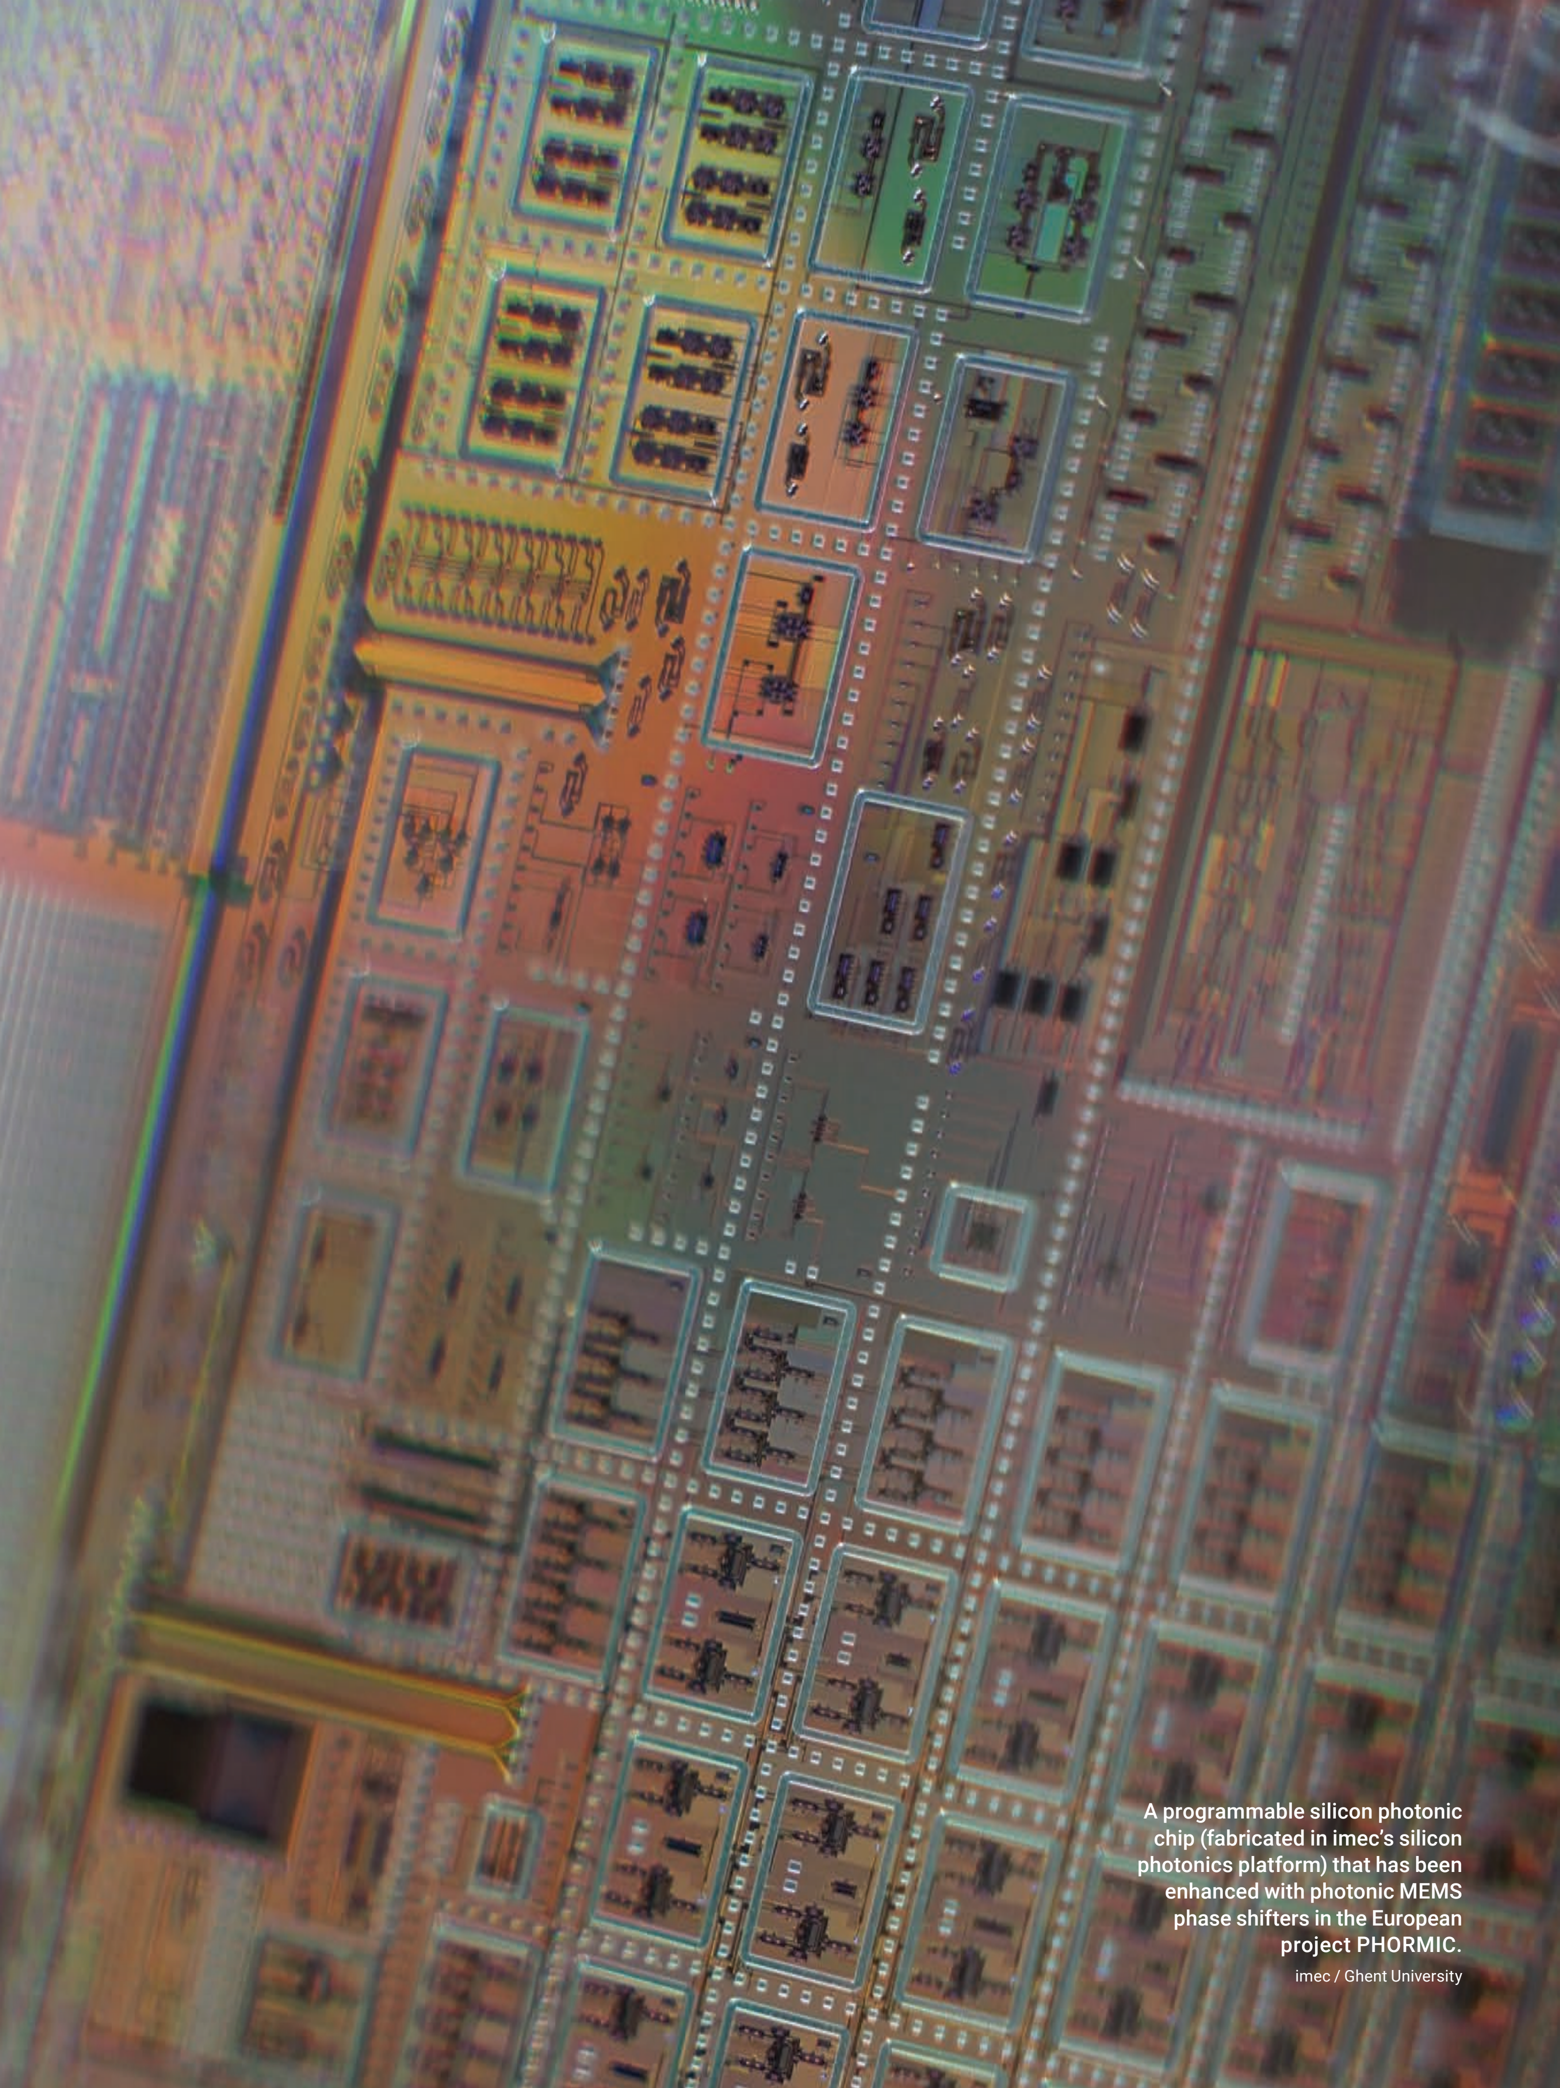

A programmable silicon photonic chip (fabricated in imec's silicon photonics platform) that has been enhanced with photonic MEMS phase shifters in the European project PHORMIC.

imec / Ghent University

Programmable Photonics

Wim Bogaerts

With the addition of programmability and software, photonic chips can enter a new era of flexibility, creativity and diverse applications.

Photonic chips, also called photonic integrated circuits (PICs), have become vital in our daily lives, often in unseen ways. Data centers wired up with millions of fiber optic connections, powered by transceivers with photonic chips, keep the internet and machine-learning models running. Light, after all, is an ideal information carrier. And photonic chips—combining light sources, electro-optic modulators, wavelength filters, photodetectors and other building blocks into an optical circuit—are the ideal technology for optical information processing, especially as the chips become increasingly complex and configurable. Yet even today, PICs are used mostly as simple converters between electrical and optical signals in fiber optic networks.

It is somewhat surprising that advanced light-based technology, central to communications, is not more pervasive in everyday applications. After all, light has many other purposes; it provides our primary window on the universe through our eyes, and our ever-present smartphones and wearables pack multiple optical sensors and displays. Still, these are mostly based on discrete optical components, not PICs. The contrast between photonics and electronics is stark: While almost every device we own contains at least one electronic chip, photonic chips have yet to break through into applications outside of communications.

In recent decades, plenty of lab demonstrations have highlighted the capabilities and usefulness of PICs for biosensing, metrology, biomedical diagnostics, lidar, quantum computing and machine learning. By integrating many optical functions on a single chip, PICs can deliver more complex functionality, lower power consumption and a smaller form factor. They can also be fabricated with the same type of processes as electronic chips, in large volumes at a potentially low cost. Why are abundant lab demonstrations that take PICs beyond transceiver applications not making the transition to commercial products?

1047-6938/24/03/34/8-$15.00 ©Optica

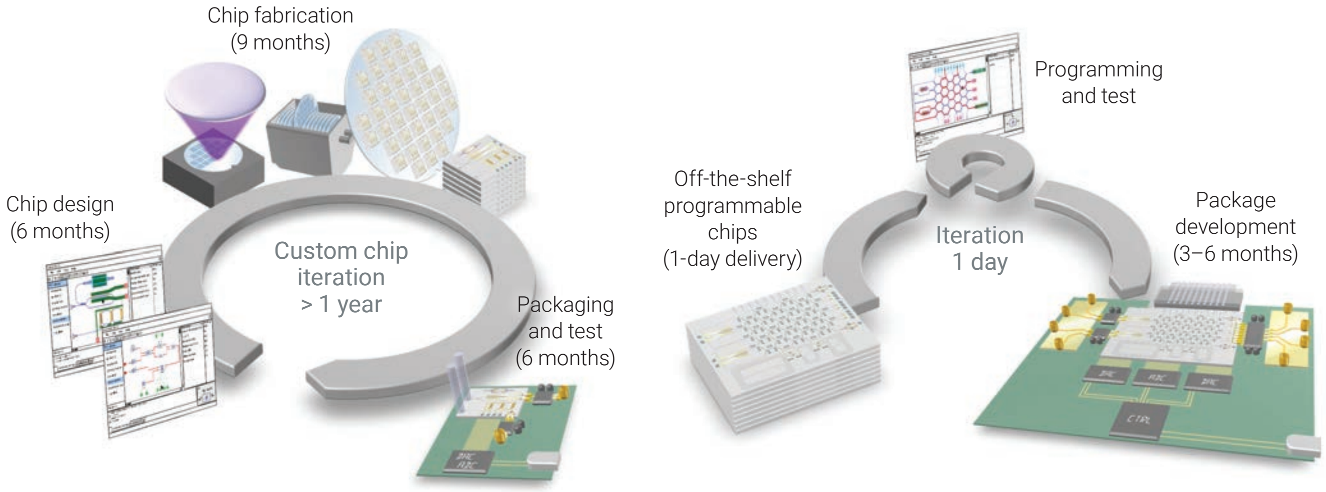

Off-the-shelf, programmable chips can drastically shorten the development timeline for new products. This concept has been proven spectacularly in electronics. Can it transfer to photonics?

Beyond chicken and egg

Part of the answer is that the PIC ecosystem faces a chicken-and-egg problem. Pervasive deployment of PICs relies on large-volume fabrication capabilities. But large-volume fabrication infrastructure is only cost-effective if there is a sizable addressable market.

Even though developments in silicon photonics now provide state-of-the-art fabrication facilities for photonic chips, the supporting ecosystem does not yet make it easy for new entrants. Circuit design and simulation methods, though rapidly improving, are not yet predictive enough to guarantee first-time-right design cycles. And there are no cost-effective standards to integrate PICs into larger systems with driver electronics, high-speed connections and free-space optics. This puts a brake on the transition from lab to product; translating an idea into a working photonic chip requires many costly fabrication cycles—and this often takes longer than a tech startup can afford.

To tear down this barrier, we should look at the mechanisms that have made electronics technology so accessible. Photonics technology today has a maturity comparable to that of electronics in the early 1990s. Around that time, two major paradigm shifts boosted the adoption of complex electronics, powering the maker revolution of the 2000s. And both can move PICs beyond their chicken-and-egg problem.

The first catalyst for electronics in the 1990s was a rapid maturing of the design methodologies and tools, with accurate predictive circuit models and the introduction of complex design libraries of validated functions. This both simplified the design process (making it possible for more people to design) and reduced the number of cycles needed to get a working prototype.

The second electronics breakthrough was programmable circuitry. The introduction of field-programmable gate arrays (FPGAs) revolutionized the prototyping of digital electronics and made it possible to innovate on a timescale of weeks, rather than the year-long design/fabricate/test cycles of application-specific integrated circuits (ASICs). Suddenly, high-performance digital electronics were accessible to millions of technically skilled creators. The FPGA ecosystem and improving design tools then facilitated the migration from working FPGA demonstrations to new ASICs, boosting the transition from low-volume products to high-volume markets.

Photonics needs similar tools that open the technology to creativity and ideas outside the small community of photonic-chip specialists. With this, the technology will proliferate to a wide diversity of applications, driving volumes up and prices down, which in turn will enable even faster adoption.

PICs are becoming programmable

Just such a technology of programmable photonics is emerging today. No clear definition exists for programmable photonics—a term used interchangeably over the past decade or so with other catchy labels, such as photonic processors, universal optics and photonic FPGAs. In general, however, a programmable photonic chip makes it possible to reconfigure optical pathways and functionality using electronic control, without the need to fabricate a new photonic chip. In practice, this means that the photonic chip has a large number of

A programmable photonic chip makes it possible to reconfigure optical pathways and functionality using electronic control, without the need to fabricate a new photonic chip.

electrically controlled actuators that locally modify the optical properties.

How does this work? Let’s start from an ordinary (nonprogrammable) PIC. The chip might contain multiple optical functions: the generation of light (lasers or LEDs); modulation of high-speed electrical signals onto a light beam (modulators); conversion of optical signals back into electrical signals (photodetectors); filters that separate optical beams based on the wavelength or polarization.

These functions are connected together into a circuit using on-chip waveguides—lithographically defined, micrometer-scale lines of high-refractive-index material that confine light using the same mechanism as an optical fiber. Adding optical splitters and power couplers allows light to be distributed over multiple paths and combined again in on-chip interferometers or resonators. With these, wavelength filters can be constructed, engineering the delays to obtain constructive or destructive interference for specific wavelengths.

PIC technology has been around for half a century, built on different material technologies (glass, polymers, III–V semiconductors, lithium niobate, silicon and others), depending on the needed functionality. Pushing to higher refractive-index contrasts (silicon nitride, silicon) has made it possible to shrink the waveguide dimensions, but at the cost of making the optical properties more sensitive to fabrication variations. To compensate for this, active tuners were incorporated, usually in the form of on-chip microheaters. The heating of such a phase shifter locally changes the refractive index in the waveguide, tuning the chip’s operating point.

Basic tuning rapidly led to active control of the flow of light in tunable couplers and switches. When a phase shifter is put into the waveguide arms of a Mach-Zehnder interferometer (consisting of a 2×2 splitter, two waveguide arms and a 2×2 combiner), we can now tune how much light ends up in one output or the other. Combining such a tunable coupler with an additional phase shifter enables control of the phase delay between the outputs—creating a so-called analog 2×2 optical gate.

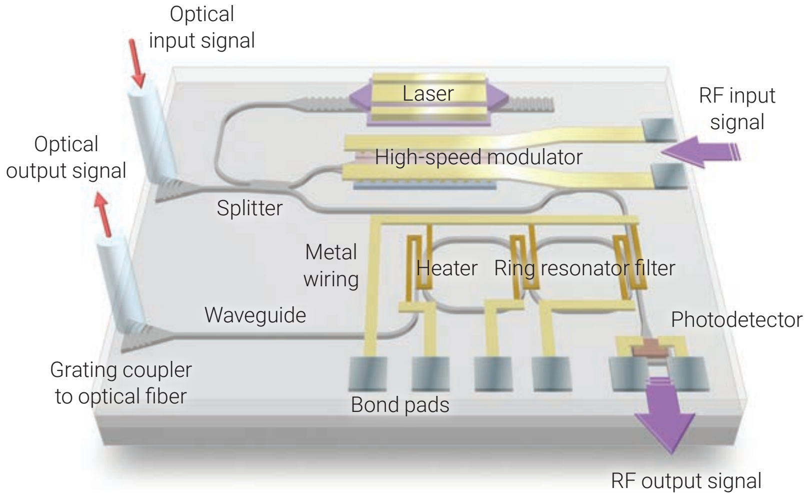

A simple photonic integrated circuit (PIC).

Forward-only waveguide meshes

Real programmable photonics came about with the realization that such 2×2 gates can be organized into a mesh (see figure, p. 37). In such a mesh, the coupling between the waveguide ports is entirely determined by the coupling states and the optical phase delays, which control constructive and destructive interference between the many paths through the mesh. The only components needed are phase shifters and tunable couplers (which can also be constructed from phase shifters).

The first proposed mesh architectures have been forward-only meshes, in which light flows in one direction between a set of input waveguide ports and a set of output waveguide ports. When properly dimensioned, such meshes can generate any possible linear combination of the light in the inputs. If the amplitudes and phases of the optical waves at the input ports represent a vector of complex numbers, the circuit itself will implement a matrix–vector multiplication (MVM), with the result represented by the amplitudes and phases of the optical output waves. This near-instantaneous arithmetic has been one of the main drivers for developing programmable photonics, as MVM operations lie at the core of many neural-network and machine-learning algorithms. The same mesh architectures have also proven very valuable for quantum information processing, their components acting as potential quantum logic gates when used with single photons.

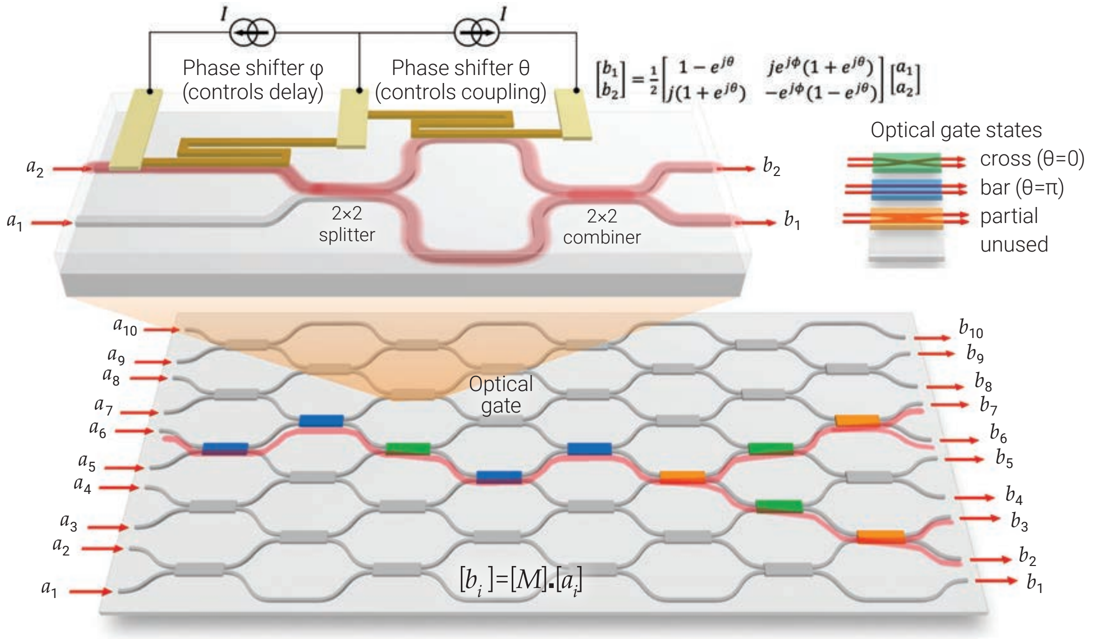

A forward-only waveguide mesh. By configuring the two phase shifters in each optical gate, any linear combination M of the input waves can be constructed, effectively implementing a real-time analog matrix–vector multiplication.

Some forward-only mesh architectures can also be combined with self-configuration routines, based on detectors either inside the mesh or at the outputs. Embedded detectors in particular have the advantage that the mesh can continue to adapt itself to changes in the input signals—for instance, to maximize the coupled power or separate multiple input signals into individual receivers.

Yet forward-only meshes also have their limitations as a general-purpose tool similar to electronic FPGAs. First, the port assignment is fixed to either input or output. Also, forward meshes are usually balanced in terms of path length. That makes them tolerant and broadband—but also complicates the implementation of wavelength filters and optical signal-processing functions such as equalization, integration or differentiation. And as forward meshes do not have recurrence, constructing optical resonators is not possible.

Recirculating waveguide meshes

Recirculating meshes might address these shortcomings by organizing the same optical gates in coupled waveguide loops. Regular arrangements for recirculating meshes use triangular, square or hexagonal unit cells, with the latter showing the most flexibility to implement different delay lines and define connectivity. More complex mesh arrangements, using unit cells of different shape or with a nonuniform distribution of optical path lengths, have been proposed but not demonstrated.

Recirculating meshes make it possible to assign every port as input, output or both, and to define all-to-all connections. They can also be configured to behave like forward-only meshes, with the caveat that they will require more optical gates than a dedicated forward-only mesh. But the key differentiator for a recirculating mesh is that it makes it possible to implement delay lines, which form the basis of on-chip interferometers. One can also define ring resonators, to build wavelength filters with a finite or infinite impulse response, simply by programming the coupling state and phase delay of the optical gates.

Though electronically tuned, the waveguide mesh provides only passive optical functions, in the form of a linear transformation on the injected signals. If we want to add active optical functionality such as gain, modulation, detection or nonlinearities, these building blocks need to be added to the circuit. The most straightforward approach is to place them at the periphery and use the mesh itself to connect the blocks into a circuit.

Because of its built-in loops, a recirculating waveguide mesh is more complicated to configure than a forward-only mesh. A large fabricated photonic circuit will not have ideal behavior, as all components deviate slightly from the ideal; this translates into imperfect phase delays and coupling values for each optical gate. In a forward-only mesh, these imperfections can be calibrated quite easily, because all optical paths are almost balanced in length and therefore the response of the entire circuit is broadband. In a recirculating

To make a PIC programmable, hundreds or thousands of phase shifters need to be integrated. And the electrical control of such a large number of devices brings along its own challenges.

mesh, light will traverse the mesh along paths with very different lengths, including resonant loops. The result is a very wavelength-dependent response, even when every building block has a broadband response. Calibration routines need to take this wavelength dependence into account.

To date, only a few recirculating meshes have been demonstrated. One compelling example has come from the Polytechnic University of Valencia, Spain, and its spinoff company iPronics. This group has demonstrated several recirculating waveguide meshes with 7 to 17 hexagonal unit cells, showing how these meshes can be used to define wavelength filters based on combinations of Mach-Zehnder interferometers and ring resonators, tunable optical delay lines, power distribution networks and linear transformations.

Many of these functions are configured in a discrete way, with light routed through an integer number of optical segments that determine the delay in an interferometer or resonator. The length of a single segment therefore also determines the granularity of the configuration and the shortest possible path-length difference that can be programmed. In terms of functionality, this translates into the free spectral range of the wavelength filters. Making the segments shorter makes it easier to process wide-bandwidth optical signals.

Chip technologies for programmable photonics

The underlying chip technology for programmable PICs is the same as for nonprogrammable ones. The key difference is in scale: To make a chip programmable, hundreds or thousands of phase shifters need to be integrated. As a result, light must pass through tens or hundreds of tuners, which need to have low optical losses and low electrical power consumption. The electrical control of such a large number of devices brings along its own challenges in co-integrating photonics with driver electronics, in crosstalk, and in the calibration,

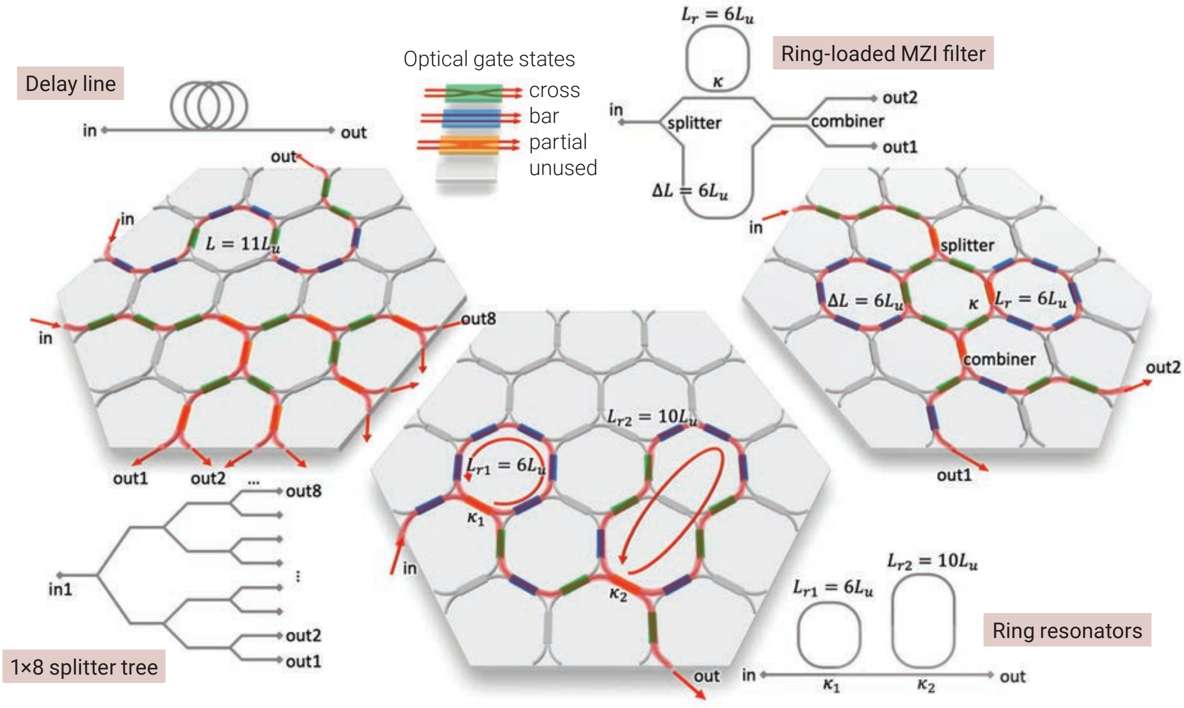

Recirculating waveguide mesh with a hexagonal lattice. Functions can be defined by setting the gates to cross, bar or partial coupling state. Left: an optical delay line and a 1×8 splitter tree. Center: two ring resonator filters with different lengths. Right: a ring-loaded MZI filter.

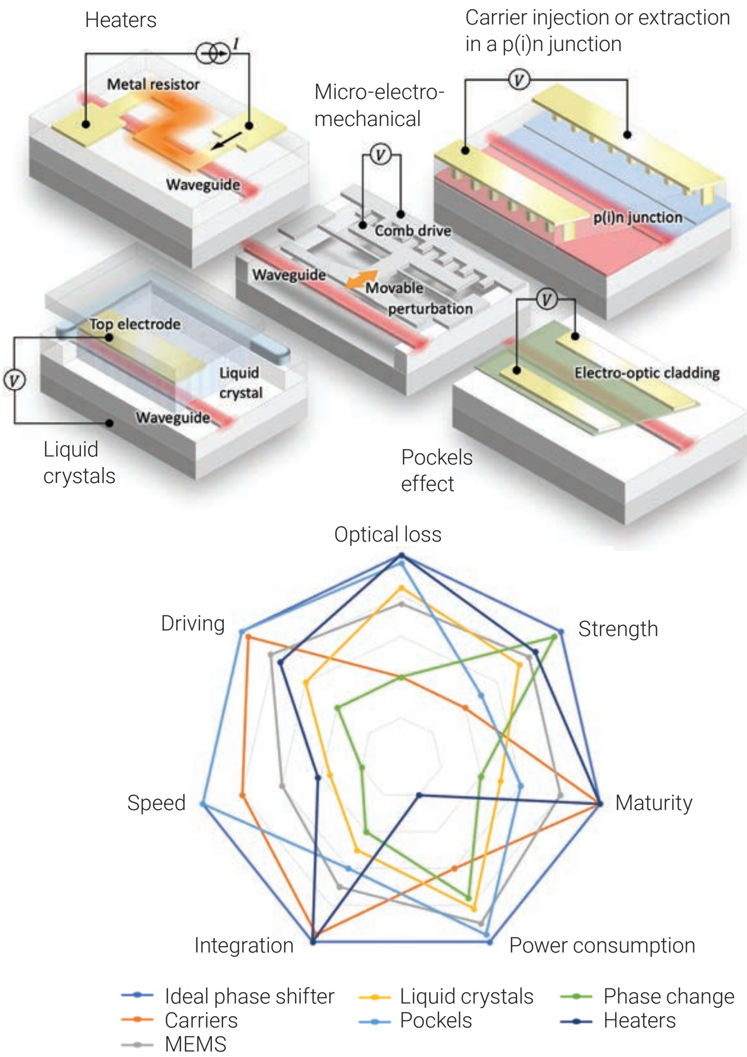

Top: Implementations of electro-optic phase shifters in silicon photonic circuits. Bottom: Comparison of different phase-shifter metrics.

configuration and control routines to make sure that the chip performs the desired function.

An "ideal phase shifter" is the dream for every photonic-circuit designer, and much research is being focused on the problem. For years, the mainstream approach has been thermo-optic, positioning microheaters near waveguides. These must continually consume power to maintain their state and will induce crosstalk when packed close together, complicating control of larger circuits. Locally removing the substrate and etching insulation trenches, thereby trapping the heat inside the small volume around the waveguide, can boost the efficiency of on-chip heaters but also slows down the phase-shifter response.

Alternative phase-shift mechanisms are thus being explored. Free carriers (electrons and holes in semiconductors) can influence the optical phase, but they introduce excessive absorption losses. Ferro-electric materials such as lithium niobate support intrinsic electro-optic actuation through the Pockels effect, and new thin-film platforms provide both high-speed modulation and efficient electro-optic phase shifters. These phase shifters have much lower power consumption than heaters, but their effect is weaker. As a result, they (as well as phase shifters based on strain and piezo-electric effects) must be physically longer, which, in a recirculating mesh, reduces the granularity for defining delays.

Still other phase shifters rely on micro-electro-mechanical devices that perturb the waveguide with a movable element, offering tuning strengths similar to heaters but with near-zero static power consumption. Similar efficiencies have been demonstrated with liquid-crystal claddings, using giant optical birefringence to adjust the propagation constant in the waveguide. So-called phase-change materials—compounds that can transition between stable states (such as crystalline and amorphous) via controlled heating and cooling—also offer large refractive-index changes. These approaches, however, are not yet sufficiently mature to be used reliably in large-scale circuits, integrated with the capabilities of existing photonics platforms.

From chip to system

The PIC constitutes the heart of a programmable photonic circuit. But to make it work, the PIC must be integrated with electronic drivers and readouts, as well as with the software routines to make it truly programmable by the user. The scale of programmable photonic circuits, much larger than a typical application-specific PIC, makes this a major challenge. Hundreds or thousands of electro-optic phase shifters need to be connected to their own electronic driver, just like every monitor photodiode. This requires a very large number of electrical wires, beyond the capacity of ordinary wire bonding or flip-chipping. 2.5D integration techniques using interposers can provide a solution here, but the complexity is still staggering.

Apart from the connections to the driver electronics, the photonic chip must also connect to the outside world. Optical interfaces are usually handled by fibers, for which coupling solutions exist. But the many fibers need to be combined with the existing electrical wiring for control as well as high-speed connections for radio-frequency input and output signals to the modulators and from the photodetectors. A full programmable photonic system thus involves a daunting complexity.

Further, the combined photonic and electronic hardware is not functional without software to make

The strongest push today for programmable photonics is in the field of computational accelerators for matrix–vector multiplication, propelled by machine learning and AI.

the circuit programmable. With the first working hardware demonstrators now operational, both for forward-only and recirculating meshes, innovation is now emerging on the software level, in calibration routines, routing and distribution algorithms, adaptive power coupling, filter synthesis and even automated optimizations for arbitrary transfer functions. The format in which these are defined and the language formalisms to describe the problems for photonics are still in their infancy, however, and it will be interesting to see how they will evolve.

Toward a programmable-PIC ecosystem

Programmable photonics comes in different flavors, driven forward by different application cases. The strongest push today is in the field of computational accelerators for MVM, propelled by machine learning and artificial intelligence. Forward-only programmable photonic circuits show considerable promise to provide order-of-magnitude performance improvements compared with traditional digital tensor processors.

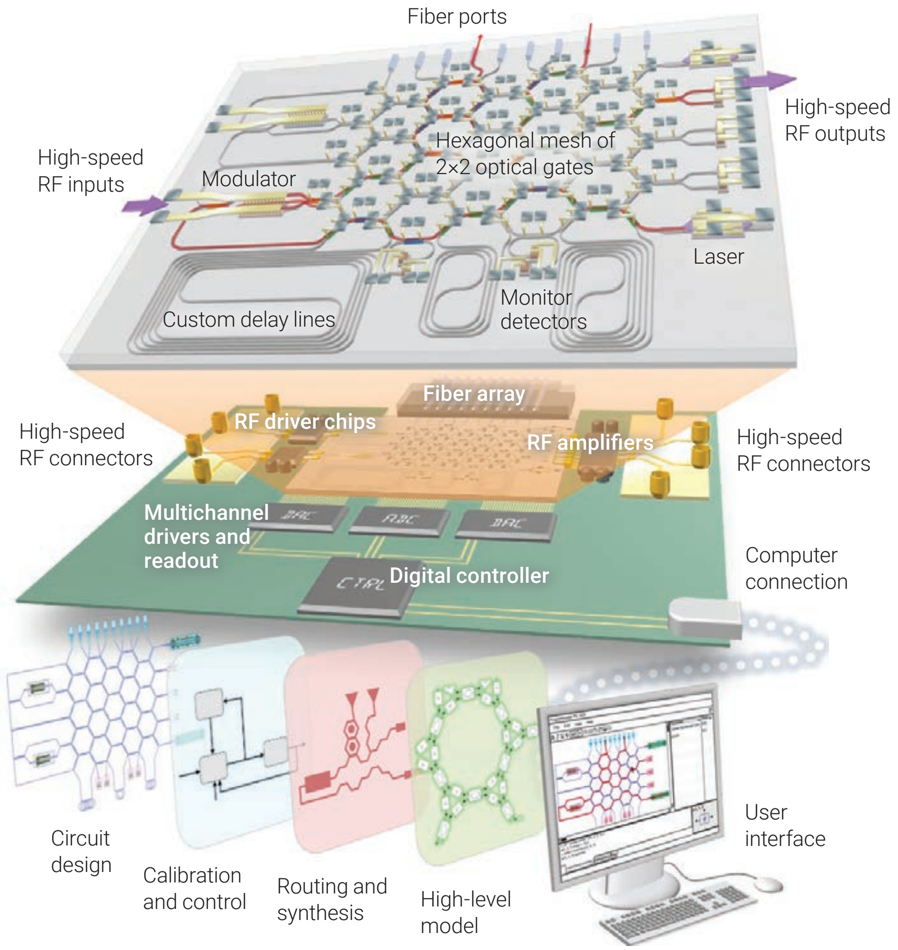

From a waveguide mesh to a full programmable photonic system. The waveguide mesh is connected on chip to lasers, high-speed modulators and detector, low-loss delay lines and monitor diodes. This chip is then interfaced to a fiber array, electronic driver and readout circuitry and a digital controller (for example, an FPGA). The user then interfaces to the chip using multiple layers of programming.

The other compelling case for programmable photonics is not tied to a particular application, but instead inspired by the promise of countless new applications—and the need for critical chip volumes to warrant the required investments today. This model is much more speculative; there is always a trade-off between flexibility and raw performance. Will programmable photonic chips become good enough that they can really be used to build actual products, as FPGAs are used today in digital electronics? If so, they can truly enable innovators, driving down the fabrication cycles for product development and even becoming a cost-effective alternative for low-volume products (for high volumes, a dedicated chip is likely to be cheaper).

The photonics ecosystem needs the boost that programmable photonics can provide. Adding a software interface to a powerful technology opens it up to a much wider community of creative minds. We can only guess where the combination will take us. OPN

Wim Bogaerts (Wim.Bogaerts@UGent.be) is with Ghent University and imec, Ghent, Belgium.

For references and resources, go online: **optica-opn.org/0324-pics**.

References and Resources

- D.A.B. Miller. “Self-configuring universal linear optical component,” Photon. Res. **1**, 1 (2013).
- S.M. Trimberger. “Three ages of FPGAs: A retrospective on the first thirty years of FPGA technology,” Proc. IEEE **103**, 318 (2015).
- J. Capmany et al. “Microwave photonics: The programmable processor,” Nat. Photon. **10**, 6 (2016).
- A. Annoni et al. “Unscrambling light—Automatically undoing strong mixing between modes,” Light Sci. Appl. **6**, page e17110 (2017).
- D. Pérez et al. “Multipurpose silicon photonics signal processor core,” Nat. Commun. **8**, 636 (2017).
- W. Bogaerts et al. “Programmable photonic circuits,” Nature **586**, 207 (2020).
- W. Bogaerts and A. Rahim. “Programmable photonics: An opportunity for an accessible large-volume PIC ecosystem,” IEEE J. Sel. Top. Quantum Electron. **26**, 1 (2020).
- J. Wang et al. “Integrated photonic quantum technologies,” Nat. Photon. **14**, 273 (2020).
- H. Zhou et al. “Photonic matrix multiplication lights up photonic accelerator and beyond,” Light Sci. Appl. **11**, 30 (2022).