Silicon photonics and III-V integration

Mike Cooke reports on recent commercial activity driven by cloud and 'big data' developments, and efforts to integrate III-V light-emitting material with silicon photonic technology.

The rise of cloud and other 'big data' services has increased the need to transfer information over larger distances to link up otherwise 'stranded' computing capacity — an application for which light is a far better carrier than charge.

In free space, light has carried information from almost the time of the big bang, the beginning of time itself. On planet Earth, data carried by photons through fibers and other materials does require some refreshment, but far less often than would be needed by electronic-based carriage.

The key challenge is to integrate optical data transfer with data processing electronics. In particular, emitters and detectors are the vital interface between silicon electronics and optical systems.

Silicon photonics offers prospects for such integration and a number of companies already have products sampling and in 'mass production'. However, the integration is somewhat incomplete in the area of providing laser or some other light power into the system. This is usually piped in from outside using fibers, although there are increased possibilities for integrating laser chips more closely into the silicon photonics using wafer- or chip-bonding. However, these techniques suffer from problems such as implementing precise alignment, which adds complexity and expense.

The problem is that the III-V compound semiconductor materials favored for light emission (and detection) don't mix well with silicon. Challenges include mismatching of the lattice structure and poisoning of silicon electronics since the III and V groups of the periodic table are the ones used for doping silicon for p- and n-type conduction, respectively.

Direct integration of III-V materials on silicon should reduce process complexity and cost, but lattice mismatching reduces material quality and reduces the possibilities for efficient light emission, particularly for lasers. Although there have been few explicit reports of cross-fertilization up to now, there could in the future be transfer of expertise between those working on photonics and III-V high-mobility transistors, where similar integration problems have been raised.

While direct growth of III-V materials on silicon would be desirable, wafer bonding is often more practical at the present stage of technical development.

The light used by silicon photonics is generally in the infrared region, being restricted by the silicon transparency window and energy bandgap to wavelengths above 1.1μm. This window covers the common optical communications wavelengths around 1.55μm.

The substrates for such work are commonly silicon on insulator. The high refractive index of 3.5 optically confines light by total internal reflection, allowing microscopic waveguides to be created.

Waveguide paths for infrared light need less extreme pattern definition of the order of a micron rather than the (tens of) nanometers of bleeding-edge electronics. This could give new life to processing facilities that would otherwise be due for running down and eventual closure.

Intel says that it has been developing silicon photonics for 16 years — in other words the inception came at the crescendo of the dotcom boom–bust and the subsequent cold dark years for optical networking development.

Prospects for silicon photonic optoelectronic integration has been improving for a while and Intel now sees opportunities in the rise of cloud computing and related developments. The company is presently promoting its small-form-factor 100G PSM4 (parallel single-mode fiber 4-lane) QSFP28 optical transceiver (Figure 1). The stated aim is applications in large-scale

Figure 1. Intel Silicon Photonics 100G PSM4 QSFP28 optical transceiver.

cloud and enterprise data centers, along with Ethernet switch, router, and client-side telecom interfaces. The specification is for 100 Gigabit [per second] Ethernet (100GbE) optical links over single-mode fiber.

Originally due in 2015, but delayed due to unspecified difficulties, Intel now says that the product is "shipping in volume". The transceiver incorporates Intel's hybrid laser technology that allows greater than 90% coupling efficiency. Combined with high-density optical interconnects and Intel's "unmatched manufacturing capabilities", the technology should deliver scaling and cost benefits.

The reach of the device is up to 2km, beyond that required by the PSM4 multi-source agreement optical interface specification (www.psm4.org). The operating temperature range is 0ºC to 70ºC. Power consumption is rated at 3.5W maximum.

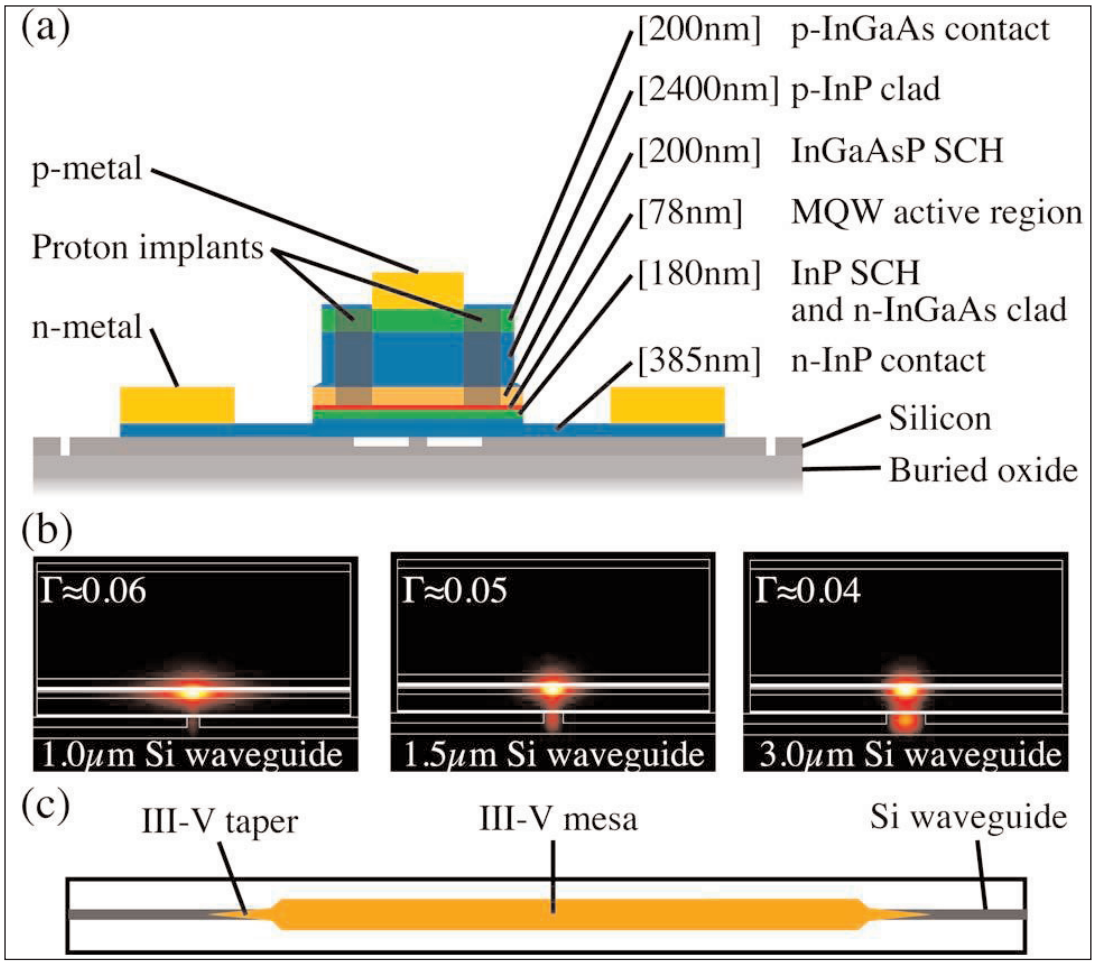

Figure 2. (a) Cross-sectional schematic of UCSB hybrid silicon active region with approximate layer thicknesses. (b) Simulated optical mode profiles and estimated MQW optical confinement factor of hybrid silicon active region for various silicon waveguide widths. (c) Top-view schematic of hybrid laser design. From [Alexander Spott et al, Optics Letters, vol40, p1480, 2015].

Beyond this specific device, Intel plans to move from pluggable to near-future embedded products at 400G, and towards integrating optical interconnects directly into switch and server electronics. Intel's roadmap also includes integrated optical application specific circuits (ASICs).

Intel has developed its hybrid silicon laser technology in partnership with University of California Santa Barbara (UCSB) [e.g. Alexander W. Fang et al, Optics Express, vol14, p9203, 2006]. The 2006 device consisted of aluminium gallium indium arsenide (AlGaInAs) quantum well epitaxial material bonded to a low-loss silicon strip waveguide (Figure 2). The cavity and contact formation was performed after bonding. The laser optical mode coupled evanescently into the silicon waveguide beneath.

UCSB continues to work on silicon photonics (http://optoelectronics.ece.ucsb.edu/publications/3), with recent work including further research on room-temperature hybrid silicon lasers, along with integrating quantum cascade lasers on silicon and more purely photonics engineering.

IBM has also been working towards integrating III-V laser and silicon photonics technology, likewise seeing cloud and 'big data' applications. In 2015, company researchers presented fully integrated wavelength-multiplexed silicon photonics chips that the company said would enable manufacturing of 100Gb/s optical transceivers. The technology links sub-100nm CMOS electronics with various optical components. Light signals are transmitted at four infrared wavelengths used for optical communications.

While the laser power for these devices seems to be fed in from outside, the company has also developed a self-aligned flip-chip assembly to couple III-V lasers into silicon photonic circuits. The solder pads of the silicon photonics and laser chips are purposely offset so that, during annealing, surface tension pulls the chip into alignment with precisely designed mechanical stops (Figure 3). The optical connections are butt-coupled and the electronic connection is through the solder pads. However, presented work seems to be at the level of demonstrations using silicon 'mock-up InP dies' in place of actual laser chips.

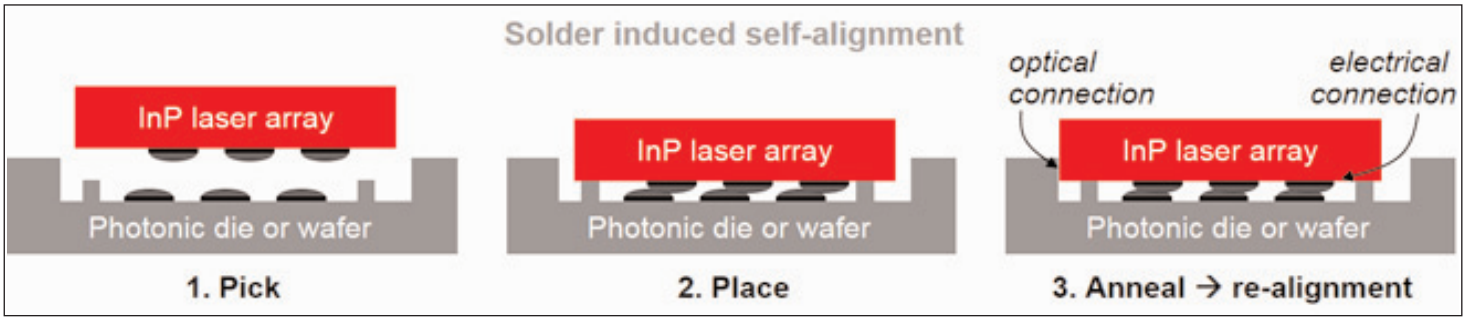

Figure 3. IBM's solder-induced self-alignment.

IBM says that it has been leading silicon photonics research for more than a decade with a series milestones starting in 2006.

In March 2016, Kaiam Corporation demonstrated what it claimed as the world's first 100Gb/s CWDM4 silicon photonics transceiver at the OFC optical networking and communication conference/exhibition. Also, the company announced sampling of devices in September 2015. Production is based at the company's facility in Livingstone, Scotland, UK.

According to Kaiam, the module is able to transmit signals through 10km of single-mode fiber. The silicon photonics component combines electronics (modulator driver, transimpedance amplifier, and clock & data recovery circuits) and silicon optical modulators and detectors in a single 3D chip stack. Wavelength multiplexing and demultiplexing is realized by glass-based planar lightwave circuits (PLCs). The continuous wave optical input is generated by indium phosphide lasers with micro-electro-mechanical coupling alignment into the PLCs. Grating couplers connect the glass PLCs to the silicon photonics IC.

InGaAs well ridges on 300mm silicon

A recent example of direct integration of III-V materials on silicon comes from Imec and Ghent University in Belgium. This team used aspect ratio trapping (ART) techniques to produce indium gallium arsenide (InGaAs) multiple quantum wells (MQWs) on 300mm-diameter silicon in a ridge format that could be used in future laser diodes [B. Kunert et al, Appl. Phys. Lett., vol109, p091101, 2016].

ART techniques have been developed to reduce defect density in III-V structures on silicon and Imec has used ART in its work on high-mobility transistors with a view to future high-speed/low-power electronics.

The Imec/Ghent team now reports: "Despite the narrow trench dimension, the grown out III/V ridge has reached a volume suitable for wave guidance,". The researchers add that they are investigating "optimized nano ridge structures with an increased GaAs volume and a confining InGaP cap layer to reduce carrier losses due to surface recombination". They expect realization of first nano ridge laser diodes using such techniques in the near future.

The III-V ridges were grown by metal-organic vapor phase epitaxy on 300mm-diameter silicon with exact (001) crystal orientation. The trenches were provided by a shallow trench isolation (STI) process used in mainstream silicon electronics, and were 10μm long with varying widths. Trenches of the same width were arranged in 600μm x 600μm fields. The trench areas were about 10% of the field. V-facets with (111) surfaces were formed by wet etching the silicon in the trench with tetramethyl-ammonium hydroxide.

The GaAs deposition started with nucleation at 360ºC using a triethyl-gallium precursor. The main growth inside and outside the trench was carried out at 580ºC with trimethyl-gallium precursor. InGaAs layers for the quantum wells were also grown at 580ºC.

The low-temperature nucleation/seeding was used to encourage two-dimensional growth, avoiding planar defects (stacking fault, twins...). The high-temperature growth gives a higher-quality GaAs crystal with misfit defect nucleation and threading dislocation gliding for full strain relaxation. Defects are trapped by the silicon dioxide trench walls. The researchers comment: "No indication of any misfit and threading dislocations was found above the STI level or close to the MQW region. Only three planar defects such as micro twins are visible in a 5.5μm-long transverse section."

The researchers add: "It is of key importance to avoid the formation of stacking faults and twins from the beginning by an optimized seed layer and sufficient pre-cleaning of the Si surface, whereas the nucleation of misfit- and threading dislocations should be initiated as deep as possible in the trench, close to the silicon surface, to benefit from a high aspect ratio."

The GaAs above the trench grew in a box-shape cross-section. The boxes were rectangular for both narrow and wide trenches. The researchers report that "the ridge courses are very straight and uniform with flat surfaces in all directions, which is essential for realizing a waveguide with low light scattering losses."

Reciprocal-space mapping using x-rays indicated that the GaAs ridges were fully relaxed. The InGaAs quantum wells were grown pseudomorphically on the GaAs buffer with ~19% indium content (Figure 4).

From bottom to top, the well thicknesses were 9.9nm, 8.9nm and 7.7nm. The periods of time for growth of the three wells was the same, as also for the three GaAs barrier layers. Although the growth is mainly on the top surface, there is some side-wall deposition.

The researchers comment: "Side-wall deposition leads to a very good lateral carrier confinement, reducing non-radiative recombination channels at the QW side surfaces and is therefore desirable. Especially, the evolving decrease in InGaAs thickness at the lateral ends of the QWs will lead to higher quantized electronic states, which induces an injection and repulsing of electrons and holes in the InGaAs layer at the box edges towards the centered QW region with constant layer thickness and lower energetic states. The InGaAs layer of less than 2nm thickness on the {110} side walls is too thin to hold bonded states with an efficient radiative recombination rate."

In photoluminescence (PL, Figure 5), a peak was seen between 1000nm and 1050nm, except for the very widest trench of 500nm. The photo-emission from the QWs was orders of magnitude higher than from the silicon substrate, which was centered around 1140nm. Some of the spectra show a shoulder around 980nm, which could be due to higher quantized states of the wells. The brightest peak was for 60nm-wide trenches. However, comparisons are complicated by factors such as ridge volume, carrier density, surface states, and so on.

The team comments: "Although it is difficult to judge about the overall impact of these uncertainties, it is very likely that the decrease in PL intensity towards wide trenches is linked to an increase in TD density. The aspect ratio of trenches wider than 100nm is quite low, leading to an insufficient defect trapping. Hence, the remaining defect density in the ridge material of wider trenches could cause high carrier losses. The intensity decline towards narrow trenches could be related to a disproportional rise in ridge surface area or less pump light absorption."

The differences in peak position are likely due to changes in indium content and/or QW width. The photoluminescence linewidth was 38meV for material grown from 100nm-wide trenches.

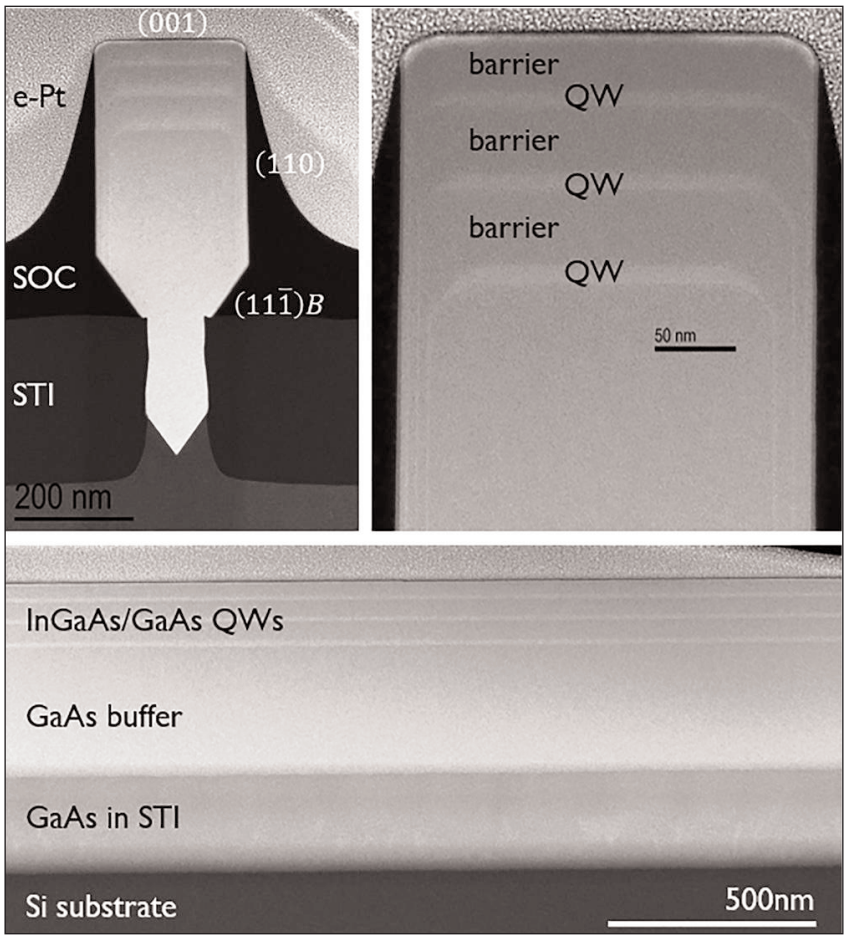

Figure 4. High-angle annular dark-field scanning transmission electron microscopy of box-shaped GaAs ridge with three InGaAs QWs. Top left: Cross-section of full ridge. Top right: Higher magnification of QW region. Bottom: Transverse section along trench. The GaAs inside STI shows slightly different contrast because some STI oxide was cut out during sample preparation.

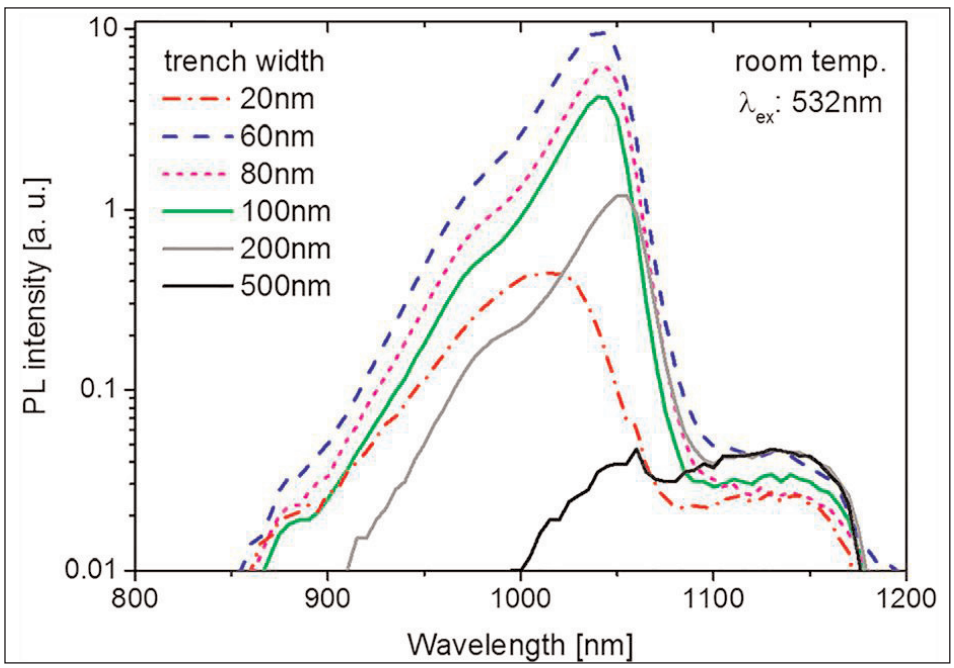

Figure 5. Room-temperature PL spectra from different trench widths in log-scale.

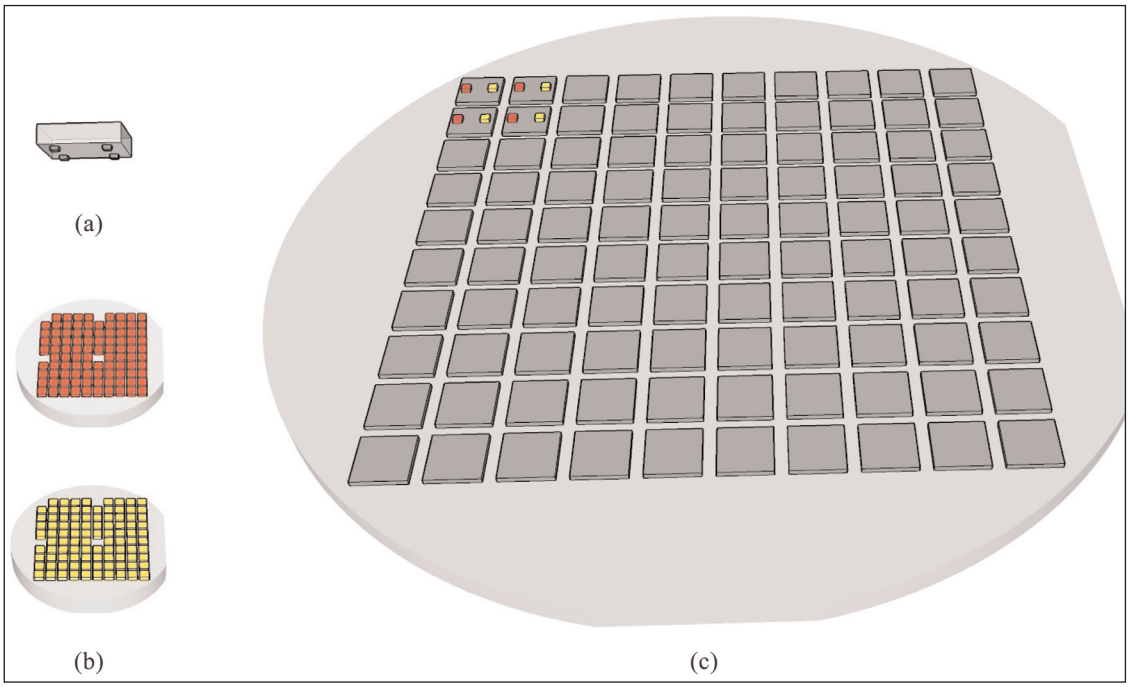

Figure 6. Area magnification in transfer printing III-V coupons from III-V source substrate to SOI target: (a) patterned stamp, (b) two source substrates with patterned coupons, and (c) SOI target substrate with four printed coupons of each source, top left.

Transfer printing

Ghent University/Imec in Belgium have also joined with X-Celeprint, Ireland/USA, to claim "the first III-V opto-electronic components transfer printed on and coupled to a silicon photonic integrated circuit". [Andreas De Groote et al, Optics Express, vol24, p13754, 2016]. "This work forms a stepping stone towards the cost-effective integration of III-V optoelectronic components onto silicon photonic integrated circuits, including lasers, semiconductor optical amplifiers and electro-absorption modulators," the team adds.

The transfer method uses a polydimethylsiloxane (PDMS) elastomer stamp to pick up and place 'coupons' from a source to target in parallel. The coupons consist of thin-film materials or even completed devices. The adhesion of the materials to the stamp is determined by the speed of contact. A fast speed picks up the coupon and a slow speed releases it. By patterning the stamp, numbers of coupons can be picked up and placed in a magnified fashion on the target substrate (Figure 6). This allows efficient use and co-integration of different expensive III-V materials and device sources.

The researchers comment: "While showing similarities with a pick-and-place technique, the main advantage of the transfer printing approach is that coupons can be transfer printed in a massively parallel way, by picking up and putting down large arrays of coupons at the same time. This allows for a high-throughput integration process."

The transferred coupons are ten to hundreds of microns in size, unlike in the usual wafer bonding methods where the dies are usually of the order of millimeters.

The coupons are transferred in a wafer-scale process, making alignment on SOI waveguide circuits non-critical, according to the researchers. Even so, with markers a three-standard-deviation (3σ) alignment accuracy of $\pm 1.5\mu m$ can be achieved.

The researchers transferred an indium gallium arsenide phosphide (InGaAsP) multi-quantum well (MQW) light-emitting device onto a silicon-on-insulator (SOI) chip with a processed waveguide circuit (Figure 7). The MQW was sandwiched between 60nm InP barriers. The total membrane thickness was 200nm.

The III-V material for the light-emitting device was grown on InP substrate. The sacrificial release layer was 1μm InGaAs. The top sacrificial layer was 100nm InGaAs, followed by 1μm InP. These top layers were designed to prevent buckling of the membrane when adhered to the transfer stamp.

The coupons were defined by a double mesa etch. A photoresist layer was used to protect the transfer layers while the release layer was under-etched. The coupons were transferred to the SOI wafer, which was coated with divinylsiloxane-bis-benzocyclobutene (DVS-BCB) polymer as an adhesive. The DVS-BCB was cured at 250ºC for an hour. After the coupon was transferred and fixed, the photoresist was removed by oxygen plasma and the top sacrificial layers by a wet etch.

Different etch chemicals were tested for the release and finishing etch processes. The researchers comment: "The process should render a sufficiently strong bond to withstand strong acids and make sure that the III-V waveguide is close enough to the silicon waveguide in order to allow optical coupling through tapers."

The best release etch was found to be with cold iron chloride that delivered both full under-cutting of the layers and high selectivity against InP etch by a factor of more than 500. The wet etch took two hours to release the 40μm-wide coupons. The top sacrificial layers were removed with hydrochloric acid for InP and sulfuric acid mixed with hydrogen peroxide for the InGaAs.

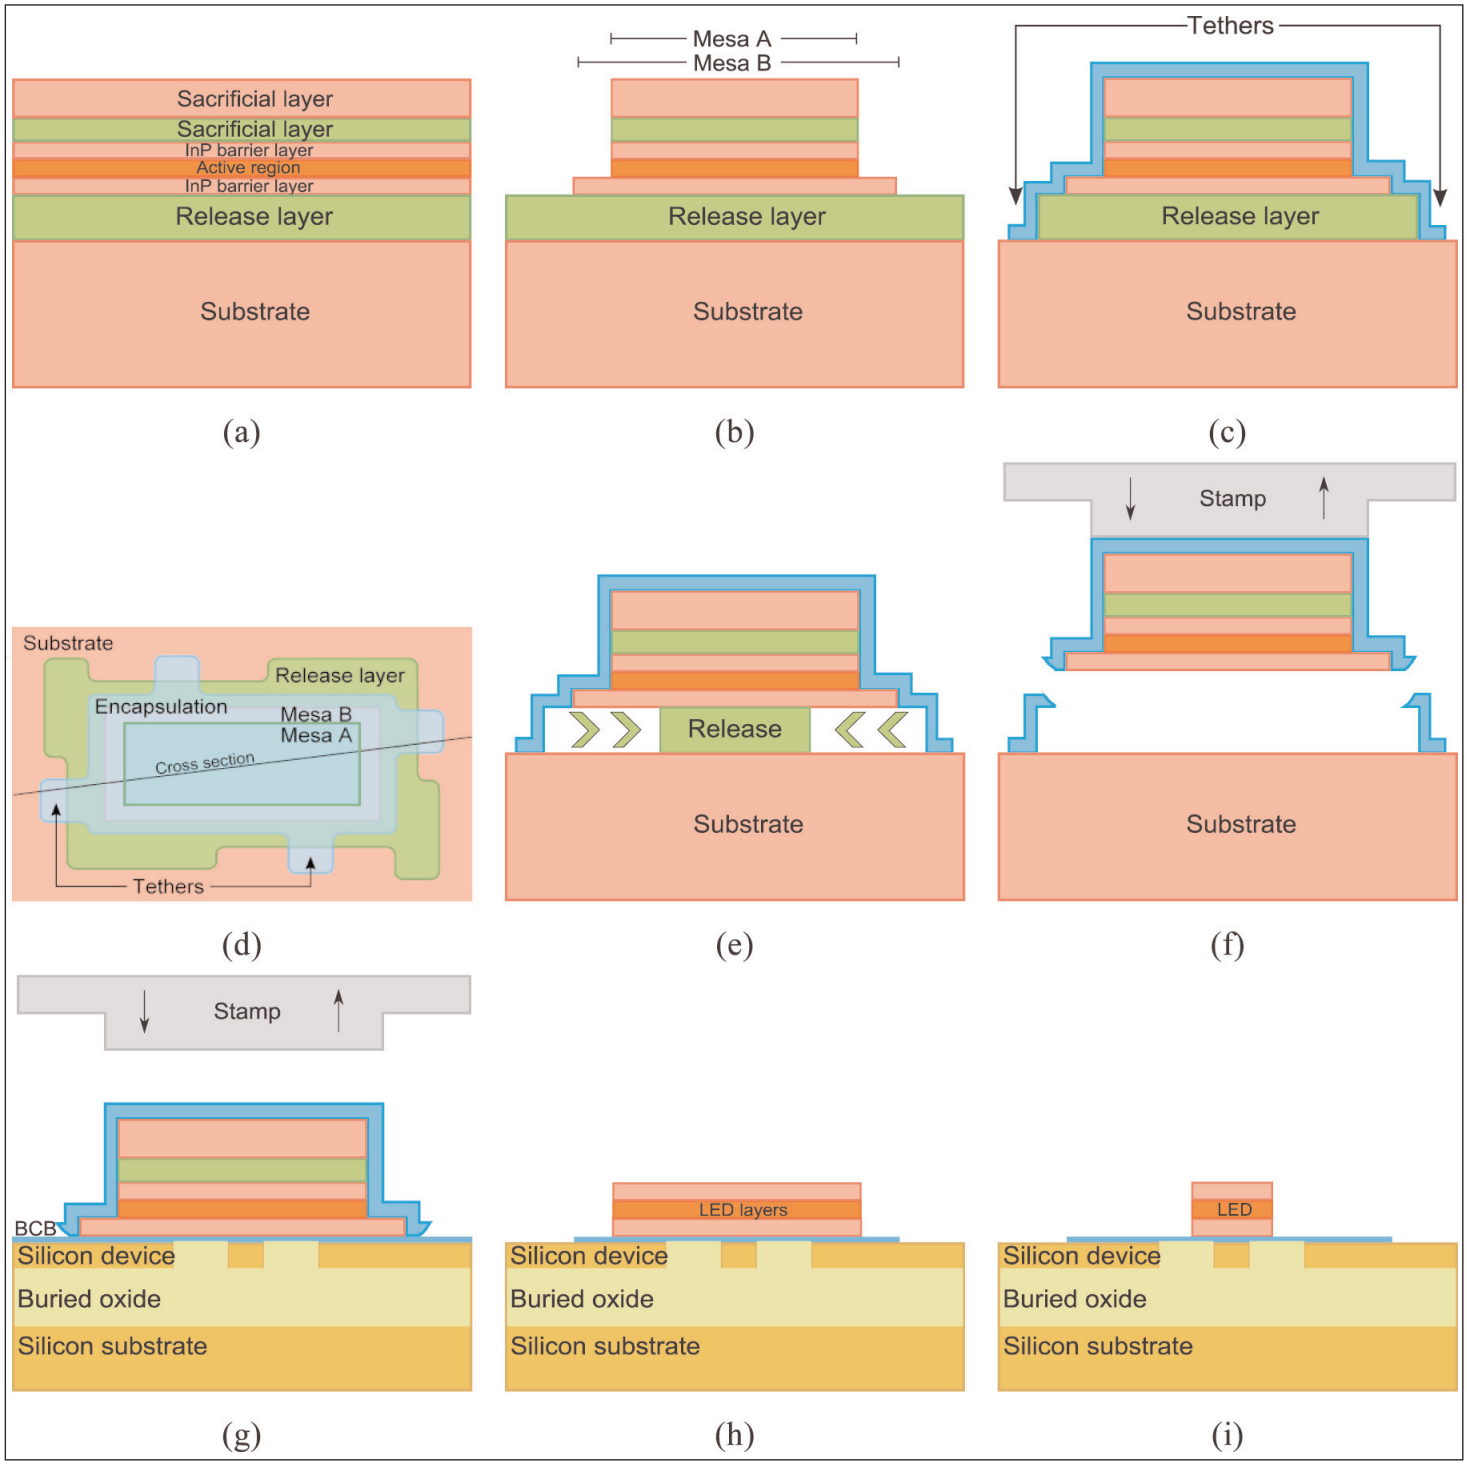

Figure 7. Process flow of transfer printed, optically pumped light-emitting device: (a) InP starting layer stack, (b) coupon patterning, (c) encapsulation and tether definition, (d) top view after encapsulation and tether definition. Further processing includes: (e) release etch, (f) pick-up of III-V coupon from source, (g) printing of coupon to SOI target, (h) removal of encapsulation and of sacrificial layers, (i) definition of light-emitting device.

The proof-of-principle setup consisted of a single-spatial-mode broadband light-emitting device integrated with SOI waveguides processed on 200mm wafers at Imec's CMOS pilot line. In this case, the III-V device was pumped with optical power from a 1310nm laser directed through the SOI waveguide system.

Measurements on the integrated III-V/SOI waveguide structure indicate a power efficiency of the emitted light at around 1550nm wavelength at the same order of magnitude as achieved with traditional III-V-on-silicon bonding. "The 3dB bandwidth of 130nm is large, and comparable as well to previously reported results," the team writes. ■

The author Mike Cooke is a freelance technology journalist who has worked in the semiconductor and advanced technology sectors since 1997.