Enabling silicon photonics through advances in III-V integration on silicon

EV Group's Dr Martin Eibelhuber discusses a wafer-level die transfer process for bonding InP laser dies to a silicon photonics wafer, allowing volume production.

Storing and processing ever increasing volumes of data at higher speeds and lower costs while decreasing energy consumption are constant goals in our communications age. Copper interconnects continue to be the standard material to enable data transport between chips, boards or racks in data centers. However, conduction-induced power losses as well as signal latency with copper are becoming more challenging to overcome, and limit increases in data communication bandwidth.

Optical data transfer — already well established for long-haul and metro networks — is able to keep up with the data transfer rates required by data centers. Silicon photonic chips designed for today's data centers have already proven transmission rates up to 40Gbit/s for a single laser [1], which can be scaled up to more than 1 terabit (Tbit)/s by multiplexing several wavelengths solely using a single optical fiber. First products on the market are already announced with transfer rates around 100Gbit/s. Thus, optical interconnects and, in particular, silicon photonics is a novel and most promising way to overcome the limitations in short-range communication. From a technological standpoint, silicon photonics combines the best of two worlds — namely the maturity of silicon (Si) technology and the superior optical properties of III-V compounds — thus enabling the integration of a fiber-optic communication platform into a single chip.

The idea of combining silicon and CMOS processing with optics fabrication was initially not seriously considered due to the indirect-bandgap properties of silicon, which inhibit efficient optical performance. Early on, it was only known that silicon has a high refractive index and is transparent at infrared wavelengths, and therefore could perform optimally as a waveguide for the desired telecommunication wavelengths of 1.3µm and 1.55µm. Subsequent years of development in CMOS-optics technology resulted in tremendous progress, and today almost all necessary elements — such as couplers, modulators and detectors — are based on the CMOS infrastructure. An overview of the performance of optical components based on different materials used in their manufacture is given in Figure 1.

Significant progress has also been made for the light source, and silicon-based lasing — once thought almost impossible to achieve — has since been proven by a Raman laser [2]. However, the performance of silicon-based lasers is still currently far behind that of III-V lasers and much work is needed to narrow the gap. As a result, today's silicon photonics still rely on III-V laser structures.

In addition, progress in epitaxial growth and new wafer bonding techniques are removing many of the restrictions of using different materials on a single chip. It has always been a desire to integrate III-V materials or other compound semiconductors on silicon in order to benefit from their different material properties.

Improvements in direct epitaxial growth on silicon have already demonstrated that the growth of high-quality crystal structures on large wafers is feasible, thus proving the concept of epitaxial heterogeneous integration of III-V materials on silicon. However, crystallo-

	Laser	Photo-detector	Modulator	Passive Devices	Wafer Level Packaging
GaAs	Good	Good	Medium	Low	X
InP	Good	Good	Medium	Low	X
$LiNbO_3$	X	X	Good	Low	X
Silicon	Poor	Medium	Low	Good	Good
Heterogeneous Integration of InP	Good	Good	Medium	Good	Good

Figure 1: Performance of optical elements related to the used material.

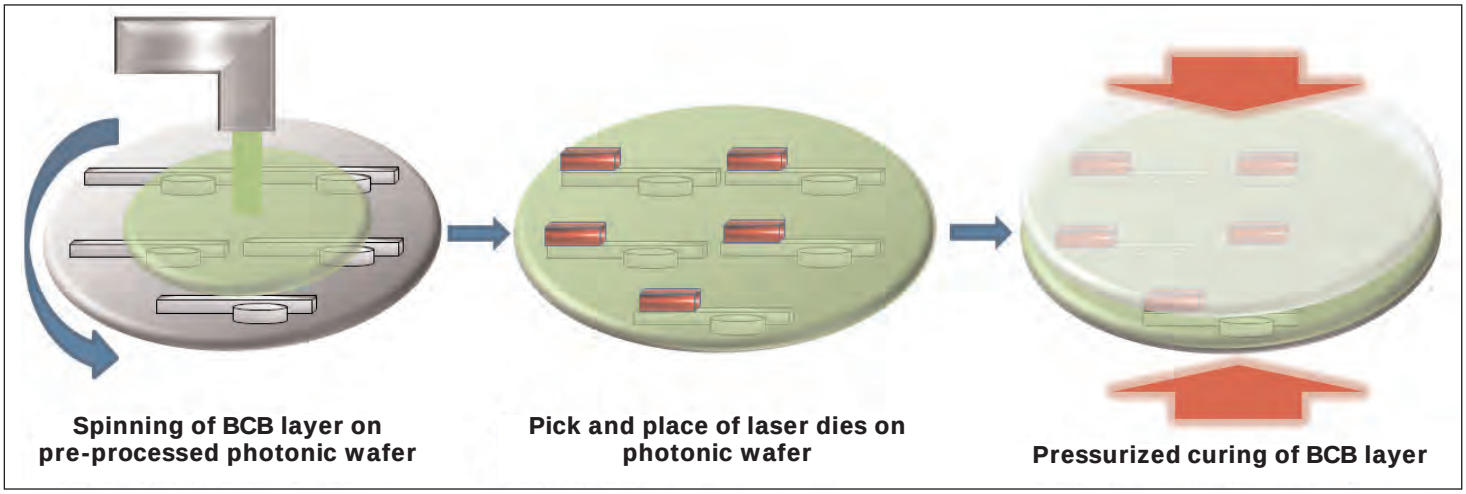

Figure 2: Process flow for adhesive bonding of vertically coupled laser dies.

graphic defects arise within the interfaces between the silicon substrate and III-V compounds as a result of the differences in lattice constant and coefficient of thermal expansion (CTE) between the silicon and III-V materials. A defect-free crystal structure can only be re-established after inserting a thick buffer layer between the materials. Efforts to reduce the thickness of this layer (such as the inclusion of dislocation filter layers) still result in a layer thickness of several microns. This is a major drawback for silicon photonics, since the most widely accepted approaches involve the vertical coupling of light to the silicon waveguide with the laser die placed directly on top of the waveguide [3] (see Figures 2 and 3). This vertical coupling requires a superior optical interface with the active region no more than a few tens of nanometers away from the silicon — at least an order of magnitude smaller than the thickness of the required buffer layer.

In contrast, wafer bonding enables optimal interfaces with few or no crystallographic defects as the materials are joined with a transparent interlayer that mitigates the effects of the lattice constant and CTE mismatch. The two leading methods of wafer bonding for this purpose are either bonding using a transparent adhesive or direct molecular bonding.

As shown in Figure 2, bonding with an adhesive has the advantage of being very simple and easy to implement. The process flow includes only a coating process to flatten the topography of the pre-processed photonic wafer, placement of the laser dies, and finally a wafer-level bonding step with thermal curing. UV curing is not an option since neither silicon nor III-V materials are transparent to UV light. In order to keep the distance between the materials constant, superior pressure and temperature uniformity are crucial to guarantee the best performance of the devices. Thus, the top bond chuck has to comply efficiently for die or wafer inhomogeneity while putting the same pressure on the single dies. Curing the adhesive at the die level carries an enormous disadvantage of requiring heating and cooling for every single die resulting in very long process times. Current roadmaps already highlight the need to economically process several thousand laser diodes on a single photonic wafer. This fast integration of the whole photonic wafer can only be enabled by wafer-level bonding.

However, even though direct bonding cannot compete with the simplicity of adhesive bonding, it is typically the preferred solution, as it enables better device performance. Despite losses within the adhesive material, the main drawback of adhesive bonding is the distance between the laser dies and the waveguide, which is the key parameter for the light coupling efficiency between the laser and the waveguide. Direct wafer bonding allows two materials to be joined with an oxide interlayer that is only a few nanometers in thickness, resulting in a higher coupling efficiency that is not achieved with adhesive interlayers.

Direct fusion bonding of indium phosphide (InP) laser diodes on silicon is challenging in several aspects. First, the topography of the pre-processed photonic wafer has to be flattened by SiO_2 deposition and chemical mechanical planarization (CMP) in order to ensure the required low surface roughness. Fusion bonding is a two-step process consisting of a room-temperature pre-bond and an annealing step. Traditional annealing processes developed for silicon-on-insulator (SOI) wafer manufacturing required an annealing temperature of 1100°C, which is far too high for CMOS and III-V technology. In particular, for bonding InP on silicon, the bond annealing temperature should not exceed 350°C, as the CTE differs significantly between these materials.

Despite these temperature restrictions, sufficient bond strength must be obtained. Low-temperature plasma-activated direct wafer bonding is a process that lowers the required annealing temperatures in order to reach high bond strength. Plasma activation of both wafer surfaces is used to modify the surface chemistry of both materials, thereby significantly reducing the annealing temperature requirements. In this way,

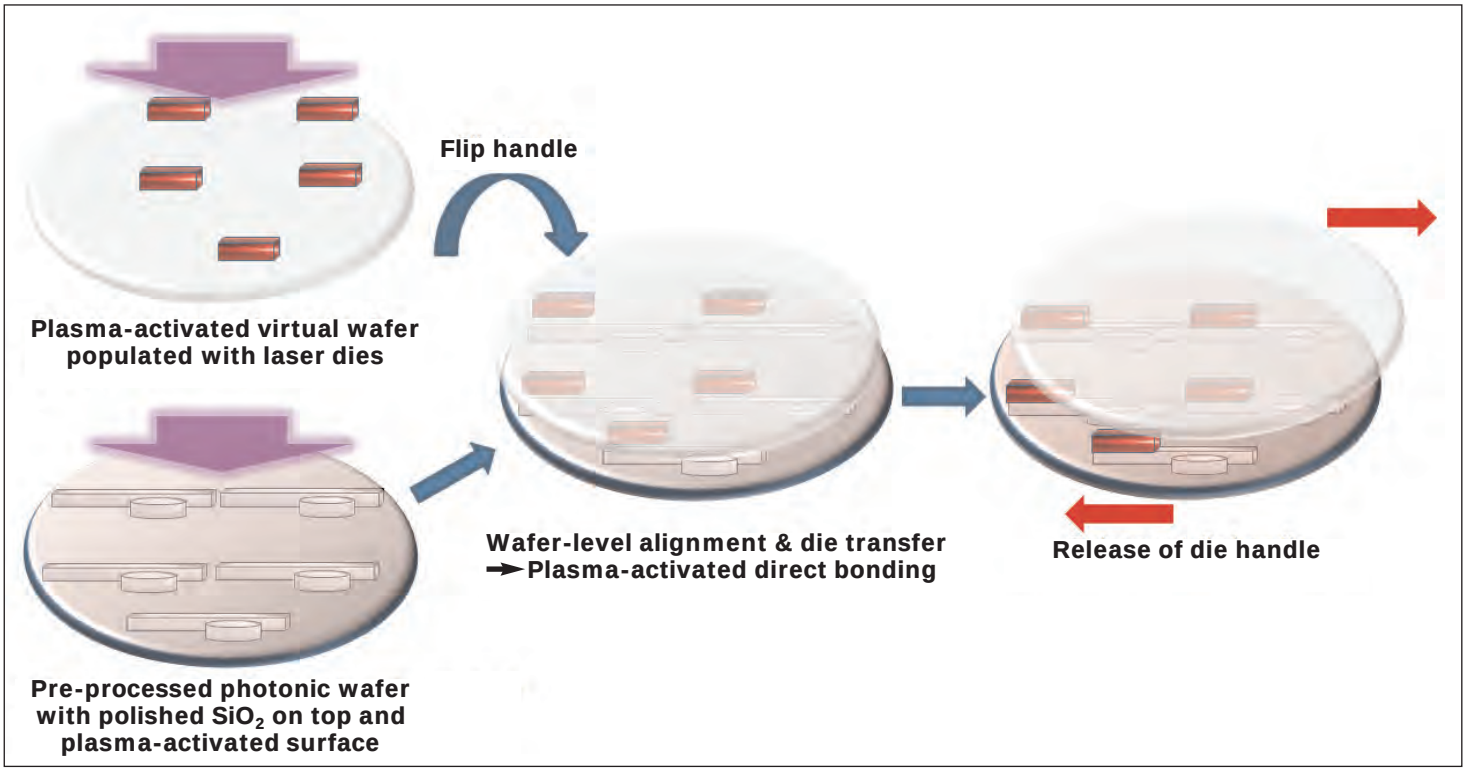

Figure 3: Wafer-level die transfer process for plasma-activated direct bonding of vertical-coupled laser dies.

materials supporting a high crystal quality of compound semiconductors can be joined with a carrier that accounts for differences in thermal expansion. Annealing temperatures below 350°C have been demonstrated using this approach for bonding InP on silicon wafers at diameters up to 150mm [4].

Since heterogeneous integration of InP is needed only on a minor part of the photonic wafer, bonding full wafers to each other is economically less attractive than placing the laser dies only in the needed areas. A die-level process enables one to overcome geometric constraints with InP wafers, since InP — which is currently only available on substrates up to 150mm in diameter — can then be scaled up to 200mm and 300mm production lines. Nonetheless, reasonable process times can only be achieved with a wafer-level process.

One way to overcome these challenges, as depicted in Figure 3, is to selectively place the laser dies on a handle wafer, creating a virtual wafer, transfer them in a single wafer-to-wafer bonding step, and finally release the handle [5]. This process flow enables dies to be placed on the handle wafer very quickly, while more time-consuming processes — such as cleaning, plasma activation and bonding at elevated temperatures — can be applied at the wafer level. In particular, cleaning of the dies is very important, since direct bonding is very sensitive to particle contamination, which can inhibit direct contact of the surfaces. As a result, this approach enables integration of III-V laser diodes directly bonded onto silicon-based photonic circuits with shorter process times and without wasting unreasonable amounts of real estate.

An alternative approach to get the light into the waveguide is edge coupling. In this case the laser diodes are placed at the same level as the photonic structures by etching pockets into the SOI substrate where the dies are placed. As a result, the facet of the laser points directly to the edge of the waveguide, which allows more power to enter the photonic circuit via a tapered waveguide. For this approach the bond interface does not need to be transparent. Hence, a eutectic bond is the preferred solution, as it allows the landing pad to be used as a bottom contact. Since the process flow of the wafer-level die transfer in Figure 3 can be applied to almost any kind of bond interface, it can also be adopted to realize edge-coupled photonic structures at the wafer level, as shown in Figure 4. For eutectic bonding, creating a virtual wafer enables one to remove undesirable oxides on the dies at the wafer level. Furthermore, the required heating and cooling ramps for eutectic bonding are far too time consuming to populate a photonic wafer with the laser dies in a reasonable time scale by serially bonding each chip to a CMOS wafer. Fast population of the handle wafer and aligned wafer-level bonding of the laser dies on the landing pads can solve this issue while still ensuring high transfer rates.

In conclusion, heterogeneous integration of InP laser diodes on silicon for silicon photonic circuits has received much attention recently. Despite the concentrated efforts to directly grow III-Vs onto silicon, dislocations are hard to control and the desired interface qualities have not been achievable. Wafer-to-wafer bonding offers an ideal solution for this problem. Nonetheless, the bonding of fully processed device wafers results in

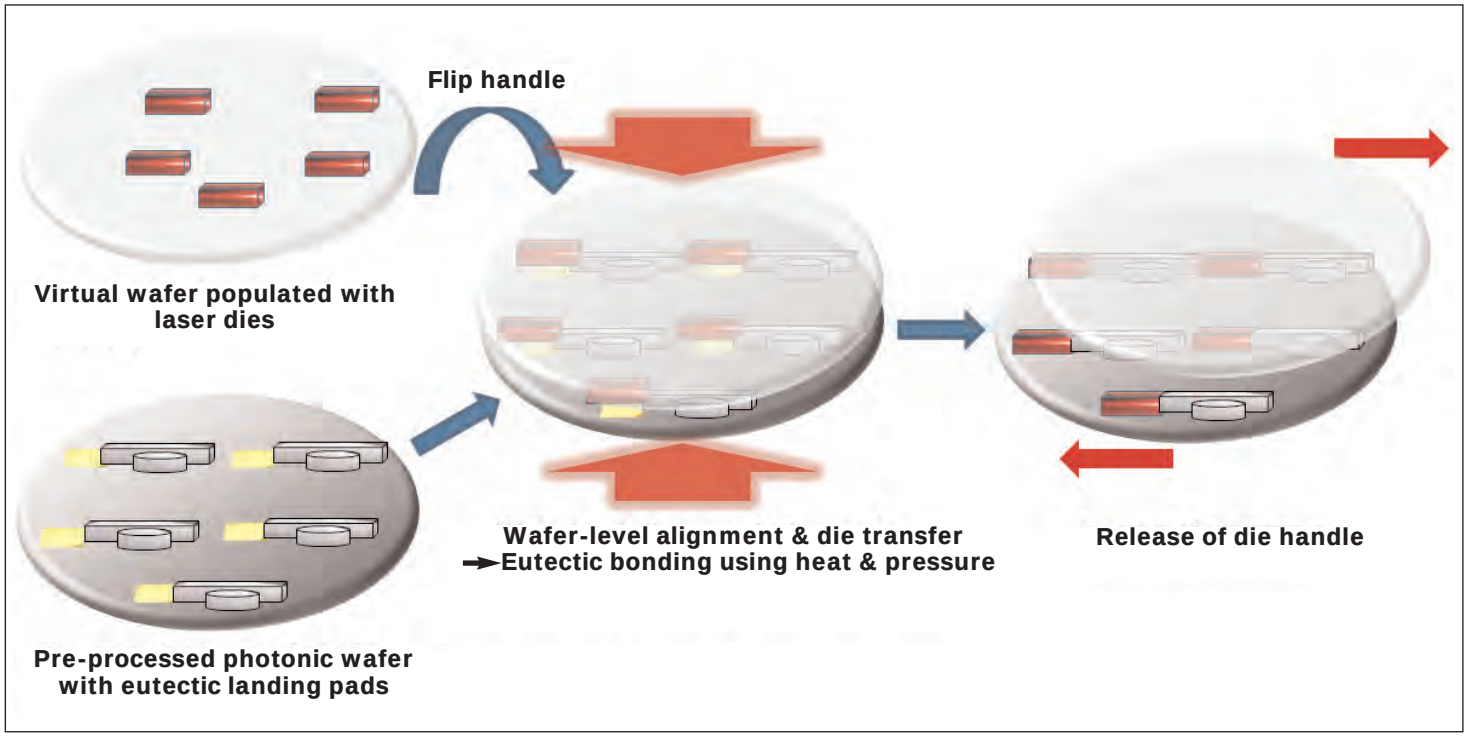

Figure 4: Wafer-level die transfer process for eutectic bonding of edge-coupled laser dies.

significant waste and loss of real estate, since only a low filling factor of the laser dies is needed. Furthermore, the usage of 200mm and 300mm CMOS production lines for InP wafers — as well as the integration of laser dies with different wavelengths in order to enable data transfer in both dedicated telecommunication bands around 1.3 and 1.55µm — would significantly aggravate this waste and lead to higher cost of ownership. Thus, in all integration schemes — be it adhesive, direct or eutectic bonding — large numbers of laser dies have to be implemented on a single photonic wafer. The key to enable reasonable time scales and to move silicon photonics into mainstream production is to implement wafer-level die transfer processes.

These advanced bonding techniques for photonic integrated circuits or heterogeneous integration in general can add tremendous benefit. The parallel processing of bonding a virtual wafer of InP die to the photonic wafer offers a substantial increase in throughput compared to serially bonding each InP die to the photonic wafer. It also solves the geometry issue, CTE mismatch, and substrate size mismatch associated with bonding an InP wafer to the photonic wafer. ■

Author: Dr Martin Eibelhuber, Business Development Manager, EV Group, E-mail: m.eibelhuber@evgroup.com
www.evgroup.com

References

[1]Y. Vlasov, W. M. Green and F. Xia, 'High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks', Nature Photonics, vol. 2, no. 4, 242–246, 2008
[2]H. Rong, A. Liu, R. Jones, O. Cohen, D. Hak, R. Nicolaescu, A. Fang and M. Paniccia, 'An all-silicon Raman laser', Nature 433, 292–294, 2005
[3]H. Park, A. Fang, S. Kodama and J. Bowers, 'Hybrid silicon evanescent laser fabricated with a silicon waveguide and III-V offset quantum wells', Opt. Express 13, 9460–9464, 2005
[4]D. Liang and J. E. Bowers, 'Photonic integration: Si or InP substrates?' Electronics Letters, vol. 45, no.12, p578, 581, June 4 2009
[5]T. Plum, 'Chip-on-Wafer Bonding — A new Fabrication Technology for High-Power Semiconductors', Annual Report RWTH Aachen, 24, 2006

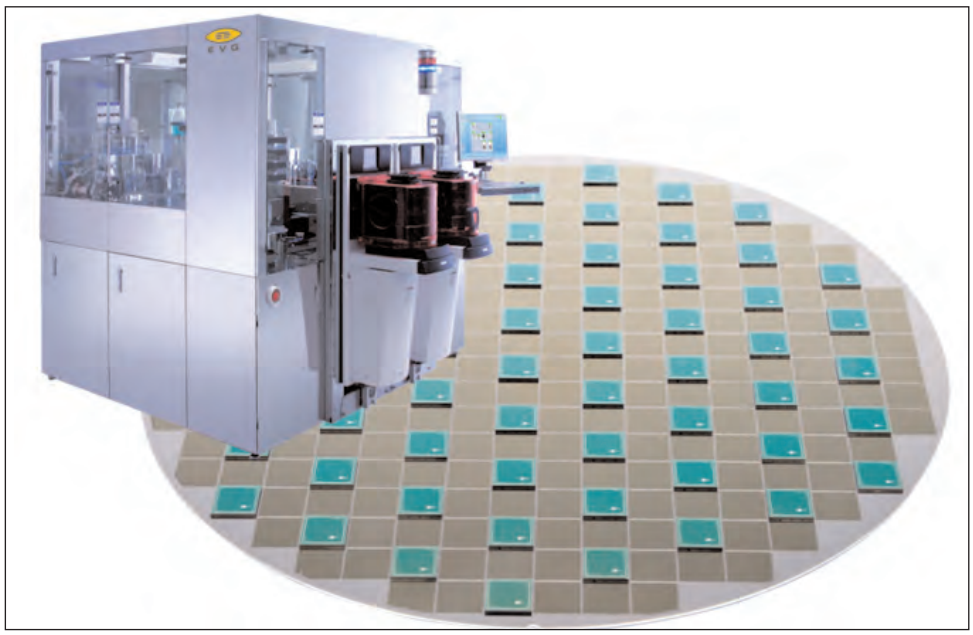

Figure 5: Advanced chip-to-wafer bonded dies on a 200mm wafer.