InP promises to turbo-charge ICs

If we stick to using copper interconnects in silicon ICs, then it's only a matter of time before we arrive at a performance-limiting data-transfer bottleneck. The solution: additional optical interconnects built from silicon waveguides and InP lasers and detectors, says **Dries Van Thourhout**.

Severe data-transfer bottlenecks between the different sections of a silicon IC are threatening to limit the performance of future generations of electronic circuits. This is because the continuing reductions in device dimensions are making it increasingly difficult to keep propagation delays at an acceptable level. Even if we take the most optimistic estimates for conductor resistivity and dielectric permittivity, it seems likely that ICs performance will not keep pace with that set out in the International Technology Roadmap for Semiconductors. It is clear that a radically different approach is needed for linking together the different parts of a silicon IC.

One promising option could be the addition of a high-speed optical interconnect layer that sits on top of, or in between, the classical copper interconnect layers of future microprocessor chips (see figure 1). Aside from the bandwidth advantage, the benefits of such an optical interconnect layer would be a reduced temperature sensitivity for the link, an immunity to electromagnetic noise, lower power consumption and the possibility for synchronous operation, both within the circuit and with other circuits.

In Europe we have been pursuing exactly this type of technology through a European Union funded project entitled "Photonic Integration layer on CMOS by wafer-scale integration (PICMOS)". This recently completed €4.5 million ($6 million) project, which ran for just over three years, focused on building an optical interconnect layer by combining a silicon nanophotonic optical waveguide circuit and microscale III-V lasers and detectors. The effort was led by researchers at the Ghent University Photonics Research Group that is affiliated to the Interuniversity Micro Electronics Center (IMEC), which are both in Belgium. It involved six other partners: the French companies ST Microelectronics and Soitec-owned TRACIT Technologies; France's INL (Institut des Nanotechnologies de Lyon) and CEA-LETI; the National Center of Scientific Research "Demokritos" in Greece; and the Technical University Eindhoven, in the Netherlands.

Optical interconnect layers for on-chip communications require very compact waveguide circuits, due to the lack of available space. Fortunately, silicon "wire" optical waveguides allow just that, and can be fabricated by classical lithography and etching on silicon-on-insulator (SOI) substrates with a 200–400 nm silicon top layer. Restricting the wire width to around 500 nm produces single-mode waveguides that are ideal for data-transfer applications, and have virtually lossless bend radii as small as 2 μm thanks to a high refractive index contrast. Spacing these waveguides close together is also possible, as they can be brought to within a micron of one another without inducing serious losses – a level of miniaturization comparable to the pitch of copper lines on the upper interconnect layers in today's CMOS chips.

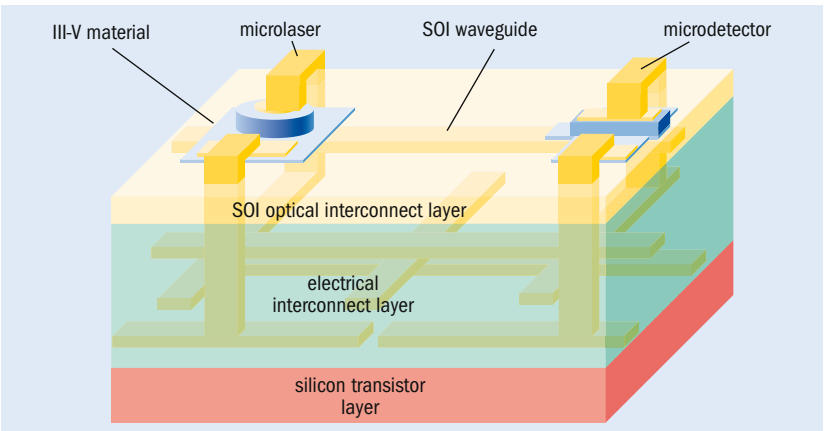

Fig. 1. On-chip optical interconnects could speed up data transfer between parts of a silicon IC.

Laser sources

Devices that have recently been built for these silicon-wire waveguides by groups in the US at Intel, Luxtera and Cornell, and a French partnership between LETI and IEF, feature very powerful filtering, wavelength selectivity and high-speed modulation. However, this approach suffers from a significant drawback: the difficulty of getting light out of silicon. Despite extensive research employing schemes using nanoscale features or silicon's nonlinear properties, it is unlikely that purely silicon-based lasers will reach an efficiency comparable to that of their III-V cousins for the foreseeable future. Consequently, we believe that it is imperative that the light-emitting function of III-V materials complements the strengths of the silicon-wire waveguide.

There has already been a large effort over the past few years to try and reach this goal by using advanced epitaxial growth techniques to form monolithically integrated structures. Our PICMOS project, however, has followed a radically different approach. This method, which is free from the lattice-matching constraints associated with producing III-V devices

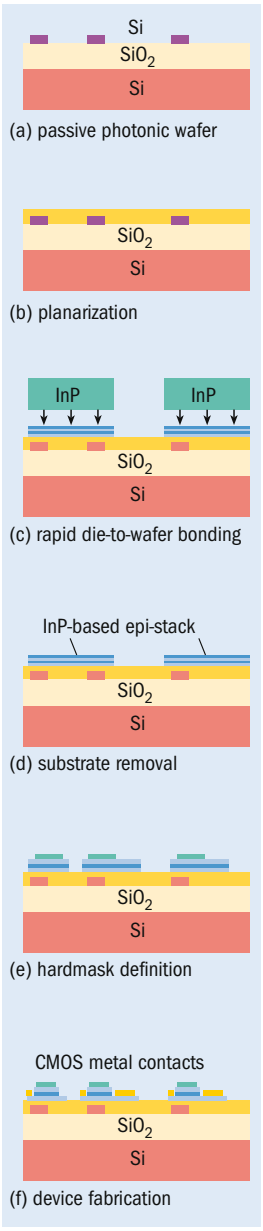

Fig. 2. (left) The PICMOS approach to building an InP microlaser coupled to a silicon-wire waveguide involves the fabrication of a passive photonic silicon wafer **(a)**, its planarization **(b)** and the addition of InP die **(c)**. The substrate is then removed from these InP structures **(d)**, before a mask defines the laser **(e)** and CMOS metal contacts are added **(f)**.

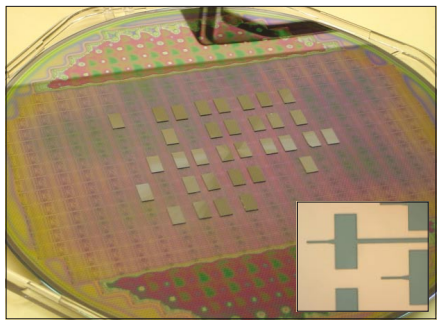

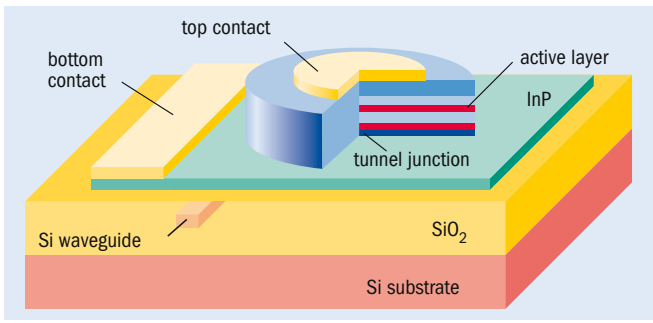

Fig. 3. (above left) A wafer bonding process unites a 200 mm CMOS silicon-on-insulator wafer with InP dies using 15 nm thick silica layers. The inset gives a magnified view of the MOCVD-grown InP dies. **Fig. 4.** (above right) Microdisk lasers have been used to provide the optical sources for the PICMOS project. These devices support whispering gallery modes and a central top contact can be added without introducing significant optical loss. **Fig. 5.** (right) An array of eight microdisk lasers, with a diameter of 7.5 μm, on top of nanophotonic silicon-wire waveguides after metallization.

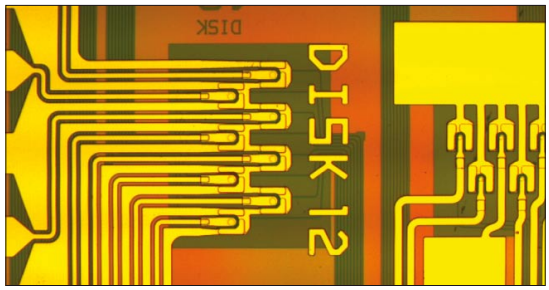

on silicon, employs an efficient die-to-wafer bonding process to transfer small III-V-based epitaxial dies onto SOI wafers. Following substrate removal, the III-V epilayers are then processed using classical wafer-scale processes (see figure 2 for details).

Our starting point for the process is the SOI wafer. A combination of deep-ultraviolet lithography and inductively coupled plasma (ICP) etching form the nanophotonic waveguiding circuits in the wafer, before SiO_2 deposition and chemical-mechanical polishing (CMP) provide planarization. We then bond InP dies containing suitable epitaxial layers epi-side down onto this SOI wafer. The sizes of these dies are application-dependent and vary from 1 mm^2 to over 1 cm^2. Since the InP dies have not been processed, their alignment tolerance is relatively large (>50 μm). This allows for a fast pick-and-place process that would be unsuitable for the transfer of preprocessed optoelectronic components. Mechanical grinding and selective wet etching then removes the InP substrate material to leave a thin layer of III-V epitaxial material on the SOI wafer. Finally, lithography, etching and metal deposition defines the lasers and detectors.

The key advantage of this approach is the wafer-scale lithography process used to align the III-V devices and the underlying silicon-wire waveguides. This approach is low-cost, but delivers an alignment accuracy of better than 0.1 μm.

Our bonding process for the transfer of InP layers to silicon is similar to wafer-to-wafer bonding processes that have already been used by many other research groups. However, in our case a full wafer-bonding process is not the most economical solution because of the high cost of InP epitaxial wafers and the significant size mismatch between commercially available InP wafers and SOI wafers, which are available in 200 mm diameters.

These drawbacks drove us to develop an alternative, efficient die-to-wafer SiO_2–SiO_2 molecular bonding process. MBE-grown wafers containing the laser structures were coated with a 15 nm silica layer using an electron cyclotron resonance deposition process and the MOCVD-grown wafers for the detectors, which had a rougher top surface, were covered with a thicker SiO_2 layer, before planarization by CMP. After both these III-V wafers were diced, the individual chips were bonded to a SOI wafer by our colleagues at CEA-LETI. The thickness and uniformity of the remaining SiO_2 layers strongly influences the operating characteristics of the final device, so TRACIT Technologies developed a specific process to optimize chip performance. Finally, the wafer-bonded dies were annealed at a low temperature to complete the transfer process (see figure 3).

Microdisk emitters

The upshot is that our process is broadly similar to that recently employed by Intel and the University of California, Santa Barbara. This US research partnership built hybrid III-V silicon lasers using a plasma-activated direct wafer-to-wafer bonding process that did not require an intermediate silica layer. However, the SiO_2–SiO_2 bonding process that we have used has the advantage of being generic: the same process and surface activation procedure can be used, independent of the wafer source.

The lasers for our optical link need to be ultra-compact, capable of high modulation speeds, produce very low threshold currents and require a power-consumption budget that is ideally lower than that needed for classical electrical paths. We have tried to address all these requirements with a microdisk that we described in a post-deadline paper at the recent Optical Fiber Communication

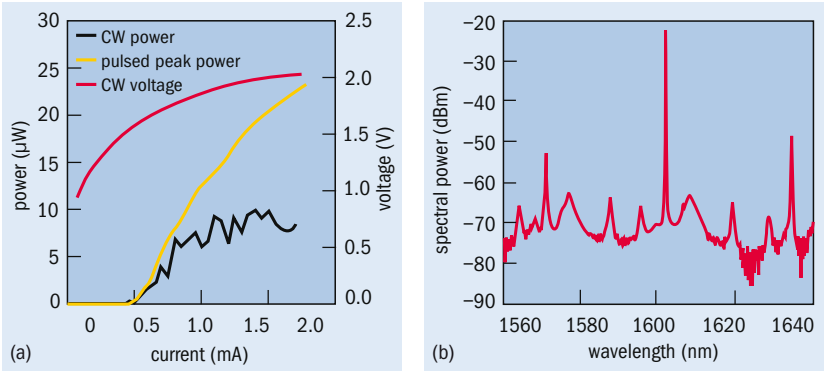

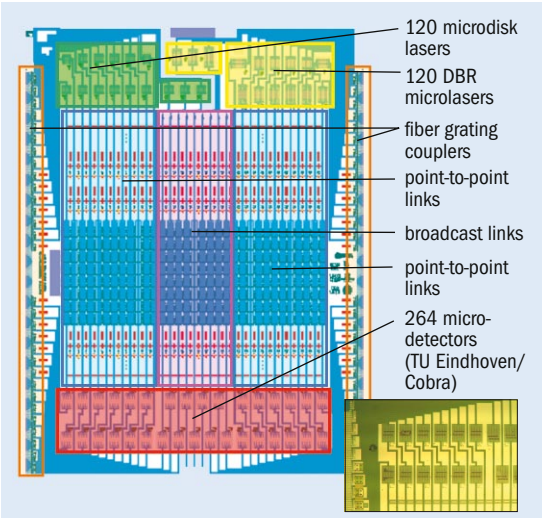

Fig. 6. (above) Performance characteristics at 20 °C for a 7.5 μm disk laser **(a)**. The continuous-wave lasing spectrum at 1.4 mA drive current, which has been normalized for the fiber-coupler efficiency and on-chip propagation loss, reveals the dominant lasing wavelength at 1.6 μm **(b)**.
Fig. 7. (right) Combining the InP lasers and detectors with silicon-wire waveguides has produced the first demonstration of a full optical link on a nanophotonic waveguide platform.

conference, which was held this March in Anaheim, CA, USA. This optical source was formed by etching the microdisk structure in a thin InP-based layer bonded on top of the silicon waveguides (see figures 4 and 5 pxx). The disk supports whispering gallery modes confined to the microdisk's edges, and a top contact can be added in the device's center without introducing additional optical losses.

Classical lasers use a heavily doped InGaAs layer as the p-side contact, but this approach would introduce unacceptable internal losses for our thin-film laser. So our partner INL developed and optimized a less-absorbing alternative structure, which contains an InGaAsP tunnel junction with a bandgap of almost 1 eV and a second n-type contact (see figure 4 for an image of the disk lasers and interconnects). Using this design our 7.5 μm diameter disk lasers produced 1.6 μm single-mode continuous wave emission at room temperature and had a threshold current of 600 μA (figure 6). Output powers up to 100 μW were produced in pulsed mode.

A comparable process has also been used to fabricate similar-sized detectors with an InGaAs-absorbing layer, which were grown by MOCVD at the Technical University Eindhoven. These devices have a 0.45 A/W efficiency, a 1.6 nA dark current, and form the final part of a full optical link.

Putting it all together

The processes for fabricating all the individual components of our photonics ICs are now established, which has put us in a position to build our first demonstrator chip, a structure containing 250 microlasers connected to a similar number of detectors with a silicon-wire waveguide array (see figure 7). Due to a processing error, the power coupled from the microsource into the waveguide was very low. However, a preliminary measurement that is still to be verified revealed that up to 60% of this coupled power was picked up by the detector. This result has provided us with the first ever demonstration of a full optical link on a nanophotonic waveguide platform, and clearly shows the potential of this technology.

Our next goals, which follow on from the recent PICMOS project, will include the integration of the completed optical layer within an actual CMOS circuit, which could be carried out using a wafer-to-wafer bonding process that is similar to those currently being developed by CMOS chip makers for three-dimensional IC stacking. In addition, we also need to perform complete fabrication of the optical layer in a CMOS fab.

We are already part way towards this goal. A CMOS pilot line and a 200 mm SOI substrate were used for both the silicon waveguide fabrication at IMEC and the die-to-wafer bonding and masking process used for the III-V mesas, which were both carried out at LETI. But the other processing steps, such as etching and metallization, were carried out in the III-V cleanrooms of IMEC, INL and Technical University Eindhoven. To refine the entire process so that it is compatible with a CMOS fab requires a better understanding of potential contamination issues and the development of a CMOS-compatible metallization process – tasks that are currently under development.

Although the PICMOS project specifically addressed on-chip optical interconnects, our electro-photonic integration platform could also serve other applications. For example, the technology could be used to unite 40 miniature disk lasers, emitting at slightly different wavelengths, with silicon-wire waveguides, and it already provides a very compact transmitter with capacity of up to 400 Gbit/s for on-chip and off-chip data links. These laser sources could also combine with compact wavelength-division-multiplexing circuits to form complex reconfigurable networks, or be used to build incredibly cheap, but very powerful optical biosensors. So the benefits of the PICMOS platform are potentially both significant and wide-ranging, and only time will tell how far this approach to building optical interconnects will permeate into tomorrow's chip technology. ●

About the author
Dries Van Thourhout (dries.vanthourhout@intec.ugent.be) was the coordinator of the PICMOS project. He is a permanent staff member of the Photonics Research Group at IMEC/Ghent University and his research interests include silicon nanophotonics, III-V/silicon heterogeneous integration and integrated photonics in general.