Fabrication-Tolerant Four-Channel Wavelength-Division-Multiplexing Filter Based on Collectively Tuned Si Microrings

Peter De Heyn, *Student Member, IEEE*, Jeroen De Coster, *Member, IEEE*, Peter Verheyen, Guy Lepage, Marianna Pantouvaki, *Member, IEEE*, Philippe Absil, Wim Bogaerts, *Member, IEEE*, Joris Van Campenhout, *Member, IEEE*, and Dries Van Thourhout, *Member, IEEE*

Abstract—We demonstrate a robust, compact and low-loss four-channel wavelength-division multiplexing (WDM) filter based on cascaded double-ring resonators (2RR) in silicon. The flat-top channel response obtained by the second-order filter design is exploited to compensate for the detrimental effects of local fabrication variations and their associated phase errors on the ring-based filter response. Full wafer-scale characterization of a cascaded, four-channel 2RR filter with channel spacing of 300 GHz shows an average worst-case insertion loss below 1.5 dB and an average worst-case crosstalk below −18 dB across the wafer, representing a substantial improvement over a first-order based ring (1RR) design. The robust 2RR filter design enables the use of a simple collective thermal tuning mechanism to compensate for global fabrication variations as well as for global temperature fluctuations of the WDM filter, the WDM laser source, or both. Highly uniform collective heating is demonstrated using integrated doped silicon heaters. The compact filter footprint of less than $50 \times 50 \ \mu m^2$ per channel enables straightforward scaling of the WDM channel count to 8 channels and beyond. Such low-loss collectively tuned ring-based WDM filters can prove beneficial in scaling the bandwidth density of chip-level silicon optical interconnects.

Index Terms—Design for manufacturing, microring resonators, silicon-on-Insulator (SOI), wavelength-division multiplexing.

I. INTRODUCTION

AGGREGATE bandwidth requirements for I/O communication in advanced CMOS nodes are expected to reach the level of TB/s in the near future [1], [2] based on projections of the International Technology Roadmap for Semiconductors (ITRS). Optical interconnects based on silicon photonics are increasingly being considered as a viable alternative to enable further I/O scaling [3]. An attractive approach to scale the bandwidth of silicon-based optical interconnects is the adoption of wavelength-division multiplexing (WDM). To meet the stringent targets for the power efficiency of chip-level optical interconnects ($\approx$1 pJ/bit), the WDM filters need to (1) have a low insertion loss (IL) and low crosstalk (XT), (2) have a compact footprint, (3) be manufacturable, (4) be thermally robust, either be design or by simple and low-power thermal tuning.

Filters realized with microring resonators (MRR) using a high-index contrast platform such as silicon-on-insulator (SOI) can be made very compact, and can be designed with large free-spectral ranges (FSR). Ring-based WDM components have been demonstrated with a considerably smaller footprint [4] compared to other filter implementations such as arrayed waveguide gratings (AWG) or echelle gratings. Besides enabling high integration density, the compactness of ring-based filters enables low-power thermal tuning. Higher-order, multi-ring filters enable a higher design freedom compared to first-order single-ring designs, allowing filters to be designed with wider channel bandwidths and steeper channel roll off characteristics [5]. The resulting flat-top channel response can be exploited to improve the crosstalk and insertion loss of the WDM filter.

One of the major issues for realizing compact WDM optical devices on SOI is the sensitivity of these components to variations in waveguide dimensions. These include linewidth variations during lithographic pattern definition as well as thickness variations of the top silicon layer of the SOI stack [6]. Linewidth variations can be present both at the device scale (1–100 μm), chip scale (1–20 mm) or wafer scale (up to 200 mm). Local variations at the local device scale are typically much smaller ($\approx$1 nm) than the global variations found across a full wafer ($\approx$10 nm and higher). Silicon thickness variations typically occur at the wafer scale, with variations of 10 nm and more across a full wafer [6]. In addition, silicon wavelength-selective devices are known to be highly sensitive to temperature variations. These perturbations directly influence the effective refractive index and give rise to a detrimental shift of the resonance wavelength of WDM filters such as microring resonator based filters [7], either on device, chip or wafer scale [6].

To make waveguides more robust to dimensional variations, one can optimize the waveguide dimensions [6] or use the less confined TM-polarization [7]. The drawback of using the TM-polarization is the larger bending radius, which limits the FSR owing to the lower confinement in the silicon core.

In this paper, we exploit a flat-top response of a cascaded second-order double-ring filter to compensate for the local fab-

rication variations present within the filter footprint. A channel spacing of 300 GHz is chosen and found to offer a good compromise between channel density and filter robustness, resulting in an average worst-case channel insertion loss below 1.5 dB and an average worst-case channel crosstalk below -18 dB across the wafer. As will be shown, this represents an improvement of close to 5 dB on both metrics compared to a design based on a cascaded first-order single-ring filter. The robust filter response allows to use a simple collective tuning mechanism to compensate for the global fabrication and temperature variations. Uniform local heating of the WDM filter is demonstrated using carefully designed p-type doped silicon heaters. Collective tuning would require only a single thermal control circuit for each WDM filter, as compared to an individually tuned WDM filter array on a 200 GHz grid as proposed in [8]. This reduced complexity for thermal control will likely reduce the power and footprint overhead of the required CMOS control circuits, which will be beneficial when scaling to higher channel counts. The idea of the paper has been demonstrated in [9] and is in this paper elaborated using wafer plots and device-count histograms providing the reader more insight.

The paper is organized as follows. We start the paper by explaining in Section II the microring filter designs used in this work, followed by the fabrication details in Section III. Then, Section IV is devoted to the methodology used to characterize the filters. The actual experimental results are presented in Section V where a comparison is made between a receiver using first-order and second-order filters regarding robustness towards resonance shifts due to waveguide dimensional variations while optimized using a collective heater. In Section VI the conclusion can be found.

II. DESIGN

The WDM filter consists of four cascaded double-ring resonators (2RR). The channel spacing is designed to fit a grid spacing of 300 GHz (2.4 nm at 1550 nm) and the free spectral range (FSR) is 12 nm. For comparison, a similar four-channel WDM filter based on cascaded single-ring resonators (1RR) is designed as well. The rings have a racetrack shape and are implemented with a 9 μm coupling length and a 5 μm radius. A small increase in round-trip length of the ring is used to establish the shift in resonance wavelength according to the defined channel spacing (CS) using following formula.

$$\Delta L = \frac{n_g}{n_{eff}} \frac{CS}{\lambda_{center}} L \qquad (1)$$

with n_g and n_{eff} respectively the group and effective refractive index of the mode, λ_{center} the center wavelength and L the total round-trip length of the MRR. In this design $\Delta L = 150\,\text{nm}$ which is implemented with an added 75 nm in both racetracks of the ring. The very small increase in coupling strength can be compensated with a 1.5 nm change in gap, which is not implemented due to the 5 nm mask grid resolution. The 2RRs and 1RRs are designed to have a -20 dB roll off at one channel spacing (2.4 nm) away from their center resonance wavelength, as well as a maximized (flat-top) bandwidth at the drop ports. This results in a bus-ring waveguide power coupling of 0.08 for the 1RR and a bus-ring and ring-ring waveguide power coupling
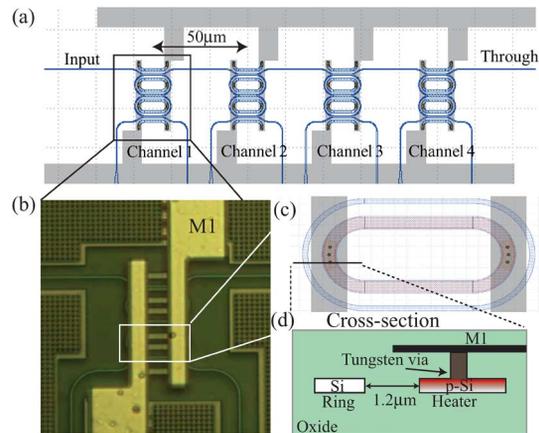

Fig. 1. A robust 4-channel WDM demultiplexing filter based on collectively tuned silicon microrings with (a) design overview, (b) microscopic picture of the filter and heater, (c) zoom on the design of 1 microring and p-doped silicon resistor and (d) cross-section of the ring waveguide and heater structure.

of respectively 0.32 and 0.04 for the 2RR. Using a 450 nm wide and 220 nm thick waveguide with oxide cladding for both the bus and ring this results in a bus-ring gap of 295 nm for 1RR and a bus-ring and a ring-ring gap of respectively 205 nm and 340 nm for 2RR.

III. FABRICATION

The WDM filter was fabricated on a 200-mm SOI wafer with 2 μm buried oxide and 220 nm top c-Si layer using a subset of processing modules from imec's Silicon-Photonics Platform (iSiPP). Two silicon patterning steps were carried out in which respectively 70 nm and 220 nm of the c-Si layer were locally etched to define fiber-grating couplers as well as the ring filter and the strip access waveguides. The heaters were implemented as highly p-type doped, 1 μm-wide silicon resistors, located at 1.2 μm away from the inner edge of the ring waveguides. Local silicide and a CMOS-like tungsten/copper back-end are used to contact the heaters.

A design schematic and microscopic picture of one channel of a 4-channel WDM filter using 2RRs with a collective heater is given in Fig. 1.

IV. METHODOLOGY

Both receiver designs are characterized using a tunable laser with a wavelength step size of 10 pm, centered around 1550 nm. The spectral responses are normalized by subtracting the Gaussian-like spectrum of the fiber grating couplers which are extracted from the fitting of the off-resonance data of the through port response. Measurements are performed using an automated alignment procedure assuring a reproducible fiber-to-chip coupling in a temperature-controlled environment. Because the coupling of the microring resonators is wavelength dependent, filter characteristics are all measured preferably around the same wavelength. Due to fabrication variations, the center wavelength of the filter bank is varying over the wafer with more than one FSR. Therefor the analysis is done on the four adjacent resonances such that on average the first channel is situated around 1540 nm.
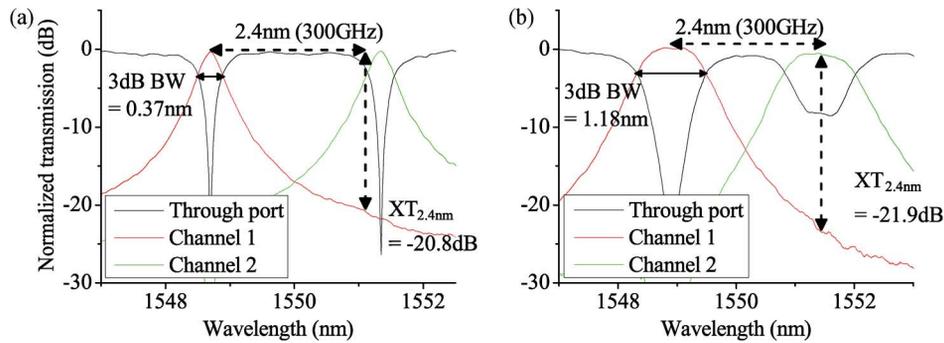
Fig. 2. A typical channel response of (a) a 1RR and (b) 2RR design, employed in a cascaded configuration. Both channel responses have a small IL and meet the XT requirement at 2.4 nm ($XT_{2.4 nm}$) of at least -20 dB.

V. EXPERIMENTAL RESULTS

To fully understand the performance of the receiver, we start in Section V-A by investigating the individual channel response of both filter designs (1RR and 2RR). Wafer-scale trends of the spectral shape, center wavelength position and the relation to their position on the wafer are examined. As discussed in the introduction, the deviation of the absolute position of the filter bank will be corrected with a collective heater, which is characterized in Section V-B regarding efficiency, uniformity and total power consumption required to tune over a full FSR. However, another important aspect is the position of the center wavelength of each channel within a device, with respect to the adjacent channels. This channel spacing analysis (Section V-C) will learn us how tolerant each filter design is against device-scale fabrication related non-uniformities. From this it will be clear that the main problem of designing a filter based on microring resonators on SOI is not the spectral shape but the variation of the center wavelength of the filter with respect to the other adjacent filters. This channel spacing variation will have direct consequences on the performance of the receiver when locked in a fixed channel grid of 300 GHz or 2.4 nm and a profound comparison is made between both filter designs. The consequences on the insertion loss and crosstalk of the receiver's worst channel are discussed in Section V-D, defining the robustness of our design against fabrication related non-uniformities.

A. Individual Channel Response

In Fig. 2(a) and (b), one can see a typical spectral response of the first channel of a 4-channel WDM filter implemented with respectively first and second-order MRRs. The 3 dB bandwidth (BW) of design 2 is indeed much larger ($3\ dB\ BW = 1.18\ nm$) than design 1 ($3\ dB\ BW = 0.37\ nm$) whereas the crosstalk (XT) is for both designs better than -20 dB. The possible but small misalignment between the resonances of the two microring resonators forming a 2RR (design 2) can cause some fluctuation in insertion loss (IL) between the different channels. This is visible in Fig. 2(b) where one can notice a difference in extinction ratio of the through port at the resonance of channel 1 and channel 2, which explains the slightly larger IL of channel 2 with respect to channel 1. We investigated nominally identical devices on different dies across the wafer and the statistics (mean, 10th and 90th percentiles) based on the mean value of

TABLE I
LIST OF FILTER SPECIFICATIONS. MEAN AND BOTH PERCENTILES (10TH, 90TH) ARE BASED ON THE MEAN VALUE OF THE 4 CHANNELS

Specification	first order			second order		
	mean	P10	P90	mean	P10	P90
1dB BW	0.22	0.2	0.24	0.87	0.8	0.94
3dB BW	0.43	0.38	0.47	1.27	1.17	1.37
XT at 2.4nm	-20.4	-21.3	-19.5	-21.4	-22.8	-20.3
IL	-0.45	-0.63	-0.25	-0.38	-0.76	-0.01

the 4 channels (device mean) are listed in Table I. One can see that the device using 2RR has slightly better mean values but a higher spreading regarding XT and IL. In Fig. 3, we plot the insertion loss (IL) of only the receiver's worst channel (largest IL) instead of the mean value listed in Table I. From this two wafer plots it is clear that this worst channel is depending on local non-uniformities because there is not a clear wafer-scale trend visible. From the device-count histogram one can see that the receiver's worst channel using the 2RR has the tendency to have a larger spreading and more skewed distribution of the IL than when 1RRs are used.

In the following we investigate the absolute position of the center wavelength of the different channels, where we use the first channel as a reference for the whole filter comb. In Fig. 4(a), the center wavelength of this first channel in function of its position on the wafer and a corresponding device-count histogram are plotted. As explained in the Section IV, the analyzed resonance wavelength is chosen within 1 FSR around the arbitrary chosen wavelength 1540 nm. It can be seen from this figure that the center wavelength is uniformly distributed between the arbitrary chosen selection borders which means that the variation between two positions of the wafer can differ with more than 1 FSR. To create more insight in these wafer-scale variations, the same data is plotted with a different selection procedure in Fig. 4. This selection procedure is based on the assumption that the difference in center wavelength between two identical designed devices on adjacent dies must be smaller than half a FSR which results in a continuous variation of the center wavelength in function of its position on the wafer. From the corresponding histogram one can notice that the maximum variation is reaching 30 nm which is more than twice the FSR. Expelling the outliers with the shortest wavelength, which are corresponding with devices on the edge of the wafer, one finds a
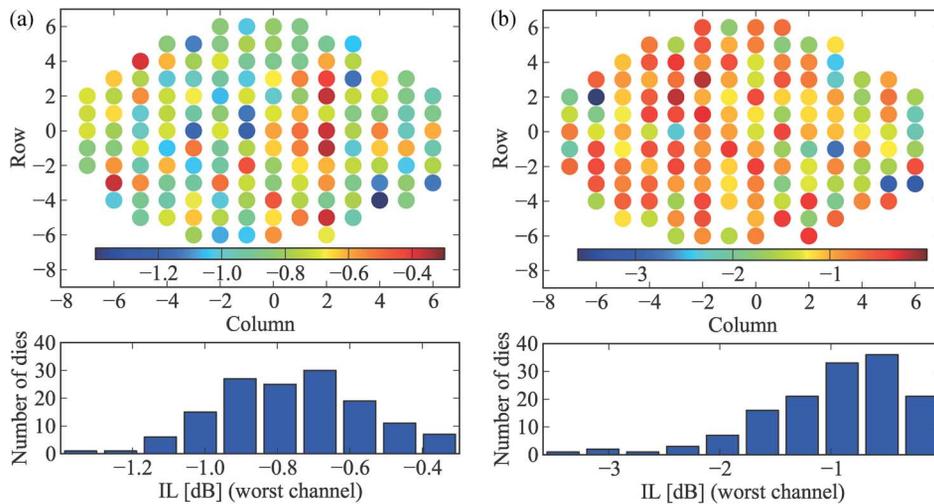
Fig. 3. Insertion loss (IL) of the receivers worst channel employed with 1RR (a) and with 2RRs (b), both in function of the receivers position on the wafer and in a device-count histogram.
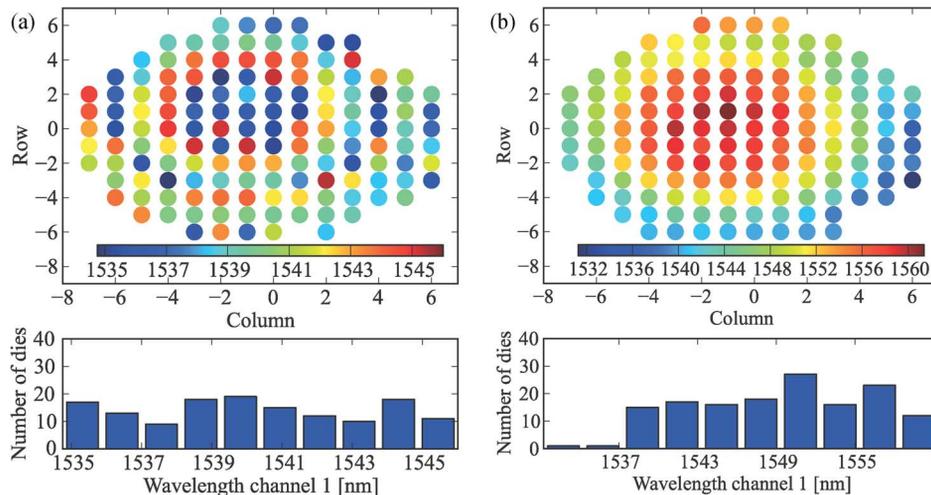

Fig. 4. The center wavelength of the first channel in function of its position on the wafer and a corresponding device-count histogram created with two selection procedures: (a) within 1 FSR around 1540 nm (b) assuming only small die-to-die shifts to create a more continuous variation.

reduced variation of 23 nm. This wafer-scale variation can partially be explained by small thickness variations of the waveguide layer within a SOI wafer where approximately 1 nm difference in thickness results in a 2 nm wavelength shift. In the following section the collective heater is characterized which will be used to compensate for these wafer-scale variations.

B. Thermal Tuning

A large thermal tuning range and low corresponding energy consumption are important for the realization of a tunable WDM receiver. These metrics are calculated by tracking the resonance red-shifts while applying a voltage to the heater. This is shown in Fig. 5(a) and (b) where this red-shift for different applied heater powers is plotted for designs using respectively 1RR and 2RR. For the sake of clarity only the first channel is plotted but the denoted power consumption and heater efficiency is calculated for the whole filter bank. The collective tuning efficiency is 0.0374 nm/mW for the design with first-order filters which is 37% more efficient then when implemented with second-order

filters (0.0273 nm/mW). This difference is explained by the larger heated area in the latter case. A necessary condition for the concept of collective tuning is a good uniformity of the heaters such that the spacing between the channels is maintained while tuning their resonances collectively. In Fig. 6 the relative shift expressed in function of power consumption is plotted for all 4 channels. From this figure it can be seen that each channel has an equal efficiency and thus a very good device-scale uniformity is achieved. Efficiency measurements were taken on a full wafer. The mean shift of all 4 channels is used to plot a device-count histogram for the power consumption needed for a full FSR shift and the efficiency itself, respectively plotted in Fig. 7(a) and (b). From this plot it is clear that there is also a very good uniformity on wafer level. Wafer-mean heater efficiency is for filter design using 1RR 0.028 nm/mW and for 2RR 0.038 nm/mW, both with a very small spread, indicating a well defined resistor used as the heater. Based on the wafer-mean values, it takes on average 296 mW for 1RR and 407 mW for 2RR to tune the 4-channel filter bank over a full FSR (calculated for each
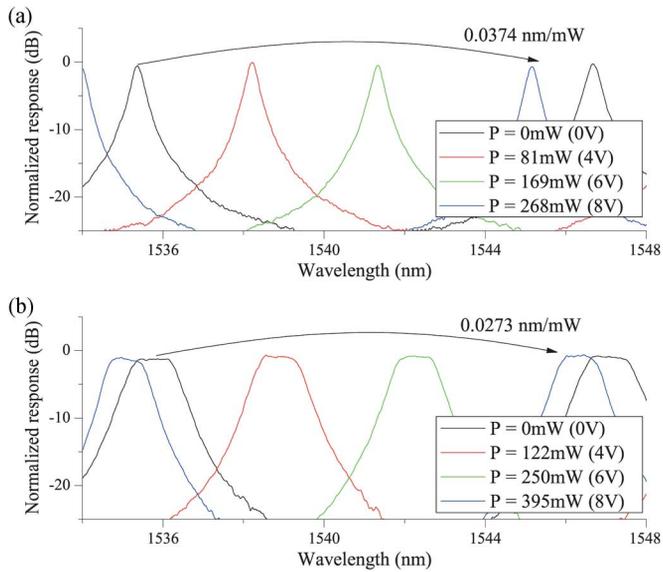
Fig. 5. Thermal tuning of the spectral response of one channel of a 4-channel WDM filter using (a) first order and (b) second-order MRRs. Tuning efficiencies indicated in the figure are for the whole filter bank.
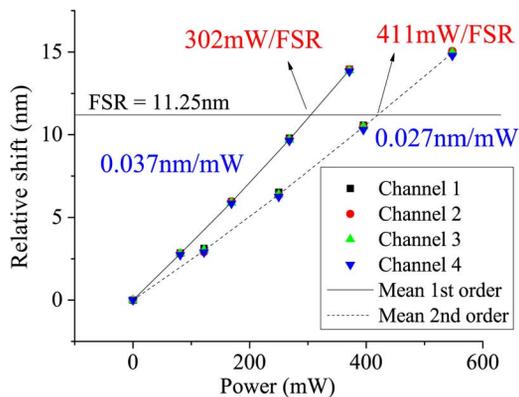

Fig. 6. The relative tuning of a 4-channel filter bank using first or second-order MRRs in function of the applied power to all four channels collectively, consuming respectively 302 mW and 411 mW to tune a full free spectral range (π-phase shift for one filter is respectively 38 mW and 51 mW).

device separately but on average 11.35 nm). For one channel this corresponds to 74 mW/FSR for the filter design with 1RR and 101 mW/FSR for 2RR, which are comparable with other reported values, e.g. 89 mW/FSR in [8]. As demonstrated in [10], these efficiencies can drastically be improved using a top-side silicon undercut-etching technique, resulting in a 2.4 mW/FSR.

To gain insight in how much power one needs to lock a receiver grid to an incoming laser grid, we can use the knowledge of the previous section where the actual position of the filter bank around a certain wavelength is uniformly distributed. Independent of the actual operation temperature, one can expect (in a worst case scenario) the need to tune the receiver filter grid up to a full FSR to lock it to an incoming laser grid. If each channel of our receiver will handle a bitrate of 20 Gb/s (following the reasoning of [2]), this would mean a worst-case power consumption of 3.7 pJ/bit for design 1 and a 5.07 pJ/bit for design 2. Note that this is a worst-case scenario, on average one device will consume half of this power (corresponding to only half a
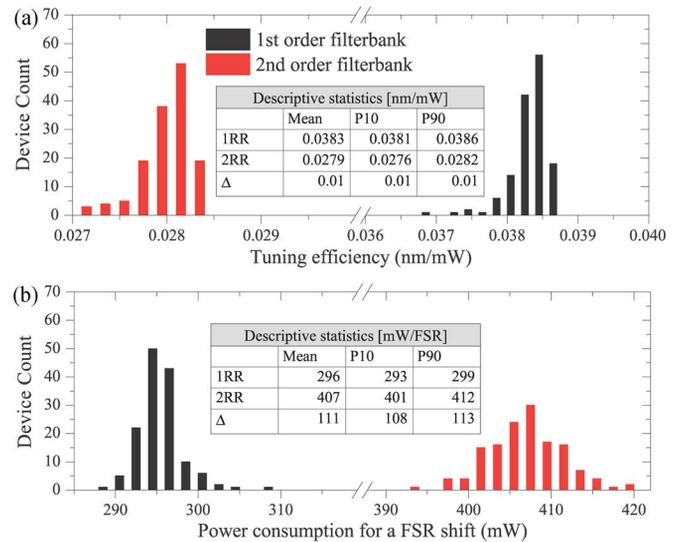

Fig. 7. Device-count histograms of (a) the efficiency and (b) the power consumption for a FSR shift of all the channels. A very good wafer-level uniformity is perceived, indicating a well defined resistor using highly p-doped silicon strip next to the optical waveguide (see Fig. 1).

FSR as the average tuning shift). This is far from the 2015 energy target for tunable WDM filters of 30 fJ/bit [2] or reported values of 15 fJ/bit [11] using under-etched waveguides and flexible wavelength registration. Another way to improve the power consumption could come from (1) fabrication, where a better control over the silicon waveguide thickness could lower the maximum tuning range required [12], (2) from design, by using designs with larger FSR and thus increasing the wavelength shift for given power consumption or (3) by an increased bitrate.

C. Channel Spacing Analysis

An important aspect of a WDM filter is a predictable and constant spacing between the center wavelength of adjacent channels. Especially when using a collective heater where these inter-channel deviations cannot be compensated, this deviation must be smaller than e.g. half the 3 dB bandwidth to keep the insertion loss tolerable. These random device-scale variations between resonances of the 50 μm-spaced filters range typically up to 1 nm. To quantify this channel spacing variability, we perform a linear regression analysis by fitting the absolute resonance position of the 4 channels (plotted in Fig. 8(a)) as function of their channel number. The slope of this fit is then the device-optimized channel spacing (CS) which is plotted in Fig. 8(b). In Fig. 8(c) the residuals representing the relative deviations away from the CS are plotted, which are in this example ≤0.2 nm. We repeat this analysis with a linear fitting with a fixed slope equal to the designed value of the CS (i.e. 2.4 nm). Using this fitting we expect larger deviations but this will better match the performance in a practical situation where the CS should equal the incoming channel grid, defined in the design phase of the component. In this way the largest deviation goes up to 0.4 nm. We can wrap this 4 deviations into a single number by taking the standard deviation (σ) and repeat both fitting analyses for all 1RR devices on a full wafer. This results in two device-count histograms for this standard deviation of the resonance deviation, calculated with both fitting analyses,
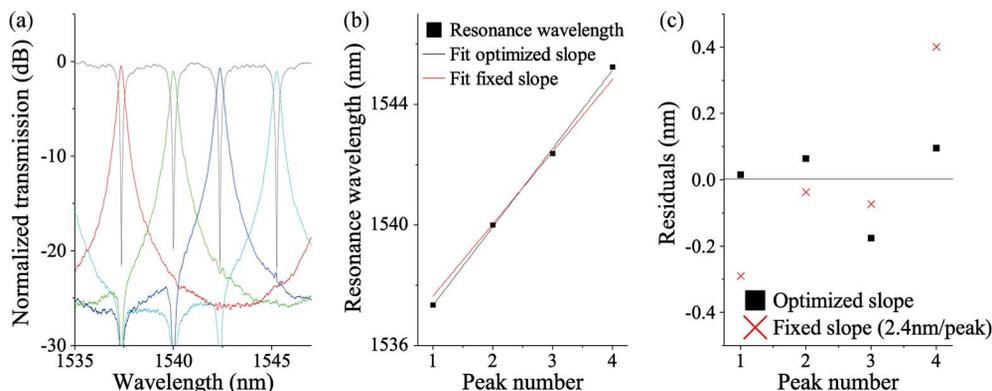
Fig. 8. A typical channel spacing analysis with (a) the normalized response of a 4-channel WDM filter (first order microring resonators). In (b) the absolute position of the resonances is plotted in function of the peak number. After fitting, the residuals are calculated and plotted in (c). Two fits are made, one where the slope is optimized for this particular device, and one with a fixed slope for all devices measured on the wafer.
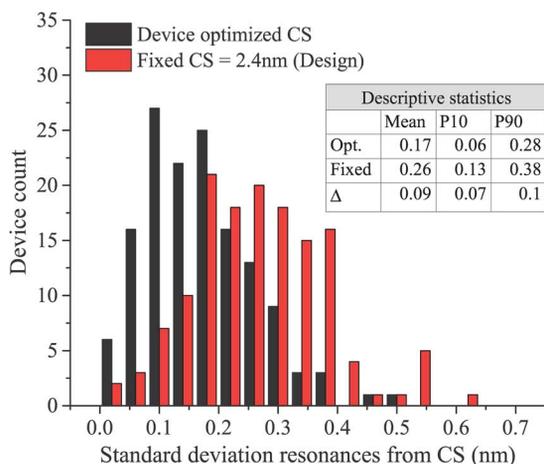

Fig. 9. The device count histogram of a full-wafer measurement describing the standard deviation of the residuals after fitting to a linear curve as described in Fig. 8. The red curve is calculated with a fixed slope, equal to the designed channel spacing (CS) of 2.4 nm, which gives large deviations when calculated with a device-optimized CS (black curve). Also basic statistics (mean, 10th and 90th percentiles) are given.

and plotted in Fig. 9. As expected the device-optimized CS gives smaller deviations with a wafer-mean value of 0.17 nm, compared to the 0.26 nm calculated with the more realistic method of a fixed as-designed CS. However, we believe that this deviation can partly be compensated by design, e.g. a better calibration of the group and effective refractive index used in the calculation of the incremental circumference of the MRR (see (1)). In this case we have found that the wafer-mean group index is 2.56, indicating a detrimental shift which could be compensated by design. In the following we investigate how robust our WDM filter designs are to these CS deviations.

D. Crosstalk and Insertion Loss

To test how tolerant both receivers are to lock into a fixed channel grid of 300 GHz or 2.4 nm, we numerically shift the measured responses of all 4 channel resonances with an appropriate multiple of the designed channel spacing on top of each other. In this way, the impact of device-scale variations

of the relative resonance frequencies away from their ideal spectral position are easily detected. The result is plotted in Fig. 10(a) and (b), respectively for a 1RR and 2RR design. It can be seen that for a 1RR design, the channel spacing deviation can be much larger than the 3 dB BW. For a 2RR design this is much better. Next, the collective tuning of the resonances of the MRRs is emulated by choosing the best position of the receiver grid which is determined by minimization of the IL of the worst channel. This procedure is also illustrated in Fig. 10 where the dashed line is the moving minimum response of all four drop channels. The maximum of this curve overlaps by definition with the minimum IL of the worst channel, which is determining the overall performance of the receiver. The resulting IL and XT of all the channels can then be calculated and is denoted in Fig. 11, where also the laser grid with a 300 GHz grid spacing is plotted. In this example, both receivers channel 1 and 2 have the largest IL and are thus limiting the performance of the receiver. We compare both receivers based on their worst channel resulting in an IL of -5.7 dB and XT of -12.6 dB for the 1RR design which makes the 2RR design superior against device-scale resonance variations with an IL of only -1.47 dB and XT of -17.3 dB. To validate our findings, we measured the same two designs on all 146 dies of a 200 mm-wafer and compared both receivers on their degraded IL and XT when collectively locked on a fixed laser grid with a 300 GHz spacing. In Fig. 12, a histogram of the IL and XT is shown for the worst channel of each design, together with the mean, 10th and 90th percentiles. From this figure it is clear that the 2RR design is showing superior characteristics regarding IL and XT, not only in the average value but also in spreading of these metrics (P10 and P90), thanks to its flat-top filter behavior. As shown in the inset tables of Fig. 12, the wafer-mean value of the IL is -6.09 dB for 1RR design, which is 4.7 dB larger than the 2RR design, which has a wafer-mean IL of only -1.39 dB. The wafer-mean XT is -13.5 dB for the 1RR design, which is 4.7 dB larger than the 2RR design, which has a wafer-mean XT of -18.2 dB.

From this tables we conclude that there is a difference in both insertion loss and crosstalk of 4.7 dB by using the 2RR design instead of the 1RR design.
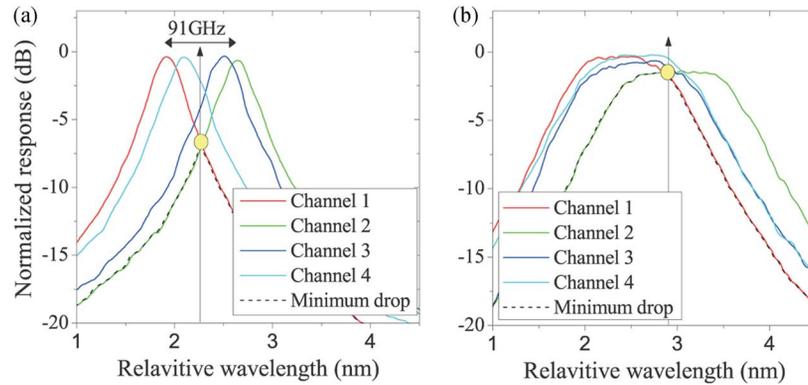
Fig. 10. Numerically shifted responses of all channel resonances, each shifted with a multiple of the designed channel spacing on top of each other for (a) 1RRs and (b) 2RRs. In this way one easily detects deviations of the designed channel spacing. Based on this curve the optimum position of the receiver grid is easily derived based on the maximum of the moving minimum of all four channels (dashed line).
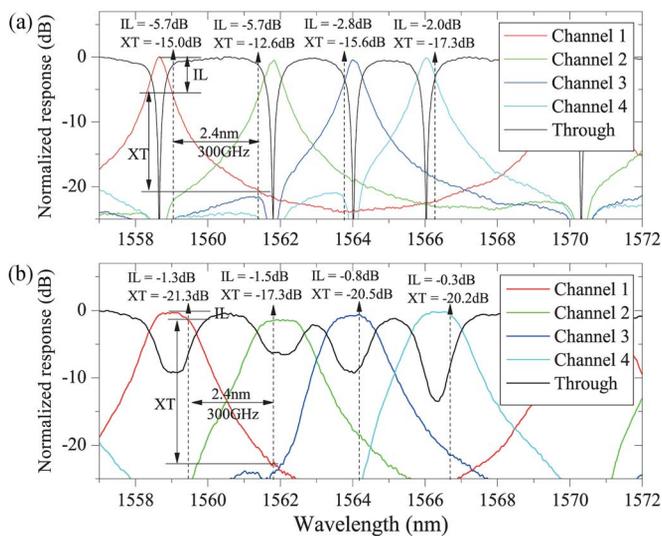

Fig. 11. The typical normalized spectral response of a 4-channel WDM filter implemented with (a) 1RRs and (b) 2RRs. The channel responses are collectively tuned to an optimum position where the insertion loss (IL) of the worst channel is minimized. The resulting IL and crosstalk (XT) per channel when collectively tuned and operated with four 300 GHz-spaced lasers (dashed arrows) is denoted in the figure as well.
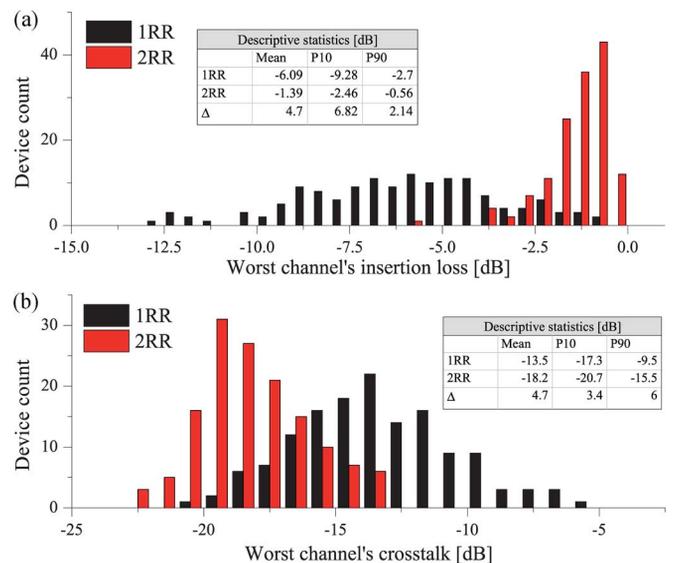

Fig. 12. Device-count histogram of (a) insertion loss and (b) crosstalk of the worst channel of the 4-channel receiver using 1RRs and 2RRs when the filter is collectively tuned. The 2RRs are must more robust against device level non-uniformities influencing the CS between the channels. In the inset table the mean, 10th and 90th percentiles are given as well.

VI. CONCLUSION

We demonstrate a robust 4-channel WDM demultiplexing filter based on cascaded silicon microring resonators. The flat-top response of second-order multi-ring filters is exploited to overcome the wafer- and device-scale irregularities on channel spacing and channel bandwidth arising from fabrication non-uniformity. The improved filter response enables a thermal control mechanism based on collective tuning to track the wavelengths of a WDM laser source on a specified channel grid. After collective tuning with integrated heaters, the 4-channel second-order filter exhibits a worst-channel insertion loss of -1.39 dB and a crosstalk of -18.2 dB (wafer-level average), an improvement of 4.7 dB on both metrics over first-order filters. The required power consumption of such a collectively tuned 4-channel WDM demultiplexing filter using with 2RRs can reach up to 406 mW/FSR (wafer-level average) to tune over a full FSR. This relatively low efficiency is expected to increase by an order of magnitude, when proven techniques on removing of the substrate are applied (e.g. [10]).

ACKNOWLEDGMENT

REFERENCES

[1] D. A. B. Miller, "Device requirements for optical interconnects to silicon chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1185, 2009.

[2] A. Krishnamoorthy, H. Schwetman, P. Koka, I. Shubin, and J. Cunningham, "Computer systems based on silicon photonic interconnects," *Proc. IEEE*, vol. 97, no. 7, pp. 1337–1361, Jul. 2009.

[3] D. Van Thourhout, T. Spuesens, S. Selvaraja, L. Liu, R. Kumar, G. Morthier, P. Rojo-romeo, F. Mandorlo, P. Regreny, O. Raz, C. Kopp, and L. Grenouillet, "Nanophotonic devices for optical interconnect," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 5, pp. 1363–1375, 2010.

[4] W. Bogaerts, S. K. Selvaraja, P. Dumon, J. Brouckaert, K. De Vos, D. Van Thourhout, and R. Baets, "Silicon-on-insulator spectral filters fabricated with CMOS technology," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 33–44, 2010.

[5] B. Little, S. Chu, P. Absil, J. Hryniewicz, F. Johnson, F. Seiferth, D. Gill, V. Van, O. King, and M. Trakalo, "Very high-order microring resonator filters for WDM applications," *IEEE Photon. Technol. Lett.*, vol. 16, no. 10, pp. 2263–2265, Oct. 2004.

[6] S. K. Selvaraja, W. Bogaerts, P. Dumon, D. V. Thourhout, and R. Baets, "Subnanometer linewidth uniformity in silicon nanophotonic waveguide devices using CMOS fabrication technology," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 316–324, 2010.

[7] S. Park, K.-J. Kim, I.-G. Kim, and G. Kim, "Si micro-ring MUX/DeMUX WDM filters," *Opt. Exp.*, vol. 19, no. 14, pp. 14 627–14 633, 2011.

[8] X. Zheng, I. Shubin, G. Li, T. Pinguet, A. Mekis, J. Yao, H. Thacker, Y. Luo, J. Costa, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "A tunable 1 × 4 silicon CMOS photonic wavelength multiplexer/demultiplexer for dense optical interconnects," *Opt. Exp.*, vol. 18, no. 5, pp. 5151–5160, Mar. 2010.

[9] P. De Heyn, M. Vanslembrouck, D. Van Thourhout, J. Van Campenhout, J. De Coster, M. Pantouvaki, and P. Absil, "Robust 4-channel WDM demultiplexing filter based on collectively tuned silicon microrings," in *Proc. 2012 IEEE 9th Int. Conf. Group IV Photonics (GFP)*, 2012, pp. 75–77.

[10] P. Dong, W. Qian, H. Liang, R. Shafiiha, D. Feng, G. Li, J. E. Cunningham, A. V. Krishnamoorthy, and M. Asghari, "Thermally tunable silicon racetrack resonators with ultralow tuning power," *Opt. Exp.*, vol. 18, no. 19, pp. 20 298–20 304, Sep. 2010.

[11] A. V. Krishnamoorthy, X. Zheng, G. Li, J. Yao, T. Pinguet, A. Mekis, H. Thacker, I. Shubin, Y. Luo, K. Raj, and J. E. Cunningham, "Exploiting CMOS manufacturing to reduce tuning requirements for resonant optical devices," *IEEE Photon. J.*, vol. 3, no. 3, pp. 567–579, 2011.

[12] S. Selvaraja, L. Fernandez, M. Vanslembrouck, J.-L. Everaert, P. Dumon, J. Van Campenhout, W. Bogaerts, and P. Absil, "Si photonic device uniformity improvement using wafer-scale location specific processing," in *Proc. IEEE Photon. Conf. 2012*, Sep. 2012, pp. 725–726.

Peter De Heyn received the masters degree in photonics engineering from Ghent University and Vrije Universiteit Brussel (VUB), Belgium in 2009. He is now working towards a PhD degree in the Photonics Research Group of Ghent University and IMEC, Belgium. His interest is optical interconnects on silicon photonics including microring resonators for filter applications, heterogeneous III-V on SOI integration, high-performance III-V and germanium photodiodes.

Jeroen De Coster received the master's degree in electrical power engineering from the KULeuven in Leuven, Belgium in 2001 and the PhD degree from the same university in 2006 after a study on the design, modelling and characterisation of RF-MEMS devices. Since 2006 he is with imec, Belgium, where he is developing measurement tools and procedures for reliability and functional testing of MEMS as well as silicon photonics devices.

Peter Verheyen received the degree in electrical engineering and the Ph.D. degree from Katholieke Universiteit Leuven, Leuven, Belgium, in 1996 and 2003, respectively. He was a Doctoral Researcher with the Interuniversity Microelectronics Center, Leuven, involved in advanced CMOS integration and MEMS integration. He is currently a member of the Silicon Photonics staff as a part of the 3D integration group.

Guy Lepage received a bachelor degree in chemistry from REGA Hogeschool, Leuven, Belgium in 1999. After starting as a research chemist at UCB Chemicals he is working with imec, Belgium, since 2004 where he is now assisting the clean room processes for photonic integration.

Marianna Pantouvaki (S'99–M'03) received the Ph.D. degree in Electronic and Electrical Engineering in 2004, from University College London (UCL), London, UK. From 2003–2006 she worked as a Research Fellow at the Ultra-Fast Photonics group, UCL, and she was a visiting researcher at the Centre for Integrated Photonics, Ipswich, UK. From 2006–2010 she worked on Cu/low-k advanced BEOL interconnects in imec, Leuven, Belgium. Since 2011 she joined the Optical IO program of imec, where she is currently involved in the development of high-speed transmitters for optical interconnects.

Philippe Absil Ph.D. is the optical IO program director at imec and is responsible for the silicon photonics technology platform development since 2010. Before that he spent seven years managing the advanced CMOS scaling program at imec. In the early 2000's he developed the passive photonics platform technology for Little Optics Inc., Maryland, USA. He earned his Ph.D. degree in 2000 from the Department of Electrical Engineering of the University of Maryland at College Park, USA. His doctoral work contributed to the early demonstrations of semiconductor micro-ring resonators.

Wim Bogaerts is professor in the Photonics Research Group at Ghent University. He completed his studies in engineering (applied physics) at Ghent University in 1998 and joined the department of information technology (INTEC) at both Ghent University and the Interuniversity Microelectronics Center (IMEC) where he received his Ph. D. degree in 2004. In the photonics research group he specialized in the modeling, design and fabrication of nanophotonic components. Currently he coordinates the development of nanophotonic components in SOI in IMEC.

Joris Van Campenhout is Program Manager of IMEC's industry-affiliation program on Optical I/O, which explores silicon-based optical interconnects enabling low-power, high-bandwidth chip I/O in future CMOS nodes. Prior to joining IMEC, he was at IBM T.J. Watson Research Center, New York, USA, where he worked on CMOS-integrated silicon optical switches. In 2007, he obtained a PhD degree in Electrical Engineering from Ghent University, Belgium, for his work on hybrid-integrated III-V microdisk lasers on silicon. Joris is a member of IEEE and OSA, and serves as a technical program committee member in OFC/NFOEC, CLEO and IEEE-IPC.

Dries Van Thourhout received the degree in physical engineering and the Ph. D. degree from Ghent University, Belgium in 1995 and 2000, respectively. From 2000 to 2002, he was with Lucent Technologies, Bell Laboratories, Crawford Hill, working on InPInGaAsP monolithically integrated devices. In 2002, he joined the Department of Information Technology (INTEC), Ghent University, continuing his work on integrated optoelectronic devices. Main interests are heterogeneous integration by wafer bonding, intrachip optical interconnect and WDM-devices.