Column-Row Addressing of Thermo-Optic Phase Shifters for Controlling Large Silicon Photonic Circuits

Antonio Ribeiro, Sibert Declercq, Umar Khan , *Member, IEEE*, Mi Wang ,
Lukas Van Iseghem, and Wim Bogaerts , *Senior Member, IEEE*

Abstract—We demonstrate a time-multiplexed row-column addressing scheme to drive thermo-optic phase shifters in a silicon photonic circuit. By integrating a diode in series with the heater, we can connect $N \times M$ heaters in an matrix topology to N row and M column lines. The heaters are digitally driven with pulse-width modulation, and time-multiplexed over different channels. This makes it possible to drive the circuit without digital-to-analog converters, and using only $M + N$ wires. We demonstrate this concept with a 1×16 power splitter tree with 15 thermo-optic phase shifters that are controlled in a 3×5 matrix, connected through 8 bond pads. This technique is especially useful in silicon photonic circuits with many tuners but limited space for electrical connections.

Index Terms—Optical phase shifters, optical waveguide components, silicon on insulator technology digital control, time division multiplexing, thermooptic effects, integrated optoelectronics.

I. INTRODUCTION

PHOTONIC integrated circuits (PIC) combine a large number of optical functions on a chip, for applications ranging from optical communications and sensing applications to microwave signal processing [1], [2]. Especially Silicon Photonics, which has a high refractive index contrast, enables large-scale integration, with hundreds to thousands of elements per mm^2 [3]. It also has a potential scaling route to volume manufacturing through existing CMOS foundries [4].

As circuits get larger, they start suffering more from variability [5], [6], which results in circuit elements that are no longer spectrally aligned with one another, degrading the performance of the overall circuit. It is possible to compensate for these deviations by tuning the phase in the individual circuit

elements, but as circuits become larger the number of tuners needed also grows. Another evolution is the recent development of programmable photonic circuits. Such circuits are conceived as a generic waveguide mesh of interconnections, where the coupling between waveguides can be tuned electrically [7]–[10]. In order to function, these meshes need a large number of tuners for both phase and coupling.

While these tuners, usually implemented as thermo-optic phase shifters (i.e. heaters close to the waveguide) can be quite compact [11], [12], all heaters need also an electrical connection to their respective electronic driver circuits. Unless the electronics are monolithically integrated on the same chip [13], every heater requires at least one surface bondpad for flip-chipping or wirebonding (assuming a common ground for all heaters). In the wirebonding approach, the number of off-chip connections is limited to the edge of the chip, which does not scale as rapidly as the chip surface.

A large number of tuners also requires a large number of electrical driving circuits. The traditional approach for driving an on-chip heater is through a direct current (DC) voltage or current drive, which is controlled by a tunable power supply and a digital-to-analog converter (DAC). These are quite costly circuits to integrate by the thousands on an electronic IC. That is why a digital driving scheme through pulse-width-modulation (PWM) is becoming more popular [14], which works well as long as the heater time constant is much longer than the repetition frequency of the PWM signal.

The digital driving also provides an opportunity to perform multiplexing on silicon photonics heaters. By incorporating a diode in series with the heater, we only make the heater respond to a positive voltage. This allows us to drive two heaters with one pair of contact pads, by connecting them with opposite polarity [14], [15]. Combining this with a digital driving scheme, we can time-multiplex the control of the heaters with a periodic digital signal where we address half the cycle in forward bias, and half the cycle in reverse bias. However, this does not dramatically increase the number of heaters we can control through a fixed number of bondpads.

In this paper, we scale up the multiplexing by connecting the heaters in an $M \times N$ matrix, enabling row-column addressing. This scheme, which has been developed in display technologies [16], connects each tuning elements to an intersection of a row and a column wire. Now we can time-multiplex the heaters

in N rows with M digital PWM signals. On top of that, we need N control signals to activate the rows. The net result of this matrix addressing scheme is that we can drive $M \times N$ heaters through $M + N$ bond-pads. This number scales with the circumference of the chip, so it can allow more heaters to be controlled with a wirebonding packaging scheme.

In Section II we discuss in more detail the digital driving scheme for thermo-optic phase shifters. We then extend this in Section III to the new concept of row-column addressing. We then demonstrate the concept experimentally in Section IV with a 1×16 multicasting circuit with 15 (3×5) phase shifters, driven through only 8 ($3 + 5$) bondpads by a custom-programmed field-programmable gate array (FPGA).

II. DRIVING OF HEATERS WITH DIODES

In silicon photonic circuits, thermo-optic phase shifters based on resistive heaters are still the most commonly used component for implementing phase tuners, tunable couplers and switches. The phase shift induced along the path of the waveguide is proportional to the length of the waveguide L_{ps} and the change in effective index Δn_{eff}:

$$\Delta \phi = 2\pi \frac{L_{ps}.\Delta n_{eff}(\lambda)}{\lambda} \qquad (1)$$

In turn, Δn_{eff} is proportional to the temperature change ΔT of the waveguide induced by a heater which is located in close proximity to the waveguide. Assuming in a first approximation that the thermal properties of the materials (thermal conductance, heat capacitance, thermo-optic coefficient) are temperature independent, the heat transport equations are linear and ΔT will be proportional to the electrical power P dissipated in the heater. The induced phase shift $\Delta \phi$ is therefore proportional to the P:

$$\Delta \phi = \pi \frac{P}{P_\pi} = \pi \frac{V \cdot I}{P_\pi} = \pi \frac{I^2 R}{P_\pi} = \pi \frac{V^2}{R \cdot P_\pi}, \qquad (2)$$

With P_π the power needed to induce a π phase shift. This power is influenced to a large extent by the thermal design that governs the heat flow from the resistor (where the heat is generated) to the phase shifter (where the temperature should be increased), and the heat flow to the surroundings. A configuration where the heater and waveguide are close together, and at the same time well insulated from the surrounding devices and the substrate, can be much more efficient, and P_π can drop from 25 mW to 3 mW [11], [12]. Because the heat is contained longer in the waveguide, this also slows down the response time of the phase shifter, which can be expressed by a thermal time constant τ_{ps}. This time constant is determined by the materials, the volume that needs to be heated or cooled, and the leakage rate at which heat is conducted to the substrate and surroundings. For silicon photonic heaters, τ_{ps} can range from 10 μs to 500 μs. This makes heaters a relatively slow tuning or switching mechanism.

A. Digital Driving

Heaters can be controlled either with voltage V or current I. In both cases, the the heater must be dimensioned with the proper series resistance R, such that enough electrical power can be delivered to reach a sufficiently large $\Delta \pi$ within the limits of the voltage or current source. Because the electrical power of the heater scales quadratically with either V or I, a DC driving scheme based on a digital-to-analog converter will induce a nonlinear discretization error in the phase shift [17].

The slow response of the heater presents the opportunity for driving the heaters in a digital way, using PWM [14]. In PWM, a fixed supply voltage is used, and the delivered power to the heater is controlled by digitally modulating the voltage on and off with a duty cycle $D.C$. In this scheme, the power in the heater is also modulated, but when the modulation cycle is much faster than τ_{ps}, both ΔT and $\Delta \phi$ are averaged out, and increase linearly proportional with $D.C$. In addition, the driving circuit does not need separate voltage levels for each tuner, and no DACs, simplifying the electronics design. PWM signals can be generated by a multitude of off-the-shelf programmable digital circuits, such as digital signal processors, microcontrollers or FPGAs, and the signal current can be boosted from single electrical voltage supply with sufficient current capacity.

B. Doped Silicon Heaters With Integrated Diodes

In this work, we used heaters consisting of a doped silicon strip running parallel with the optical waveguide. This brings the heater in close proximity with the waveguide. The availability of several doping processes in the platform for the definition of fast p-n modulators [18] makes it possible to engineer the resistance of the heater. The resistance can be further configured by dividing the strip in multiple segments which are connected in parallel.

The doping also allows us to incorporate a p-n junction within the heater strip, which effectively connects a diode in series with the heater resistor [14]. As a result, the heater will only draw current when forward-biased, while in reverse bias the diode effectively blocks the response. However, the threshold voltage of the diode makes the response curve more nonlinear, and this further complicates the DC driving. The diode IV characteristic has limited effect on digital driving with PWM, as this driving scheme uses only a single fixed voltage and it is only during the transient that the nonlinear I-V curve plays a role.

We can exploit the asymmetry in the response curve to multiplex the heaters. As we previously demonstrated in [14], by connecting two diode heaters in opposite polarity, we can drive two heaters from a single pair of pads by using either negative or positive voltage. We can even time-multiplex the PWM signals to the two heaters, relying on the large time constant to keep the waveguide at a constant temperature over the entire cycle. Of course, because the heaters are driven only half of the time, we need to run them at twice the power, and their series resistances need to dimensioned accordingly. It is also important to consider the limits of the heaters, as the higher electrical peak power could damage either the diode junction or the heater itself, even through the average power is well within the limits of the heater.
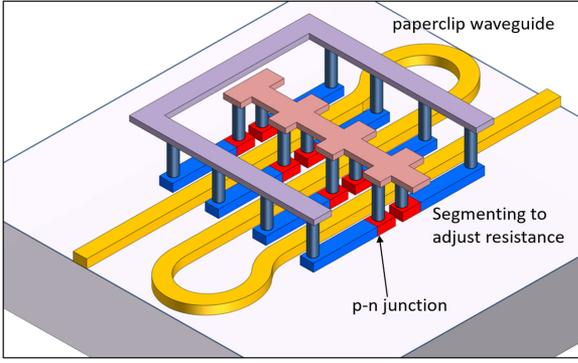
Fig. 1. Thermo-optic phase shifter using doped silicon resistors with an embedded p-n junction next to the waveguide [14]. The resistors are segmented to tune the overall resistance (and power that can be delivered for a given voltage), and the waveguide runs 3 times in between the resistors for maximum thermal efficiency.

For the results we discuss further, we have used a thermo-optic phase shifter which uses an eight-segment doped silicon stripes connected in parallel, each incorporating diode heaters. For maximum efficiency of the phase shifter, we run the same waveguide three times in between the heaters in a paperclip fashion. This results in the heater geometry sketched in Fig. 1. Without additional process modifications, such as substrate undercut and insulation trenches, this results in $P_\pi = 15\text{--}20$ mW.

III. COLUMN-ROW ADDRESSING OF DIODE-HEATER MATRICES

To extend the multiplexing concept to larger circuits and allow even more time-sharing, it is possible to electrically connect the diode-based phase shifters in a matrix topology, grouping the heaters in sets of M columns and N rows, as shown in Fig. 2(b). In this arrangement we connect the anodes of the diode-heaters that are in a same row together, while the cathodes that are in a same column are connected together. We call the rows of the matrix *control lines* and the columns *driving channels*.

It is possible to address one specific phase shifter in the matrix by setting the voltage level at its correspondent control line at a low level (Gnd), while setting the voltage of its correspondent driving channel at high level ($V+$). That puts the desired phase shifter diode in forward bias, allowing it to conduct. Because of this arrangement, the control lines acts as *enablers* in the circuit, and only a phase shifter that belongs to an enabled control line can be activated. To guarantee that only the desired phase shifter is actuated when applying a signal to its correspondent driving channel, only one control line is enabled at a time, while the remaining control lines remain disabled. Notice that for this architecture, a row is enabled at low level (Gnd), and disabled at high level ($V+$). Once we enable a row in the matrix we can drive any phase shifter in that row by setting the column to the high level ($V+$).

To allow simultaneous access to any arbitrary phase shifter in the matrix, we implemented a time-sharing actuation of the N control lines, synchronized with the driving signal applied to the M driving channels. This means that the actuation in each driving channel cycles between the N driving values for the N
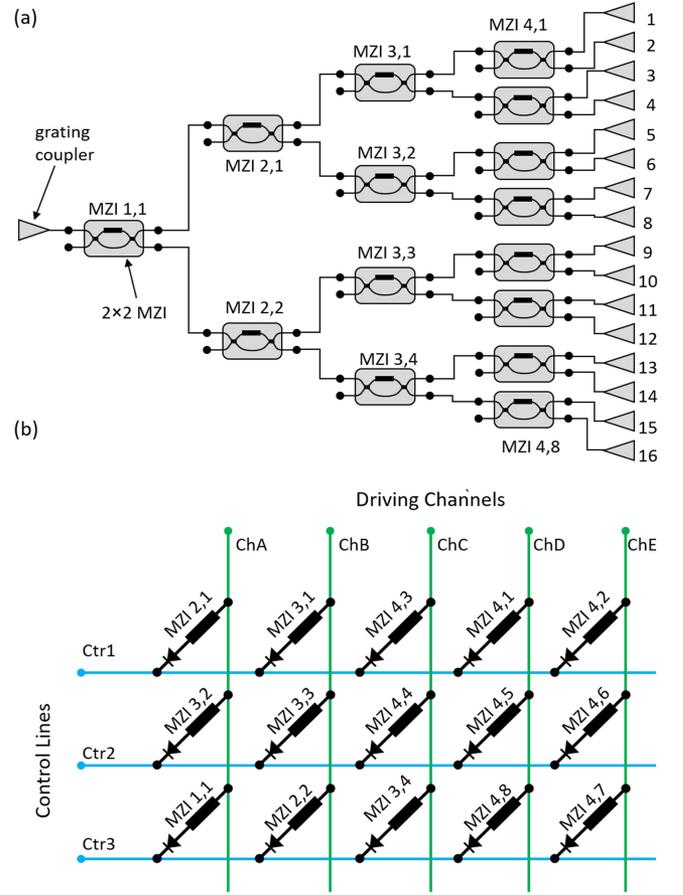
Fig. 2. (a) 1×16 splitter tree circuit consisting of 15 tunable couplers, each of which is a balanced Mach-Zehnder interferometer with a thermo-optic phase shifter in one arm. (b) Electrical connectivity of the 15 thermo-optic phase shifters. They contain a diode in series and are connected to 3 control lines and 5 driver channels.

heaters in the column. We can place M parallel channels (each with its own driving signal) on the N control lines to implement a $N \times M$ matrix-sized circuit, with N the number of control lines (rows) and M the number of channels (columns). This arrangement allows driving $N \times M$ phase shifters with only $N + M$ contact pads (and digital driving sources), as opposed to $N \times M$ contact pads and power sources needed if using traditional driving schemes.

Time-sharing means that each heater will be driven during the equivalent time of its correspondent subchannel. For a three-fold multiplexing (as illustrated in Fig. 2(b) and Fig. 3), each phase shifter will be powered, at most, during one-third of the cycle, while the remaining period it stays idle. Such a scheme only works if we rely on the large time-constant of the thermo-optical phase shifters to average out the power in the phase shifter during the driving cycle. When the refresh cycle is fast enough, the waveguide in each phase shifter will maintain a constant temperature.

A. A 1×16 Power Distribution Network

We demonstrate this concept with the 1×16 power distribution network shown in Fig. 2(a). The four binary splitting
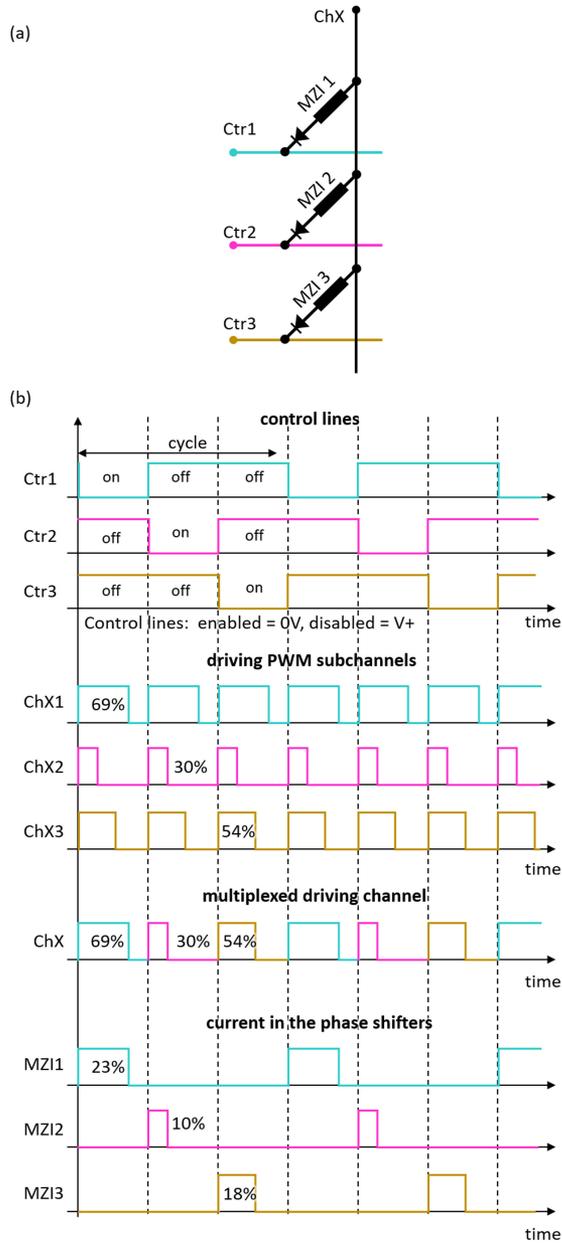
Fig. 3. Time multiplexing the signals on the driving channels. (a) A single driving channel ChX connecting three MZIs to control lines. (b) The control lines are each enabled during one third of a cycle. The subchannels $ChX1 - ChX3$ are modulated with a PWM duty cycle corresponding to the power that needs to be delivered to the corresponding MZIs. The subchannels are multiplexed onto the driving channel ChX synchronized with the control signals. The power is delivered to the MZIs during the corresponding parts of the cycle. Note that the control lines are at high voltage $V+$ when disabled, and at 0 V when enabled.

levels contain in total 15 tunable beam splitters (each with one phase shifter) which are connected in the 5×3 matrix shown in Fig. 2(b). This matrix circuit can be divided into 5 identical subcircuits along a single column, each containing one driving channel and 3 control lines, as shown in Fig. 3(a). This requires a total of eight contact pads (and PWM driving sources) to drive all fifteen phase shifters needed to operate the circuit simultaneously.

Fig. 3(b) shows the time traces of a single driving channel in a circuit with $N = 3$ control lines. Notice that only one control line is enabled at any given time, which guarantees that only one phase shifter per driving channel will draw current at that time. Now the channel time signal (e.g., ChX) is subdivided in three subchannels ($ScA1$, $ScA2$, and $ScA3$), and each subchannel is synchronized with the 'enable' signal of one control line ($ScX1$ with $Ctr1$, $ScX2$ with $Ctr2$, and $ScX3$ with $Ctr3$). This makes each phase shifter in the channel to be addressed only during the duration of its respective subchannel, allowing the driving signal to be time-shared to control all three phase shifters in the column.

As the signal is time-shared between three subchannels, and the MZI is only accessible during one-third of the time, the effective duty cycle seen by each MZI is the duty cycle of the subchannel divided by the total number of subchannels. That means that each MZI can be effectively driven by a PWM signal with a maximum duty cycle $D.C. = 1/N$, which corresponds to $1/3 = 33.3\%$ in this circuit.

B. Limitations and Considerations

Time-sharing actuation is only possible due to the very long time constant of the thermo-optical phase shifters, as we are driving the phase shifters with a non-constant signal. That means that the frequency of the enabling signal (used to select which row is enabled) has to satisfy the requirements imposed by reaction time of the phase shifters [14]. The same applies to the frequency of the PWM signal applied to the driving channels. For our heaters with a $\tau_{ps} = 80$ μs, this means we have to use a PWM clock of at least 200 kHz [14].

A second important characteristic of time-multiplexed access to the heaters is that only a fraction $1/N$ of the total cycle can be used by a phase shifter, as we can only enable one row at a time. That means that the phase shifter receives at most 33% of the electrical power corresponding to the driving voltage V of the PWM signal. To guarantee that this does not become a limiting factor we have to design the phase shifters for $3\times$ the electrical power which they would normally draw if they have an entire channel to themselves. If we want every phase shifter to cover the full range of $0 - 2\pi$ for a driving voltage of $V+$ while being multiplexed, we should actually design it to induce a maximum phase shift $\Delta\phi = N \cdot 2\pi$rad at $V+$ when operated in DC. This overdimensioning means that the thermo-optic phase shifters will also be larger and have a longer optical length to cover a full 2π phase shift, compared to the non-multiplexed situation. Increasing the time multiplexing factor N requires proportionally larger heaters that draw a larger electrical peak power.

The number of thermo-optic phase shifters that can be time-multiplexed in a same driving channel is directly proportional to the time constant of the phase shifters: slower heaters will allow multiplexing more control lines N. It also depends on the frequency of the driving signal and the power efficiency of the phase shifters. For multiplexed applications it is desirable to have phase shifters that are slow enough, so low-frequency PWM signals can be used, avoiding the use of complex electronics for generating (and amplifying) high speed signals. At the same time it is desirable to have phase shifter with high power efficiency, as more phase shifters sharing a driving channel reduces the time fraction for each phase shifter in this channel. From the point of view of time multiplexing, it is a good thing that mechanisms that
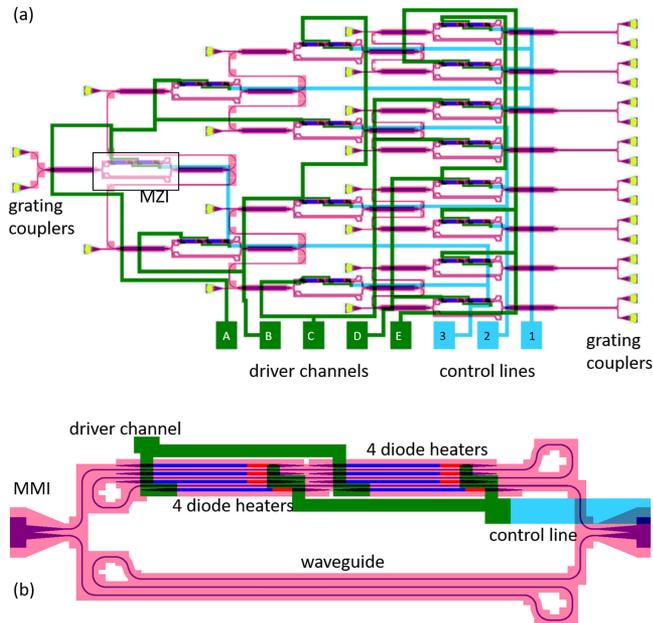
Fig. 4. (a) Mask layout of the 1×16 multicast network. (b) Zoom of the arms of each 2×2 MZI, with a diode heater in the top arm.

make a heater more efficient (e.g., undercutting the substrate) also make the heater slower [11], which makes it possible to drive the PWM signal at a lower frequency, or multiplex more heaters at the same frequency.

IV. EXPERIMENTAL DEMONSTRATION

To experimentally demonstrate the column-row matrix addressing driving capabilities we implemented the 1×16 multicast network shown in Fig. 2, with the 3×5 matrix arrangement.

A. Chip Design and Fabrication

The silicon photonic chip was fabricated in IMEC's iSiPP50 G process [18] through the Europractice multi-project wafer service [19]. The silicon photonics process supports 3 waveguide etch depths (we used a full 220 nm etch for waveguides and heaters), multiple p-type and n-type dopant implants, and 2 metallization layers to wire up the heaters.

The optical unit cell is a 2×2 MZI (where one input port is unused), using the standard 2×2 multimode-interferometer from the design kit as a beam splitter. Each MZI is balanced, so designed to couple all the light to the cross port when the heater is unpowered. The circuit layout is shown in Fig. 4. The thermo-optic phase shifter is implemented as doped silicon stripes next to the optical waveguide, as shown in Fig. 1. The heater is contacted to the overlying Metal1 and Metal2 routing layers that connect it to the bondpads. Within the strip we embedded a p-n junction to implement a diode. To increase the efficiency of the heater, we route the waveguide three times between four heater lines in a paperclip fashion. With this configuration, we get an experimental heater efficiency of $P_\pi = 20.9$ mW. This is somewhat worse than the expected value for P_π of 15–17 mW, and as a result we are not able drive the phase shifters over a full 2π through the $3\times$ multiplexing.
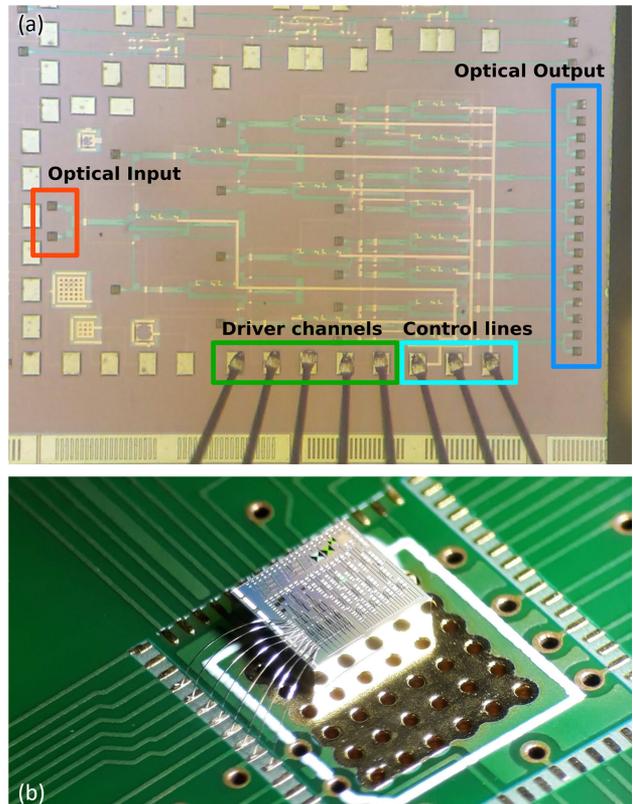
Fig. 5. (a) Microscope picture of the 3×15 multicast network shown in Fig. 2(b). (b) The mounted and wirebonded chip.

We designed the lengths of the phase shifter to operate in the voltage range of ± 7 V, and to induce a maximum phase shift $\Delta\phi = 6\pi$ rad at 7 V. This requires 2 segments of 50 μm for each of the four stripes connected in parallel and with a designed resistance of 410 Ω. This would give us a peak electrical power of 120 mW, good for an effective phase shift $\Delta\phi = 2\pi$ at duty cycle $D.C. = 33.3\%$, for a signal amplitude of 7 V, which would make this phase shifter fit to be used in a $N = 3$ time-sharing arrangement. The measured resistance is somewhat higher at ≈ 640 Ω, and as a result the current will be lower for the same voltage. At the peak value, when a channel applies the maximum phase shift in all subchannels, we draw a current of 11 mA, which corresponds with an electrical power of 77 mW. This is sufficient for a phase shift of $\approx 1.33\pi$ per channel. We measured a time constant for the heater of $\tau_{ps} = 80$ μs, which means the driving cycle should be repeated with a frequency of at least 200 kHz to make sure the signal is averaged out properly.

The matrix of 15 heaters is connected to 8 bond pads. The driver channels are wired on Metal1, while the control lines are wired on Metal2, using 10 μm wide lines.

B. Packaging and Measurements

To perform the measurements the PIC was mounted on a printed circuit board (PCB) and its electrical circuit was wirebonded to the board to allow the control of the heaters in the circuit. Fig. 5(b) shows a microscope picture of the optical circuit wirebonded to the PCB.
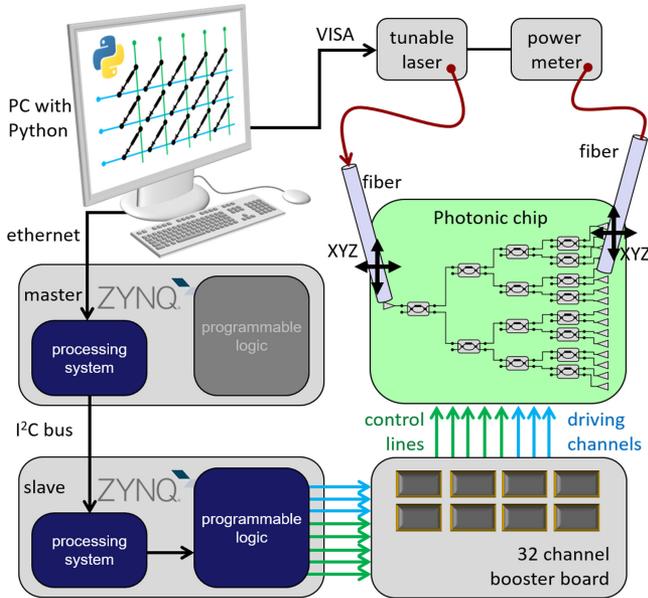
Fig. 6. Schematic of control and measurement architecture. Using Python, the desired driving power for each of the 15 heaters is transmitted to the master Zynq board, which dispatches it to one of its slaves which can control up to 100 PWM channels. 8 channels are actuated (3 control lines and 5 driving channels) and amplified through a current booster circuit which can handle up to 32 channels. The transmission of the photonic chip is measured using a tunable laser and power meter, coupled to the chip with grating couplers.

To measure the transmission of the chip, we used a pair of fibers on independent XYZ stages, oriented at a $10°$ degree angle to the vertical. Light from a tunable laser, polarization-controlled using manual polarization paddles, is coupled from one fiber into the input waveguide using a grating coupler. From the other side, the light can be collected from one of the output grating couplers. As we designed the grating couplers with a spacing of 50 μm, we cannot make use of a fiber array to measure all outputs simultaneously. This somewhat limits the functionality we can demonstrate with this chip.

C. Generating the Driving Signals

While the driving signals of the control lines are fairly simple, the driving signals of the channels, with their N subchannels, are more complicated to generate. In addition, all M subchannels should be synchronized with the enabling signal of the control lines. Because it is hard to perform such synchronization in software with a repetition rate of 500 kHz, we implemented the matrix driving routine in an FPGA.

We used a Xilinx Zynq 7010 system-on-chip on a MicroZed developer board [20]. This can drive up to 100 digital output channels independently. To synchronize a larger number of channels, multiple boards can be slaved to a single master board. The Zynq 7010 contains a processing system with a dual core ARM cortex-A9 microprocessor which runs FreeR-TOS, a multi-threaded real-time operating system. The code on the microcontroller performs the bookkeeping of which phase shifters are associated with which control lines and channels, and translates the desired phase shift into the required $D.C.$ for

each subchannel, which is then passed on to the programmable logic, which contains the FPGA. The routines programmed in the FPGA hardware generates the 256-level (8-bit) PWM signals for each subchannel and time-multiplexes the M sets of N subchannels into M independent channel signals, synchronized to the N control signals.

The output signals of the Zynq board are at 3.3 V and require a high-impedance load. Therefore, we placed an additional booster board in between the FPGA outputs and silicon chip to raise the output voltage to 7 V and buffer up to 100 mA of output current for the phase shifters in the circuit.

The complete measurement process is controlled from a PC through python scripts that access the external instruments with a VISA interface and communicate with the Zynq microcontroller over Ethernet.

D. Routing Light Through the Network

To demonstrate our control over multiple heaters, we configure the 1×16 network to direct all the input power to a single output. First we select an output by positioning the output fiber over one grating coupler. This is done using passive alignment with a camera as a visual reference, so the fiber position is not optimized with high precision. Then we we run a simple configuration algorithm to demonstrate the operation and multiplexing of the heaters in matrix addressing mode. We sweep the $D.C.$ of the leftmost, first-stage phase shifter. After a full sweep, we position the heater in the optimum position (i.e. where the transmission to the output fiber was maximum.) Then we proceed to the next level in the tree and perform the same sweep. Note that for each level, the previous levels continue to be controlled, as they maintain their optimum setting.

The power output time trace for this procedure is shown in Fig. 7 for two different outputs. After 4 levels, the 1×16 switch is in its optimal position. At this point, we are driving 4 heaters simultaneously, divided in 2 subcircuits with two heaters each.

Because the heaters cannot cover a full 2π phase shift, there could be situations where we cannot reach the optimum coupling. Better dimensioning of the heaters, or a slight increase of the driving voltage (here limited by our current booster board) can solve this.

V. SCALING

In this demonstration we used $N = 3$ multiplexing to demonstrate the matrix addressing. how many channels can be multiplexed is limited by a number of factors:

- The time constant of the heaters: Rather than faster phase shifters, we would benefit from a heater with a larger time constant. This can be accomplished by better insulation, for instance through undercutting the substrate [11], [12]. This would allow more control lines within a PWM cycle.
- The driving electronics and photonic-electronic integration: One can speed up the PWM clock to multiplex more channels. Here we used a 200 kHz clock, and we experimented with clock speeds up to 500 kHz. However, then the effect of parasitics and crosstalk become more pronounced, especially as the square-wave PWM signals contain much
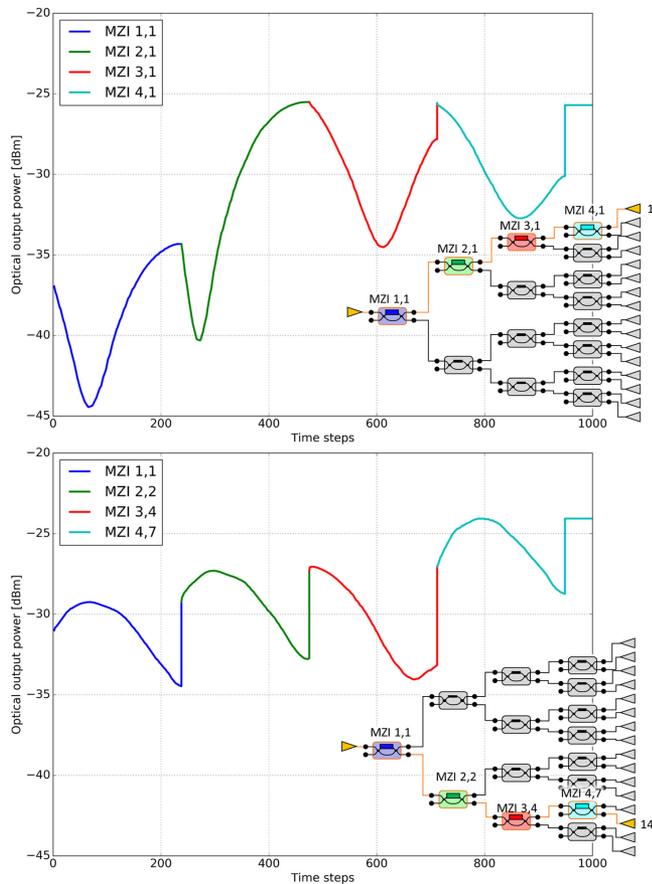
Fig. 7. Time trace of the powers in the 4 heaters as we configure the 1×16 circuit to direct the input to one particular output. We perform a sweep of the $D.C.$ along the path where the light has to go. After each sweep, the heater is set to its optimum value, and the next level in the tree is swept.

higher frequencies. In that case, the integration/packaging strategy of the electronics and photonics becomes critical.

- The efficiency of the heaters: More efficient heaters would reduce the reqower for a 2π phase shift, and therefore a larger number of phase shifters can be driven with the same driver power.
- The maximum current in the heater. As the multiplexing is increased, the heater will draw a larger peak current. This should not exceed the maximum current that would destroy the heater or the diode. This can be tuned by properly segmenting the heaters in parallel resistors.
- Increasing the drive voltage. The 7 V we used here could be increased to 9 V, as we the reverse breakdown voltage of the diodes is somewhat larger. As heaters respond quadratically to voltage, this would result in an increase in electrical power of 65%. Unfortunately, our driving circuit was designed to operate at a fixed supply voltage.

By simultaneously optimizing several of these variables using already demonstrated engineering approaches, we think the multiplexing can be boosted to $N = 10 - 15$. This would make it possible to control a circuit with 1000 phase shifters using only ≈ 100 pads, which can be easily accommodated with low-cost wirebonding techniques.

One criterion which we did not discuss this demonstration is the optical path length of the heater. When used in interferometric circuits, it can be important to have a short optical path length. While our paper-clip shaped heater is quite efficient, it has a long optical path length. Using more efficient heaters, e.g., by using undercut, would also shorten the path.

VI. CONCLUSION

We introduce a technique for driving large numbers of thermo-optical phase shifters in silicon photonics circuits. Taking advantage of the high time constant of thermo-optical phase shifters ($\tau_{ps} \approx 80$ μs), we use a digital PWM signal to control the phase shift. The use of a digital signal has advantages such as the simplification of the power source (as we can use a fixed-voltage source instead of a variable one) and the better linearity of the response of the phase shifter.

We then use the PWM signal to implement multiplexed control, connecting the heaters in a matrix arrangement that enables the driving of $N \times M$ phase shifters with only $N + M$ bondpads and power sources. For this, we need to time-multiplex the driving signal, which is again possible thanks to the high time constant of the thermo-optical phase shifters. The multiplexing itself is enabled by p-n junctions embedded in the silicon resistors, effective placing a diode in series with the heater. The asymmetry between forward and reverse bias of the diode allows us to select any heater in the $M \times N$ matrix through its control line and driving channel.

We demonstrated this using a 1×16 multicast circuit which incorporates 15 MZIs with a single thermo-optic phase shifter each. Organized in $M = 5$ subcircuits connected to $N = 3$ control lines, we can control the 15 phase shifters with only 8 pads. Using an FPGA to synthesize the multiplexed PWM signals for the driving channels, we show that we can arbitrarily direct the light from the input to any output.

The scaling of this technique is limited by the maximum phase shift the heater can induce (at full electrical power) and the time constant, as well as several electrical design constraints. While in this example we limit ourselves to $N = 3$ multiplexing, more efficient heaters (which in general also have a longer time constant) could boost this with another factor of $3 - 5\times$.

REFERENCES

[1] X. Chen *et al.*, "The emergence of Silicon Photonics as a flexible technology platform," *Proc. IEEE*, vol. 106, no. 12, pp. 2101–2116, Dec. 2018.

[2] M. Smit, K. Williams, and J. van der Tol, "Past, present, and future of InP-based photonic integration," *APL Photon.*, vol. 4, no. 5, May 2019, Art. no. 050901.

[3] A. Khanna *et al.*, "Complexity scaling in silicon photonics," *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2017, pp. 11–13.

[4] A. Rahim, T. Spuesens, R. Baets, and W. Bogaerts, "Open-access silicon photonics: Current status and emerging initiatives," *Proc. IEEE*, vol. 106, no. 12, pp. 2313–2330, Dec. 2018.

[5] Z. Lu *et al.*, "Performance prediction for silicon photonics integrated circuits with layout-dependent correlated manufacturing variability," *Opt. Express*, vol. 25, no. 9, pp. 9712–9733, 2017.

[6] W. Bogaerts, Y. Xing, and M. U. Khan, "Layout-aware variability analysis, yield prediction and optimization in photonic integrated circuits," *IEEE J. Sel. Topics Quantum Electron.*, vol. 25, no. 5, Sep./Oct. 2019, Art. no. 6100413.

[7] D. Pérez, I. Gasulla, and J. Capmany, "Programmable multifunctional integrated nanophotonics - Review article," *De Gruyter - Nanophoton.*, vol. 7, no. 8, pp. 1351–1371, 2018.

[8] A. Ribeiro, A. Ruocco, L. Vanacker, and W. Bogaerts, "Demonstration of a 4 × 4-port universal linear circuit," *Optica*, vol. 3, no. 12, 2016, Art. no. 1348.

[9] D. A. B. Miller, "Self-aligning universal beam coupler." *Opt. Express*, vol. 21, no. 5, pp. 6360–6370, 2013.

[10] J. Carolan *et al.*, "Universal linear optics," *Science*, vol. 349, no. 6249, pp. 711–716, 2015.

[11] A. Masood *et al.*, "Comparison of heater architectures for thermal control of silicon photonic circuits," in *Proc. 10th Int. Conf. Group IV Photon.*, 2013, pp. 83–84.

[12] P. Dong *et al.*, "Thermally tunable silicon racetrack resonators with ultralow tuning power." *Opt. Express*, vol. 18, no. 19, pp. 20 298–20 304, 2010.

[13] V. Stojanović *et al.*, "Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes [Invited]," *Opt. Express*, vol. 26, no. 10, pp. 13106–13121, 2018.

[14] A. Ribeiro and W. Bogaerts, "Digitally controlled multiplexed silicon photonics phase shifter using heaters with integrated diodes," *Opt. Express*, vol. 25, no. 24, pp. 29778–29787, 2017.

[15] D. J. Goodwill, L. Chrostowski, and H. Mehrvar, "Photonic elements driven by common electrical driver," Patent US20170276969A1, 2015.

[16] B. J. Lechner, F. J. Marlowe, E. O. Nester, and J. Tults, "Liquid crystal matrix displays," *Proc. IEEE*, vol. 59, no. 11, pp. 1566–1579, Nov. 1971.

[17] I. Zand and W. Bogaerts, "MP6 - discretization effects of digital control of thermally tunable 2x2 MZI couplers," in *Proc. IEEE Photonics Soc. Summer Topical Meeting Ser.*, 2019, pp. 1–2.

[18] M. Pantouvaki *et al.*, "Active components for 50 Gb/s NRZ-OOK optical interconnects in a silicon photonics platform," *J. Lightw. Technol.*, vol. 35, no. 4, pp. 631–638, 2017.

[19] Europractice, "Europractice multi-project wafer service," [Online]. Available: http://www.europractice-ic.com/. Accessed on: 2020.

[20] V. Boppana *et al.*, "UltraScale+ MPSoC and FPGA families," in *Proc. IEEE Hot Chips 27 Symp.*, 2016, pp. 1–37.

Umar Khan (Member, IEEE) received the Ph.D. degree from Tyndall National Institute, University College Cork, Cork, Ireland for his work on "Design and implementation of micro-structures with refractive index contrast for optical interconnects and sensing applications." After completing the Ph.D. degree, he joined Photonics Research Group with Ghent University—IMEC, Belgium as a Postdoctoral Researcher. His current research interests include compact models, variability aware design, parameters extraction and programmable circuits.

Mi Wang received the B.Eng. degree in optoelectronics from the Department of Electronics and Engineering, Sichuan University, Chengdu, China, in 2013, and the M.Sc. degree from Friedrich–Alexander University, Erlangen, Germany in 2016. She is currently working toward the Ph.D. degree with Photonics Research Group advised by Wim Bogaerts from Ghent University, Ghent, Belgium. She joined Huawei European Research Center after her master's degree. Her main research interest was on LCoS SLM study and its application in wavelength selective switch, econfigurable photonic circuits, and photoniccomponent design.

Lukas Van Iseghem received the master's degree of industrial engineering in electronics, KU Leuven, Leuven, Belgium, in 2016. He is currently working toward the Ph.D. degree in the Photonics Research Group with Ghent University—IMEC. His current research interests include programmable integrated photonics, optical material integration, new approaches to tune optical devices and analog computing.

Antonio Ribeiro received the Ph.D. degree from the Photonics Research Group—IMEC, Ghent University, Ghent, Belgium for his work on programmable integrated silicon photonics circuits, where he published his thesis entitled "Building Blocks and Subcircuits for Programmable silicon Photonics Circuits. He is currently employed as a Silicon Photonics Engineer with Brolis Semiconductors, Vilnius, Lithuania.

Wim Bogaerts (Senior Member, IEEE) received the Ph.D. degree in the modeling, design, and fabrication of silicon nanophotonic components from Ghent University, Ghent, Belgium, in 2004. He is a Professor in the Photonics Research Group with Ghent University—imec. During this work, he started the first silicon photonics process on imec's 200 mm pilot line, which formed the basis of the multi-project-wafer service ePIXfab.

In 2014, he Co-Founded Luceda Photonics, a spin-off company of Ghent University, IMEC and the University of Brussels (VUB). Luceda Photonics develops unique software solutions for silicon photonics design, using the IPKISS design framework.

Since 2016, he has been again a Full-Time Professor with Ghent University, looking into novel topologies for large-scale programmable photonic circuits, supported by a consolidator grant of the European Research Council. He also coordinates the H2020 project MORPHIC on MEMS-enables programmable photonic ICs.

His current research focuses on the challenges for large-scale silicon photonics: Design methodologies and controllability of complex photonic circuits. He has a strong interest in telecommunications, information technology and applied sciences. He is a member of Optical Society of America OSA and SPIE.

Sibert Declercq received the graduate degree from Ghent University, Ghent, Belgium, in 2019, and received the master's degree in industrial engineering in embedded systems. In his masters thesis, he researched the usage of heaters to control on-chip silicon waveguides. This research is a practical step to configurable silicon photonic chips. He is currently employed as a Development Engineer with an engineering firm, which develops custom embedded systems for other companies.