing for 3D System-on-Package. In *Proc. Design, Automation and Test in Europe*, 2006.
- [10] Assaf Shacham et al. Photonic Networks-on-Chip for Future Generation Chip Multiprocessors. In *IEEE Trans. on Computers*, 2008.
- [11] Duo Ding et al. O-Router: An Optical Routing Framework for Low Power On-Chip Silicon Nano-photonics Integration. In *Proc. Design Automation Conf.*, 2009.
- [12] Duo Ding et al. OIL: A Nanophotonic Optical Interconnect Library for a New Photonic Networks-on-Chip Architecture. In *Proc. System Level Interconnect Prediction*, 2009.
- [13] Yan Pan et al. Firefly: Illuminating Future Network-on-Chip with Nanophotonics. In *Proc. Int. Symp. on Computer Architecture*, 2009.
- [14] Johnnie Chan et al. PhoenixSim: A Simulator for Physical-Layer Analysis of Chip-Scale Photonic Interconnection Networks. In *Proc. Design, Automation and Test in Europe*, 2010.
- [15] Ajay Joshi et al. Silicon-Photonic Clos Networks for Global On-Chip Communication. In *Int. Symp. on Networks-on-Chip*, 2009.
- [16] Jong-Moo Lee et al. Controlling Temperature Dependence of Silicon Waveguide Using Slot Structure. In *Optics Express*, 2008.
- [17] Biswajeet Guha et al. CMOS-Compatible Athermal Silicon Microring Resonators. In *Optics Express*, 2010.
- [18] Moustafa Mohamed et al. Power-Efficient Variation-Aware Photonic On-Chip Network Management. In *Int. Symp. on Low Power Electronics Design*, 2010.
- [19] 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.
- [20] Po Dong et al. Low Vpp, Ultralow-energy, Compact, High-speed Silicon Electro-Optic Modulator. In *OPTICS EXPRESS*, 2009.
- [21] Predictive Technology Model, <http://ptm.asu.edu>.
- [22] Payam Rabiei et al. Polymer Micro-Ring Filters and Modulators. In *J. of Lightwave Technology*, 2002.
- [23] Rsoft is a CAD Suite for Photonics Device Simulation.