Bidirectional Tuning of Microring-Based Silicon Photonic Transceivers for Optimal Energy Efficiency

Yuyang Wang
University of California
Santa Barbara, CA, U.S.A.
wyy@ece.ucsb.edu

M. Ashkan Seyedi
Hewlett Packard Labs
Hewlett Packard Enterprise
Palo Alto, CA, U.S.A.

Jared Hulme
Hewlett Packard Labs
Hewlett Packard Enterprise
Palo Alto, CA, U.S.A.

Marco Fiorentino
Hewlett Packard Labs
Hewlett Packard Enterprise
Palo Alto, CA, U.S.A.

Raymond G. Beausoleil
Hewlett Packard Labs
Hewlett Packard Enterprise
Palo Alto, CA, U.S.A.

Kwang-Ting Cheng
Hong Kong University of
Science and Technology
Hong Kong

ABSTRACT

Microring-based silicon photonic transceivers are promising to resolve the communication bottleneck of future high-performance computing systems. To rectify process variations in microring resonance wavelengths, thermal tuning is usually preferred over electrical tuning due to its preservation of extinction ratios and quality factors. However, the low energy efficiency of resistive thermal tuners results in nontrivial tuning cost and overall energy consumption of the transceiver. In this study, we propose a hybrid tuning strategy which involves both thermal and electrical tuning. Our strategy determines the tuning direction of each resonance wavelength with the goal of optimizing the transceiver energy efficiency without compromising signal integrity. Formulated as an integer programming problem and solved by a genetic algorithm, our tuning strategy yields 32%∼53% savings of overall energy per bit for measured data of 5-channel transceivers at 5∼10 Gb/s per channel, and up to 24% saving for synthetic data of 30-channel transceivers, generated based on the process variation models built upon measured data. We further investigated a polynomial-time approximation method which achieves over 100x speedup in tuning scheme computation, while still maintaining considerable energy-per-bit savings.

KEYWORDS

Silicon photonics, optical interconnects, microring tuning, energy efficiency.

ACM Reference Format:
Yuyang Wang, M. Ashkan Seyedi, Jared Hulme, Marco Fiorentino, Raymond G. Beausoleil, and Kwang-Ting Cheng. 2019. Bidirectional Tuning of Microring-Based Silicon Photonic Transceivers for Optimal Energy Efficiency. In *24th Asia and South Pacific Design Automation Conference (ASPDAC '19), January 21–24, 2019, Tokyo, Japan.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3287624.3287649

ASPDAC '19, January 21–24, 2019, Tokyo, Japan

ACM ISBN 978-1-4503-6007-4/19/01…$15.00
https://doi.org/10.1145/3287624.3287649

1 INTRODUCTION

With the benefits in throughput and energy efficiency brought by dense wavelength division multiplexing (DWDM), optical interconnects have been proposed to accommodate traffic-intensive applications in future high-performance computing systems [1]. Silicon photonics is emerging as a cost-effective and scalable solution to optical interconnects by taking advantage of large-scale CMOS-compatible integration [2]. A promising architecture incorporates innovations in quantum-dot (QD) comb lasers [3] and silicon photonic microring resonators [4] to achieve concurrent DWDM [5, 6].

Microring resonators are highly wavelength-selective devices which can be used to modulate or filter optical signals at their resonance wavelengths [7, 8]. Due to process variations, the fabricated resonance wavelengths can deviate significantly from design values [9], and thus require active post-fabrication tuning to align with the carrier wavelengths. For p-i-n junction-based, carrier-injection microring resonators [4], the electro-optic (EO) and thermal-optic (TO) effects shift the resonance wavelengths in opposite directions [10, 11], which provides such devices with inherent capability of bidirectional tuning. Despite better energy-efficiency of electrical tuning as compared to thermal tuning, the increase of electron-hole pair density introduced into the cavity by electrical tuning results in degradations of extinction ratios (ER) and quality factors (Q) of the transmission spectra [4, 12, 13], which makes electrical tuning less preferable. However, if the required tuning distance is small and thus the degradations of ER and Q are limited, such degradations can be compensated by the increase of laser power to maintain transmission quality, if the overall power consumption of the transceiver (TRx) is still lower than that of thermal tuning. Nevertheless, such opportunities have not been explored in microring-based DWDM applications which are currently dominated by thermal tuning schemes [14–18].

One of the greatest challenges faced by thermal tuning schemes is the nontrivial power consumption [19]. Techniques based on channel remapping [9, 20] and channel redundancies [21–23] were proposed to mitigate the expected tuning power. Wafer-level variations were also exploited in [24] to further reduce the average tuning power when a large number of transceivers are available. However, the inherent low energy efficiency of resistive thermal tuners limits the effectiveness of these techniques as long as all-thermal tuning schemes are adopted.

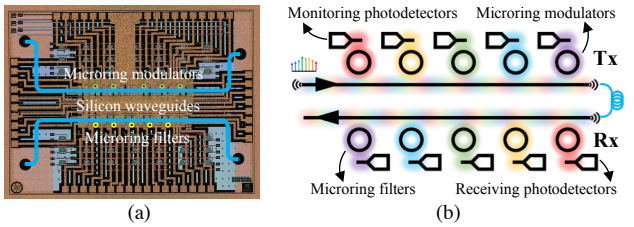

(a) (b)

Figure 1: (a) Microscopic image and (b) architecture illustration of a fabricated microring-based transceiver.

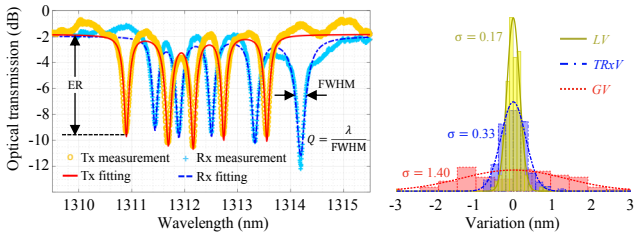

(a) (b)

Figure 2: (a) Measured and fitted spectra of a 5-channel transceiver, where FWHM: *full width at half maximum*. (b) Variation characterization for resonance wavelengths.

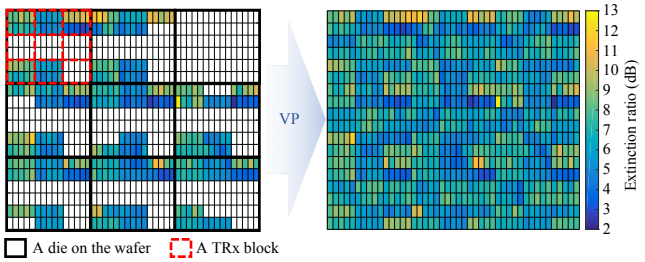

Figure 3: Relative locations of measured ER (left) and predicted values for unmeasured microrings by Virtual Probe (right).

In this study, we propose a hybrid tuning strategy for carrier-injection microring-based transceivers which involves both thermal and electrical tuning. Our strategy accounts for the potential degradations of ER and Q induced by electrical tuning, and determines the tuning direction of each resonance wavelength with the goal of optimizing the overall energy efficiency of the transceiver link. We solved for the optimal tuning scheme using a genetic algorithm (GA) and a polynomial-time approximation method. Tuning schemes given by both methods were evaluated on measured data from a batch of 5-channel transceivers, as well as synthetic data, generated based on process variation models built from measured data, of 5~30-channel transceivers, showing considerable savings of overall energy per bit.

2 MICRORING-BASED TRANSCEIVER

2.1 Overview

Our bidirectional tuning strategy targets the optical transceivers made of p-i-n junction-based, carrier-injection microring resonators. The transceivers that we used for variation characterization were fabricated by CEA-Leti on a 200 mm silicon-on-insulator (SOI) wafer. As illustrated in Fig. 1, each fabricated transceiver consists of five high-speed microring modulators in the transmitter (Tx) for on-off keying (OOK) modulation, and five corresponding microring filters in the receiver (Rx) for demultiplexing. The channel spacing of the transceiver is designed to be 80 GHz in the 1310 nm regime (O-band). The free spectrum range (FSR) of the microring is measured as ~13.9 nm, which can accommodate up to 30 multiplexed channels. In total, 45 transceivers located on 9 dies were sampled for measurement, in which 5 transceivers showed no clear resonance dips, and another 6 transceivers failed in either Tx or Rx. So 34 transceivers were successfully measured and characterized.

2.2 Variation Characterization

The optical signal travels through the *through port* of each cascaded microring in the Tx, and is coupled out at the *drop port* of a microring in the Rx, before sensed by the receiving photodetector. Process variations of the through/drop-port transmission spectra are characterized as follows.

2.2.1 Through-port spectrum. We model the through-port transmission spectrum of microring #i with a Lorentzian function-based equation, characterized by its resonance wavelength $\lambda_{r,i}$, extinction ratio $ER_{\text{thru},i}$, quality factor $Q_{\text{thru},i}$, and a through-port loss $\alpha_{\text{thru},i}$ [12]:

$$T_{\text{thru},i}(\lambda) = \alpha_{\text{thru},i} \cdot \left(1 - \frac{1 - 1/ER_{\text{thru},i}}{1 + \left(2Q_{\text{thru},i} \cdot (\lambda - \lambda_{r,i})/\lambda_{r,i}\right)^2}\right) \quad (1)$$

The through-port spectrum of an n-channel Tx/Rx is modeled as the product of the through-port spectrum of each cascaded microring:

$$T_n(\lambda) = \prod_{i=1}^{n} T_{\text{thru},i}(\lambda) = \alpha \cdot \prod_{i=1}^{n} \left(1 - \frac{1 - 1/ER_{\text{thru},i}}{1 + \left(2Q_{\text{thru},i} \cdot (\lambda - \lambda_{r,i})/\lambda_{r,i}\right)^2}\right) \quad (2)$$

where $\alpha = \prod_{i=1}^{n} \alpha_{\text{thru},i}$. Fig. 2(a) shows an example of the measured and fitted transceiver spectra, with variations observed in resonance wavelengths, extinction ratios, quality factors, as well as overall through-port losses.

We adopted the process variation model for resonance wavelengths proposed in [23], where each resonance wavelength is decomposed into the design value and three independent variation components, namely the global variation (GV), the local variation (LV), and the Tx-Rx offset (TRxV). All of the three variation components are approximated by normal distributions with zero mean, as shown in Fig. 2(b). The loss factor α is also approximated by a normal distribution with a mean of 0.68 and a standard deviation of 0.06.

Contrarily, location dependencies were observed in the process variations of ER and Q, prohibiting them from being modeled with independent variation components. Instead, we employed Virtual Probe (VP) [25, 26] to capture the spatial patterns from measured data, and predict the values at other locations as our synthetic data for ER and Q. Fig. 3 shows the application of VP on measured ER as an example.

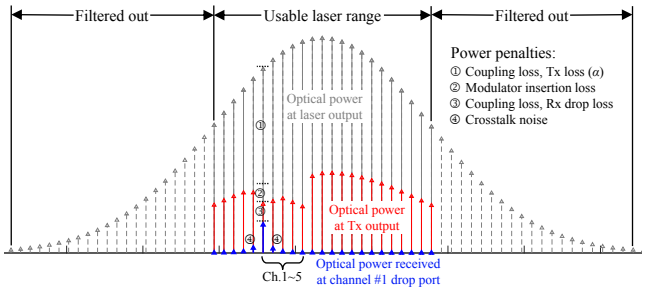

Figure 4: Illustration of the laser spectrum and power losses in a 5-channel transceiver.

2.2.2 Drop-port spectrum. The drop-port transmission spectrum of microring #i is also modeled as a Lorentzian function characterized by its resonance wavelength, extinction ratio, quality factor, and a drop-port loss $\alpha_{\text{drop},i}$:

$$T_{\text{drop},i}(\lambda) = \alpha_{\text{drop},i} \cdot \frac{1 - 1/ER_{\text{drop},i}}{1 + \left(2Q_{\text{drop},i} \cdot (\lambda - \lambda_{r,i})/\lambda_{r,i}\right)^2} \tag{3}$$

Process variations for each parameter are characterized using the same method described in Section 2.2.1. Finally, synthetic data of transceivers with up to 30 channels are generated and used, together with measured data, to evaluate our bidirectional tuning strategy.

3 TRANSCEIVER LINK POWER MODELS

The total power consumption of a transceiver link for data communication includes those consumed by the comb laser, resonance wavelength tuning, modulator drivers, receiver transimpedance amplifiers (TIA), and serializer/deserializer (SerDes) circuitry. The energy efficiency of the transceiver is measured as the power consumption divided by the aggregated data rate, in the unit of *p*J/bit, for which the smaller the value the better.

3.1 Laser Power Budget

The QD comb laser is capable of generating a group of evenly-spaced frequency combs. However, the optical power of each comb line is usually different [3]. A Gaussian-shaped comb spectrum is assumed, with a spectrum efficiency $\eta = P_{\text{usable}}/P_{\text{total}} \approx -3.2$ dB [27], as illustrated in Fig. 4. The optical power provided at the laser output should be high enough so that the following power budget equation holds for any TRx channel [10]:

$$P_{\text{comb line}} \cdot \prod_i PL_i \geqslant P_{\text{sensitivity}} \tag{4}$$

Here, $P_{\text{comb line}}$ is the optical power of the laser comb line for a particular channel; $PL_i \in (0, 1)$ denotes the various power losses along the transceiver link, including coupling loss, Tx through-port loss (α), microring insertion loss, microring drop-port loss, crosstalk penalties, etc (Fig. 4); $P_{\text{sensitivity}}$ is the receiver sensitivity requirement, of which a relationship with the target data rate is given in [28]. The laser power budgets computed from Eq. (4) at various data rates show good consistency with both actual measurements [6, 29] and technology projections [3, 30].

Recent high-efficiency QD comb lasers have achieved up to 7-9 dBm per comb line at ~20% wall-plug efficiency (WPE) [3]. It is noteworthy that the optical nonlinearities of the microrings and the silicon waveguides limit the optical power that can be injected into the transceiver [31]. In this study, any transceiver that requires over 5 dBm optical power per channel or over 20 dBm optical power per waveguide is considered unusable.

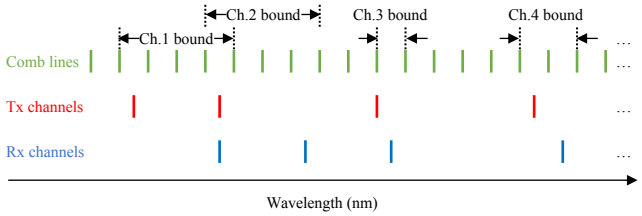

Figure 5: Bounds of candidate carrier wavelengths for each TRx channel.

3.2 Microring Wavelength Tuning

The microring tuning power is computed according to the tuning direction of the channel. For thermal tuning, a tuner efficiency of 0.15 nm/mW is assumed [32]. For electrical tuning, the EO effect model for carrier-injection microrings proposed in [12] is utilized to convert the tuning distance into current. Then, the electrical tuning power can be computed based on the microring DC model [12]. Meanwhile, the degradations of ER and Q of the microring can also be derived, and compensated by the increase of laser power budget based on the inequality of Eq. (4) described in Section 3.1.

3.3 Power Models for Other Components

The power consumptions of modulator drivers, receiver TIAs and SerDes circuitry mainly depend on the data rate per channel. In this study, lookup tables of power consumptions of these components are made for different data rates, based on the values reported in [19].

4 PROBLEM FORMULATION

Our bidirectional tuning strategy selects the carrier wavelengths to which the TRx channels are to be aligned. The difference between the resonance wavelengths before and after tuning determines the degradations of ER and Q, and ultimately the shape of the TRx spectra. Therefore, given the spectrum-related parameters of a TRx before tuning, the overall energy-per-bit consumption of the TRx after tuning can be seen as a function of the selected carrier wavelengths:

$$\text{Energy per bit} = E\left(\vec{\lambda}_{\text{carrier}}\right) \tag{5}$$

The pool of candidate carrier wavelengths is a set of comb lines provided by the QD comb laser:

$$\text{Combs} = \{\lambda_1, \lambda_2, \ldots, \lambda_L\} \tag{6}$$

where L is the number of usable comb lines. For an n-channel transceiver, the overall energy-per-bit value can only be computed once n carrier wavelengths are all determined. The goal is then to find $\vec{x} = [x_1, x_2, \ldots, x_n]$ where $x_i \in \mathbb{Z} \cap [1, L]$, such that:

$$E\left([\lambda_{x_1}, \lambda_{x_2}, \ldots, \lambda_{x_n}]\right) \text{ is minimized.}$$

The vector $\vec{x}$, with all integer elements, is in fact the indices of the laser comb lines that are selected as carrier wavelengths. Our strategy preserves the order of TRx channels before and after tuning, by forcing $x_i < x_{i+1}$ during the optimization.

To shrink the size of the problem space, we set the lower and upper bounds of candidate carrier wavelengths for each transceiver

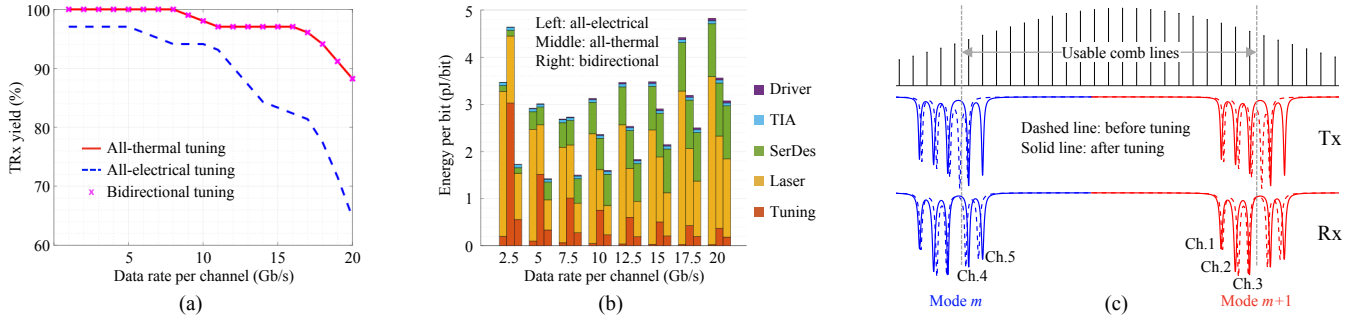

Figure 6: Evaluation of bidirectional tuning on measured data: (a) yield comparison; (b) energy per bit comparison; and (c) bidirectional tuning opportunities identified by our strategy.

channel, as shown in Fig. 5. For the ith channel, the lower (upper) bound of its candidate carrier wavelengths lb_i (ub_i) is the index of the nearest laser comb line to the left (right) of both resonance wavelengths of the Tx and Rx microrings. Finally, our bidirectional tuning strategy is formulated as an integer programming problem of finding $\vec{x}$ subject to a lower bound $\overrightarrow{lb}$, an upper bound $\overrightarrow{ub}$, and a linear inequality constraint $A\vec{x} < \vec{b}$ (Eq. (7)), with the goal of minimizing the overall energy per bit E.

$$\overbrace{\begin{bmatrix} 1 & -1 & 0 & \dots & 0 & 0 \\ 0 & 1 & -1 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & -1 \end{bmatrix}}^{A} \overbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{bmatrix}}^{\vec{x}} < \overbrace{\begin{bmatrix} -1 \\ -1 \\ -1 \\ \vdots \\ -1 \end{bmatrix}}^{\vec{b}} \tag{7}$$

The NP-completeness of integer programming problems has been proven in [33]. In this study, we employed an implementation of the genetic algorithm in MATLAB to solve for the optimal wavelength tuning scheme.

5 BIDIRECTIONAL TUNING EVALUATION

5.1 Evaluation on Measured Data

We first evaluate our bidirectional tuning strategy on measured data of 5-channel transceivers. The tuning scheme given by our strategy is compared to all-electrical and all-thermal tuning schemes, as shown in Fig. 6(a) and Fig. 6(b). Due to severe degradations of ER and Q and thus explosive increase in laser power budgets, all-electrical tuning results in fewer usable transceivers (and thus a lower yield) and a higher energy-per-bit consumption compared to all-thermal tuning, especially at a higher data rate per channel. On the contrary, our bidirectional tuning scheme does not compromise the yield of the TRx at any data rate per channel compared to all-thermal tuning. Since all-thermal tuning is essentially a special case of bidirectional tuning, and is automatically chosen by our strategy when all other tuning schemes result in either worse energy-per-bit consumption or device failure, it serves as the worst case of our bidirectional tuning strategy.

The application of our strategy on measured data also brings significant energy-per-bit savings, as shown in Fig. 6(b). For commonly used channel rates in the range of 5~10 Gb/s, which lead to relatively low energy per bit consumptions, our strategy (the right bar of each 3-bar group) results in 32%~53% savings of overall energy per bit compared to all-thermal tuning (the middle bar), with ~90% of measured transceivers identified to benefit from bidirectional tuning.

Non-intuitively, our tuning strategy does not necessarily lead to a higher laser power consumption compared to all-thermal tuning, despite some degradations of ER and Q induced by electrical tuning in selected channels. The reason lies in the non-uniform laser power spectrum, where laser comb lines near the center inherently provide higher optical power than that of the outer ones. Fig. 6(c) shows an example of bidirectional tuning opportunities identified by our strategy. Electrical tuning on channel 1 and 2 aligns them to laser comb lines with higher optical power, with only minor degradations of ER and Q. It is noteworthy that our bidirectional tuning strategy has no knowledge of the Gaussian-shaped laser spectrum, and should be able to capture such opportunities for other laser spectrum models (e.g. the flat-comb model in [27]).

5.2 Evaluation on Synthetic Data

We further evaluate our bidirectional tuning strategy on synthetic data of transceivers with 5~30 channels. Synthetic spectra of 81 transceivers, based on the variation models derived from measured data, are generated for each configuration. Fig. 7 summarizes the energy-per-bit savings of our bidirectional tuning strategy over all-thermal tuning. Up to 56% saving can be attained for 5-channel transceivers at 5~10 Gb/s per channel, showing consistency with the simulation results on measured data. As the channel count increases, the free spectrum range of the microring is packed with more channels, which reduces the flexibility in carrier wavelength selection due to the inequality constraint expressed in Eq. (7) of Section 4. As a result, the attainable energy-per-bit saving of our bidirectional tuning strategy decreases for transceivers with higher channel counts. However, ~24% energy-per-bit saving is still observed for 30-channel transceivers at 5 Gb/s per channel when bidirectional tuning is applied.

The problem space of bidirectional tuning grows exponentially with the number of channels, and thus requires a longer time for the genetic algorithm to converge. One of the stopping criteria is *MaxStallGenerations*, which sets the maximum number of generations over which the best value of the objective function stalls before the algorithm terminates. In Fig. 8, we explored the attainable energy-per-bit saving of our strategy and the computation time of

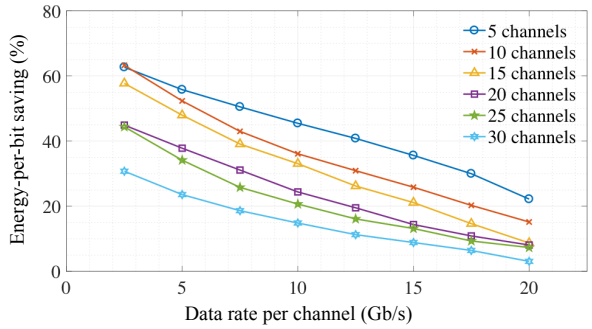

Figure 7: Evaluation of energy-per-bit savings of bidirectional tuning vs. all-thermal tuning on synthetic data of 5~30-channel transceivers.

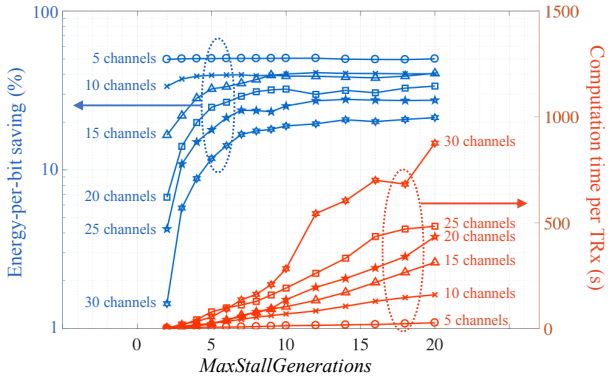

Figure 8: Energy-per-bit savings vs. computation time for bidirectional tuning.

the algorithm for various channel counts and *MaxStallGenerations*. The data rate is set at 7.5 Gb/s per channel. It is observed that a *MaxStallGenerations* of ~10 is large enough to approximate the maximum energy-per-bit saving, without incurring unnecessarily long computation time.

It is worth mentioning that in practice, thermal tuners are fabricated for every microring, and electrical tuning is essentially a DC bias applied through the driver. Since the optimal tuning scheme can be computed once the transceiver is fabricated and measured, our bidirectional tuning scheme does not increase the design complexity of the transceiver compared to all-thermal tuning, nor does it require sophisticated run-time reconfiguration support from the driving circuitry.

6 A POLYNOMIAL-TIME APPROXIMATION

The computation of the optimal tuning scheme adds to the already long testing time, for which any extra time cost, if nontrivial, should be minimized. In this section, we investigate a polynomial-time approximation method to shorten the computation time of our bidirectional tuning strategy. Instead of computing all carrier wavelengths as a whole, our approximation method simplifies the problem by assuming that wavelength tuning of one channel has minor influence on its neighboring channels, and thus determines the carrier

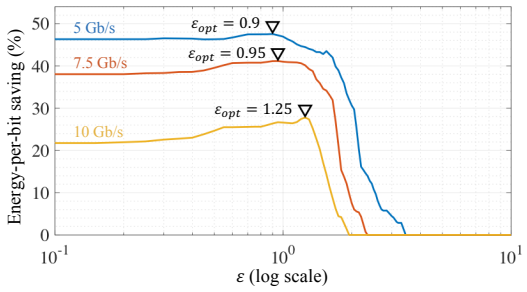

Figure 9: Effect of ε on the energy-per-bit savings of the polynomial-time approximation method.

Listing 1: A polynomial-time approximation method for bidirectional tuning

```
// n is the number of channels
Find lb⃗ and ub⃗ of each carrier wavelength;
// Determine carrier wavelengths one channel at a time
carrier = [];
for j = 1 to n do {
  Initialize best energy consumption Emin = Inf;
  for i = lb(j) to ub(j) do {
    Compute energy E if comb line i is selected;
    if E <= ε · Emin do {
      Emin = E; // Update best energy consumption
      carrier(j) = i; // Select this comb line
      if j < n do {
        // Update bounds for next channel
        lb(j+1) = max(lb(j+1), i+1);
        ub(j+1) = max(ub(j+1), lb(j+1));
      }
    }
  }
}
Compute energy consumption based on selected comb lines;
If worse than all-thermal tuning, use all-thermal tuning;
```

wavelength for each channel one at a time. Listing 1 outlines the main steps of the approximation method. The energy per bit computation takes $O\left(n^2\right)$ time, resulting in an overall time complexity of $O\left(n^3\right)$ for this method.

Note that the approximation method greedily selects the carrier wavelength for each channel, which may lead to a local minimum. In order to explore a larger solution space, we introduce a parameter ε which, intuitively, is the tendency of the algorithm to choose thermal tuning over electrical tuning for a channel when both schemes have comparable energy-per-bit consumptions. An ε smaller than 1 encourages the algorithm to consider electrical tuning even if it is suboptimal for a channel, yet to make room for subsequent channels on its right which may actually benefit from electrical tuning.

Fig. 9 shows the energy-per-bit savings of the tuning schemes given by this approximation method compared to all-thermal tuning, evaluated on measured data of 5-channel transceivers at various data rates of 5~10 Gb/s per channel. The energy-per-bit savings at optimal ε for this range of data rates were identified to be 28%~48%, which is quite comparable to the 32%~53% savings derived by the

Table 1: Polynomial-time approximation vs. genetic algorithm evaluated on synthetic data at 7.5 Gb/s per channel.

Number of Channels	Energy-per-bit saving (%)		Computation time per TRx (s)		Speedup
	GA	Approx.	GA	Approx.	
5	51	46	12.817	0.107	119x
10	43	37	67.362	0.304	221x
15	39	35	101.195	0.495	204x
20	31	27	190.273	0.893	213x
25	26	23	108.929	1.027	106x
30	17	15	310.796	2.805	110x

genetic algorithm in Section 5.1. However, the average computation time for the tuning scheme of each TRx reduces drastically compared to the GA method.

We further evaluate the approximation method on synthetic data of transceivers with higher channel counts to explore the trade-off between the speedup and the attainable energy-per-bit saving. A data rate of 7.5 Gb/s per channel is assumed in the experiments. The computation time of the GA method is measured with a *MaxStallGenerations* of 10. Table 1 lists the comparison between the approximation method and the GA method evaluated across various channel counts. As observed from the simulation results, the polynomial-time approximation method is able to derive a tuning scheme over 100x faster than the GA method, while losing only 2∼6 percentage points of the attainable energy-per-bit saving.

7 CONCLUSION

To break the inherent bottleneck of energy-inefficient thermal tuners in microring-based optical links, we propose a bidirectional tuning strategy for carrier-injection microring-based transceivers which incorporates a hybrid of both thermal and electrical tuning. Our strategy determines the tuning direction of each TRx channel with the goal of minimizing the overall energy-per-bit consumption. The potential degradations of ER and Q are well accounted for, thus incurring no compromise in signal integrity. By formulating our strategy as an integer programming problem, we applied a genetic algorithm (GA) to solve for the optimal tuning schemes under various system configurations. The derived schemes lead to significant savings of overall energy per bit for measured and synthetic data of transceivers with various channel counts, based on commonly used data rates ranging from 5 to 10 Gb/s per channel. We further explored a polynomial-time approximation method which speeds up the computation of bidirectional tuning schemes for over 100x, which adds negligible overhead on device testing time, while losing only minor percentage points of attainable energy-per-bit savings.

ACKNOWLEDGMENT

This study is partially sponsored by American Institute for Manufacturing (AIM) Integrated Photonics.

REFERENCES

[1] R. G. Beausoleil, M. McLaren, and N. P. Jouppi, "Photonic architectures for high-performance data centers," *IEEE J. Sel. Top. Quantum Electron.*, vol. 19, no. 2, pp. 3 700 109–3 700 109, Mar. 2013.
[2] R. G. Beausoleil, "Large-scale integrated photonics for high-performance interconnects," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 7, no. 2, pp. 1–54, Jun. 2011.
[3] D. Livshits *et al.*, "High efficiency diode comb-laser for dwdm optical interconnects," *2014 IEEE Opt. Interconnects Conf. OI 2014*, vol. 6, pp. 83–84, 2014.
[4] Q. Xu *et al.*, "Micrometre-scale silicon electro-optic modulator," *Nature*, vol. 435, no. 7040, pp. 325–327, May 2005.
[5] C. H. Chen *et al.*, "Concurrent multi-channel transmission of a dwdm silicon photonic transmitter based on a comb laser and microring modulators," in *2015 International Conference on Photonics in Switching (PS)*, 2015, pp. 175–177.
[6] M. A. Seyedi *et al.*, "Concurrent dwdm transmission with ring modulators driven by a comb laser with 50ghz channel spacing," in *2016 21st OptoElectronics and Communications Conference (OECC) held jointly with 2016 International Conference on Photonics in Switching (PS)*, 2016, pp. 1–3.
[7] Q. Xu *et al.*, "Cascaded silicon micro-ring modulators for wdm optical interconnection." *Opt. Express*, vol. 14, no. 20, pp. 689 812–689 812–10, 2006.
[8] S. Manipatruni, L. Chen, and M. Lipson, "Ultra high bandwidth wdm using silicon microring modulators," *Opt. Express*, vol. 18, no. 16, p. 16858, Aug. 2010.
[9] A. V. Krishnamoorthy *et al.*, "Exploiting cmos manufacturing to reduce tuning requirements for resonant optical devices," *IEEE Photonics J.*, vol. 3, no. 3, pp. 567–579, 2011.
[10] R. Wu *et al.*, "Variation-aware adaptive tuning for nanophotonic interconnects," in *2015 IEEE/ACM Int. Conf. Comput. Des.* IEEE, Nov. 2015, pp. 487–493.
[11] P. Dong *et al.*, "Low power and compact reconfigurable multiplexing devices based on silicon microring resonators," *Opt. Express*, vol. 18, no. 10, p. 9852, 2010.
[12] R. Wu *et al.*, "Compact models for carrier-injection silicon microring modulators," *Opt. Express*, vol. 23, no. 12, p. 15545, Jun. 2015.
[13] B. Szelag *et al.*, "Integration and modeling of photonic devices suitable for high performance computing and data center applications," in *Proc. SPIE*, vol. 10108, Feb. 2017, p. 19.
[14] Y. Liu *et al.*, "Ultra-compact 320 gb / s and 160 gb / s wdm transmitters based on silicon microrings," *Opt. Fiber Commun. Conf.*, pp. 6–8, 2014.
[15] X. Zheng *et al.*, "Ultralow power 80 gb/s arrayed cmos silicon photonic transceivers for wdm optical links," *J. Light. Technol.*, vol. 30, no. 4, pp. 641–650, 2012.
[16] H. Jayatilleka *et al.*, "Wavelength tuning and stabilization of microring-based filters using silicon in-resonator photoconductive heaters," *Opt. Express*, vol. 23, no. 19, p. 25084, 2015.
[17] J. F. Buckwalter *et al.*, "A monolithic 25-gb/s transceiver with photonic ring modulators and ge detectors in a 130-nm cmos soi process," *IEEE J. Solid-State Circuits*, vol. 47, no. 6, pp. 1309–1322, 2012.
[18] C. Sun *et al.*, "A 45 nm cmos-soi monolithic photonics platform with bit-statistics-based resonant microring thermal tuning," *IEEE J. Solid-State Circuits*, vol. 51, no. 4, pp. 893–907, 2016.
[19] R. Polster *et al.*, "Efficiency optimization of silicon photonic links in 65-nm cmos and 28-nm fdsoi technology nodes," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 24, no. 12, pp. 3450–3459, 2016.
[20] M. Georgas *et al.*, "Addressing link-level design tradeoffs for integrated photonic interconnects," in *2011 IEEE Cust. Integr. Circuits Conf.* IEEE, Sep. 2011, pp. 1–8.
[21] Y. Zheng *et al.*, "Post-fabrication reconfiguration for power-optimized tuning of optically connected multi-core systems," in *17th Asia South Pacific Des. Autom. Conf.* IEEE, Jan. 2012, pp. 615–620.
[22] Y. Zheng *et al.*, "Power-efficient calibration and reconfiguration for optical network-on-chip," *J. Opt. Commun. Netw.*, vol. 4, no. 12, p. 955, Dec. 2012.
[23] Y. Wang *et al.*, "Energy-efficient channel alignment of dwdm silicon photonic transceivers," *Des. Autom. Test Eur. Conf. Exhib. (DATE), 2018*, Mar. 2018.
[24] R. Wu *et al.*, "Pairing of microring-based silicon photonic transceivers for tuning power optimization," in *2018 23rd Asia South Pacific Des. Autom. Conf.* IEEE, Jan. 2018, pp. 135–140.
[25] W. Zhang *et al.*, "Virtual probe : A statistical framework for low-cost silicon characterization of nanoscale integrated circuits," *IEEE Trans. Comput. Des. Integr. Circuits Syst.*, vol. 30, no. 12, pp. 1814–1827, 2011.
[26] S. Zhang *et al.*, "Joint virtual probe: Joint exploration of multiple test items' spatial patterns for efficient silicon characterization and test prediction," in *Des. Autom. Test Eur. Conf. Exhib. (DATE), 2014*, Mar. 2014, pp. 1–6.
[27] M. J. Heck and J. E. Bowers, "Energy efficient and energy proportional optical interconnects for multi-core processors: Driving the need for on-chip sources," *IEEE J. Sel. Top. Quantum Electron.*, vol. 20, no. 4, 2014.
[28] M. Bahadori *et al.*, "Energy-performance optimized design of silicon photonic interconnection networks for high-performance computing," in *Des. Autom. Test Eur. Conf. Exhib. (DATE), 2017.* IEEE, Mar. 2017, pp. 326–331.
[29] C.-H. Chen *et al.*, "A comb laser-driven dwdm silicon photonic transmitter based on microring modulators," *Opt. Express*, vol. 23, no. 16, p. 21541, 2015.
[30] K. Bergman, "Nanophotonic interconnection networks for performance-energy optimized computing," in *IEEE International Interconnect Technology Conference*, 2012.
[31] Q. Li *et al.*, "Experimental characterization of the optical-power upper bound in a silicon microring modulator," *2012 Opt. Interconnects Conf. OIC 2012*, vol. 5, pp. 38–39, 2012.
[32] K. Yu *et al.*, "A 25 gb/s hybrid-integrated silicon photonic source-synchronous receiver with microring wavelength stabilization," *IEEE J. Solid-State Circuits*, vol. 51, no. 9, pp. 2129–2141, Sep. 2016.
[33] R. M. Karp, *Reducibility among Combinatorial Problems.* Boston, MA: Springer US, 1972, pp. 85–103.