A VISION FOR THERMALLY INTEGRATED PHOTONICS SYSTEMS

Ryan Enright, Shenghui Lei, Kevin Nolan, Ian Mathews, Alexandre Shen, Guillaume Levaufre, Ronan Frizzell, Guang-Hua Duan, and Domhnaill Hernon

Introduction

We live in a rapidly changing world where advances in consumer technology are occurring so fast that the telecommunications network is struggling to keep pace. The rapid uptake of smartphones/tablets and the widespread wireless streaming of high definition video and games are putting the current network architecture under immense strain, and the problem is compounded as more applications come online on a daily basis. Key examples include high definition 4K video streaming, the online health monitoring required now and in the future for an aging population, the Internet of Things, cloud computing, and the prospect of providing education completely online. Due to the emergence of new web-based applications our projections, based on data from [1], show that international bandwidth requirements will exceed 1 Pbps by 2020, as illustrated in Figure 1. Growth rates for data transmission vary across the network depending on the source (see [2] and [3] for detailed reviews), however, it is widely accepted that yearly performance improvements in products placed within the network lag considerably behind insatiable consumer demand.

To keep pace with these rapid changes, equipment manufacturers need to innovate at speed, at scale, and at low cost. This means that device and component integration is critical to enable the next many generations of efficient and scalable telecommunications products. The level of integration required has severe implications for hardware design in general but even more considerable challenges from a thermal perspective. The thermal challenge grows with ever-increasing levels of integration, as the designer struggles to build more functionality into shrinking package space. Packing so much functionality (e.g., devices and components) into smaller package footprints will lead to substantially increased thermal densities which in turn will require deployment of new thermal solutions.

The efficient nature of transmitting high-speed data optically over long distances makes photonic devices and components key enablers to support the massive growth in data traffic. Silicon (Si) photonics [4] is currently an active area of research due to the advantages of using Si/SiO_2 for passive photonic components to achieve, for example, low loss waveguides and compact arrayed waveguides, the potential for tighter integration with Si-based electronics, as well as a significant processing infrastructure built up by investment in complementary metal oxide semiconductor (CMOS) technology. Arguing economies of scale, this last

Alcatel·Lucent

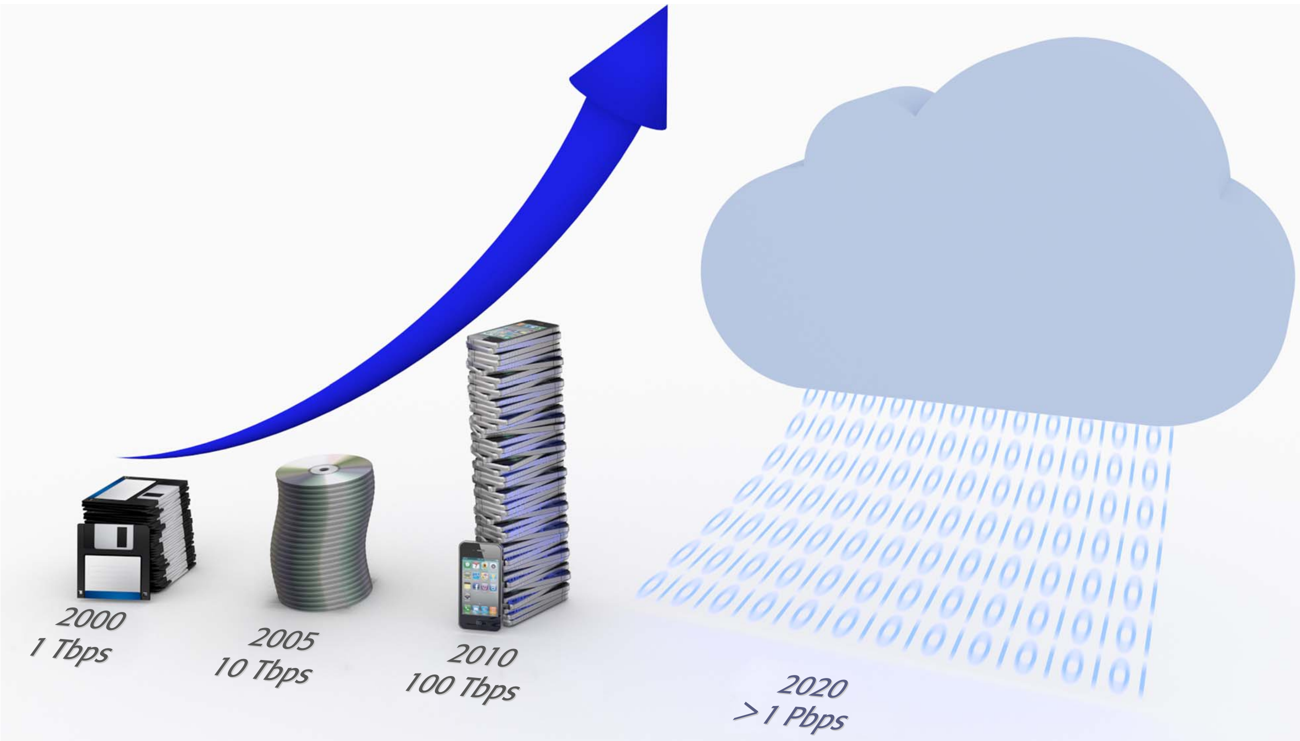

FIGURE 1. Ubiquitous wireless connectivity and the rapid emergence of cloud computing and storage are driving massive growth in data usage thereby straining the current network architecture.

point—the potential of leveraging CMOS technology—offers promise for reducing the cost of silicon-based photonic integrated circuits (PIC) in comparison to their monolithic III-V material counterparts. However, this level of integration also poses significant challenges due to the very different thermal profiles associated with photonics and electronics. In addition, since Si is an indirect bandgap material, an integrated light source that performs well under the required conditions has remained elusive [5]. This has led to the development of heterogeneously integrated laser devices such as III-V/Si hybrid evanescent lasers, where discrete III-V material active photonic devices are placed on a silicon-on-insulator (SOI) substrate [6]. From a thermal perspective, silicon photonics is challenging for several reasons. The fact that the architecture is essentially junction-side-up makes it difficult to remove the heat load; the poor thermal properties of the typical oxide isolation layer on SOI wafers separate the device from the heat sink; and the ridge-like geometry of waveguides in the active device increases the thermal spreading resistance. Further, the distinctly different absolute temperatures and temperature control requirements for photonic and electronic devices present challenges as we move towards opto-electronic integration.

Thermal management has typically been relegated to the last step in the design process and traditionally considered only to ensure long-term operation of the electronic and photonic devices. Today, thermal management represents one of the biggest bottlenecks to releasing next-generation equipment across the entire network and at every scale from the nano (transistor/photonic active region) to the macro (datacenter/telecom network).

In the discussion that follows, we first focus on developing cooling concepts within the optical package where temperature control and the ability to remove extremely large local heat fluxes are the primary challenges addressed at the device/package level. Next, we look at techniques for cooling photonics components ranging from the device to the package and all the way out to system level and report on active areas of thermal research within Bell Labs. Some of these areas require significant research and development both beyond the state-of-the-art and beyond what is presented in this paper. Our goal is to share a vision and early stage results for future Thermally Integrated Photonics Systems (TIPS) that will enable energy-efficient and scalable growth in the network.

State-of-the-Art Thermal Solution for Photonic Packages

Figure 2 shows a schematic representation of a typical solution for the thermal management of a laser array within a photonics integrated circuit (PIC) package. The important thermal aspects shown are the use of a large macro-TEC to cool the lasers. In close proximity to the lasers, resistive heaters are employed to tune the temperature of the lasers thus doubling the local thermal load. The heat load is coupled to the photonics package via many intermediary layers and the entire package is air cooled. If implemented within future highly-integrated PIC devices,

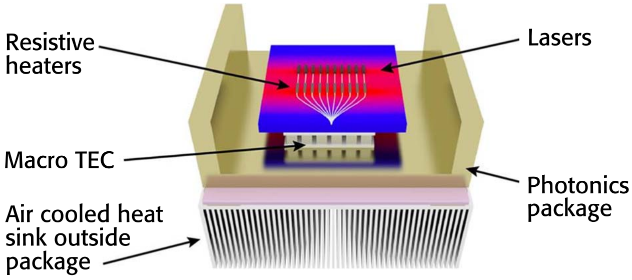

FIGURE 2. Cut away schematic of the current state-of-the-art thermal solution within a photonics integrated circuit package.

today's thermal design approach will lead to significant limitations on the number of components per package and drastically increase OpEx costs for providers due to high power demands. It is our opinion that the current state-of-the-art solution cannot be scaled to meet the thermal requirements for highly integrated PIC devices.

The main design requirement that differentiates photonics from electronics is temperature control, for instance, in high bandwidth multiplexing systems the lasers must be maintained within ± 0.1 °C of their operating temperature to maintain emission wavelengths within design specifications (i.e., to ensure that data is transmitted on-grid). Today's state-of-the-art solution to achieve this tight temperature control combines a large-scale solid-state thermoelectric cooler (macroTEC), resistive heating elements for tuning individual lasers, and large thermal spreaders and sinking to the ambient air via many intermediate layers that increase the overall thermal resistance. Essentially, lasers are cooled with a macroTEC just to be heated up again with resistive heaters. This introduces losses in the design and increases the thermal load locally around the laser.

The temperature control approach leads to a number of thermal design issues that must be addressed:

1) In a temperature-regulated pluggable photonics package today, a macroTEC with a footprint of about 1 cm^2 is employed to cool temperature-sensitive photonic components that dissipate around 1 Watt in total, e.g., an array of 10 lasers each dissipating 100 mW. Thermoelectric cooling is generally inefficient due to material performance constraints and becomes especially so when placed within a hot package that sits on a board in an elevated ambient temperature. In such an environment, substantial temperature differences are required to achieve the sub-ambient operating temperatures of the photonic devices being cooled. The overall thermal design for a system deployed today can lead to a situation where the macroTEC is operating with a coefficient of performance (COP) of < 0.3. This means that upwards of 3 Watts of electrical energy must be supplied to the macroTEC in order to remove that 1 Watt of thermal energy. Compared to the lasers, the macroTEC is much larger in size and also located quite a distance away. This poses limits on the density in which components may be placed within a package and introduces parasitic thermal resistances that further increase the system losses. The use of macroTECs is one of the most significant losses in the current design and as such represents a considerable barrier to achieving greater data rates in smaller packages.
2) Typical design conditions posit that ambient air may enter the system at a temperature of 55 °C. However, due to heating from upstream shelves and other components, a photonics package may regularly encounter a local ambient temperature of 80 °C or more. The macroTEC is employed to reduce the temperature of all lasers below their operating temperature, which is generally well below this ambient. Resistive heaters are then used to heat each laser back to the required temperature. The higher the case temperature of the photonics package, the more electrical power the macroTEC will have to consume to achieve the required laser operating temperature(s).
3) An individual laser dissipates a small thermal load of approximately 100 mW; however, because the lasers are very small in size, the device-level heat flux dissipated per laser is large: around 1 kW/cm^2. Removing this level of heat flux locally through an acceptable thermal resistence is extremely challenging and today thermal spreaders are employed over large areas. The necessity of spreading the heat load locally using poor thermal spreaders implies that the lasers cannot be closely spaced. This adds cost and consumes space within the package, which is of particular concern given the levels of integration required in silicon photonics to meet future network data usage.
4) Finally, the photonic device mounted atop the macroTEC is placed within a large air-cooled package. The space available for heat sinking outside the package is limited due to the constrained slot widths within the product chassis. Furthermore, air-cooling is thermofluidically limited in its ability to remove heat and there are several parasitic thermal resistances between the heat source (laser) and the heat sink that all contribute to inefficiency in the current thermal design.

The current approach to PIC design demonstrates significant thermal inefficiencies, and these inefficiencies limit the scaling of integrated photonics to meet network demands. If integration is key then scalable, efficient thermal solutions are critical. Specifically, several limitations must be addressed including:

- Limited thermal design in and around the photonic device,
- Use of inefficient resistive heaters to tune the laser operating temperature,
- Use of large-scale inefficient thermoelectric modules that consume significant space and power within the package, and
- Use of air cooling both within and external to the package.

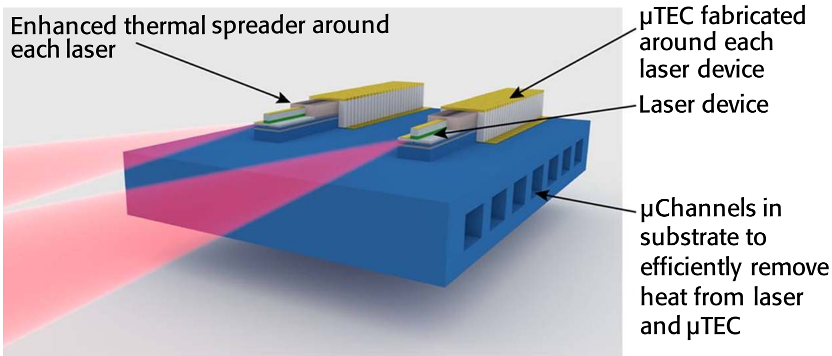

FIGURE 3. High-level schematic of the Thermally Integrated Photonics System architecture. Note that the image shown here focuses on the thermal solution around the heat-generating device (laser) and does not show the overall package design.

The TIPS vision presented below addresses current inefficiencies and provides a path to scalable and efficient thermal solutions for integrated photonics.

TIPS Vision

Figure 3 shows a high-level schematic of the Thermally Integrated Photonics Systems (TIPS) architecture. The TIPS philosophy is to leverage integrated thermal solutions to enable much greater levels of component density than what is possible today. We aim to enable the full potential of silicon photonics, incorporating both photonic and electronic functionality. At a high level, the TIPS concept focuses first on improving the thermal design of the active photonic devices, i.e., lasers, by introducing materials with improved solid-state thermal performance. Second, we aim to provide a more efficient approach to temperature control by removing the large-footprint macroTEC and individual resistive heaters and introducing targeted cooling/temperature control using device-integrated microTECs (μTECs). Third, we introduce microchannels (μChannels) for cooling the hot side of the μTECs (and other chip-level components) to move the heat load from the Si PIC to the package extremity for transfer to the ambient air with minimal thermal resistance. Finally, since a photonic device is embedded within a package and that package resides within a system, we base our approach on the premise that the entire thermal chain must be considered. Ultimately, TIPS represents bringing together best practices across a range of areas in advanced thermal management to provide a source-to-sink solution to enable highly integrated, high performance communications with reduced energy costs. The following sections look at each level of thermal design within TIPS in more detail.

Leverage integrated thermal solutions to enable greater component density.

Level 1—Thermal Design of the Laser

One of the most effective ways to improve thermal performance is to introduce optimized thermal design as close as possible to the location where heat is generated. Various approaches have been used to reduce the thermal resistance of hybrid laser structures: Sysak et al. mitigate the impact of the buried oxide layer by introducing poly-Si (with thermal conductivities of $k_{\mathrm{pSi}} = 15 - 60$ W/m.K) thermal “shunts” by etching away the SiO_2 and silicon epi-layers in the SOI substrate to either side of the waveguide and backfilling with poly-Si [7]. The SiO_2 can be completely replaced with higher k materials such as nanocrystalline diamond, which has been used to produce silicon-on-diamond substrates [8]. In order to overcome the large thermal resistance of the benzocyclobutene (BCB) bonding layer used to attach the III-V material to the SI substrate, Stankovic et al. introduced the concept of extending the p-contact beyond the III-V material to bring it into contact with the SOI substrate near the laser waveguide [9]. This approach was shown to reduce the thermal resistance of the device by 20% to 30%, depending on the thickness of the BCB bonding layer. While all these approaches can reduce the thermal resistance of hybrid lasers, our approach is to enable increased thermal spreading at the heat source itself by placing a highly thermally conductive material in intimate contact with the laser ridge comprised of a p-type region and the multiple quantum well (MQW) active region. This allows heat to be efficiently removed at the source, which will have a greater impact on the device's thermal performance. In particular, we discuss initial results on replacing BCB—a poor thermal conducting polymer currently employed as a waveguide cladding material in some existing hybrid laser designs—with aluminum nitride (AlN), which can have up to a 500X higher thermal conductivity.

In addition to seeking a material with increased thermal conductivity relative to BCB, we also considered the optical confinement and mechanical attributes required to

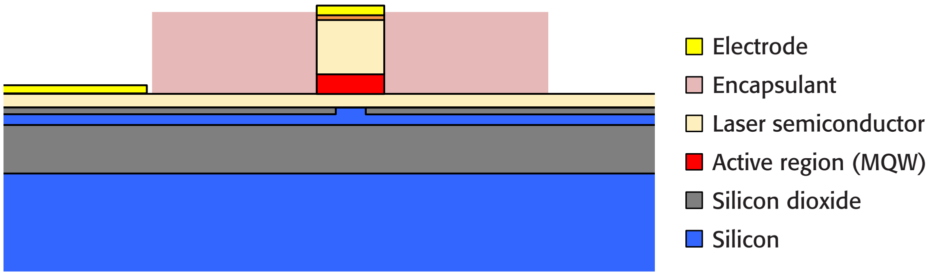

FIGURE 4. Schematic cross-section of a typical hybrid III-V on Si laser architecture fabricated by Bell Labs and its III-V lab (not to scale).

ensure the effective and reliable operation of the laser device. In general, a potential candidate material should have a refractive index less than that of the waveguide ($< \sim 3.2$ in the case of an InP waveguide) to ensure lasing performance is maintained. Based on a survey of non-metallic (i.e., non-electrically conducting) materials demonstrating large intrinsic thermal conductivities [10], we chose aluminum nitride (AlN) as a potential candidate due to its high intrinsic thermal conductivity (320 W/m.K) [11], suitable refractive index (2.12 @ 1550 nm), a coefficient of thermal expansion $(\mathrm{CTE}_{\mathrm{AlN}} \approx 4.5 \times 10^{-6}\ \mathrm{K}^{-1})$ closely matched to InP $(4.6 \times 10^{-6}\ \mathrm{K}^{-1})$ [12], low toxicity, and the fact that it can be deposited as a poly-crystalline film with acceptable thermal conductivity (≥ 100 W/m.K) using low-temperature (< 300 °C) reactive DC magnetron sputtering [13–16], thus ensuring back-end compatibility with existing low temperature

Thermal resistance has considerable impact on a laser's performance.

processes used in hybrid laser device fabrication. These additional considerations are important to prevent failure of the hybrid lasers due to the stresses placed on the laser device during processing and subsequent operation at elevated temperatures.

Thermal resistance is the key thermal parameter with considerable impact on a laser's performance in continuous wave (CW) operation. The thermal resistance of a laser, $R_{th} = \Delta T / P_{\mathrm{d}}$(K/W), describes the temperature increase in the laser active region above a reference substrate temperature (ΔT) as a function of the power dissipated by the laser (P_{d}) and represents a standard measure of impedance to heat flow. Minimizing the laser thermal resistance is desirable from the point of view of reducing or even eliminating the need for inefficient thermoelectric cooling as well as reducing the requirements on the overall system thermal design.

Figure 4 shows a schematic cross-section of a typical hybrid III-V on Si laser architecture fabricated by the Bell Labs III-V lab [6]. The CW optical performance (L-I curve) of the BCB-encapsulated laser as a function of substrate temperature is shown in Figure 5a. The laser shows decreased performance with increasing substrate temperature and roll-over behavior at higher power dissipation, symptoms of excessive temperature in the active region due to self-heating. The laser's thermal resistance was determined by measuring the ratio of the change in output wavelength versus power dissipated in the laser $d\lambda/dP$ (nm/W) during CW operation and the change in output wavelength versus the temperature in the active region $d\lambda/dT$ (nm/K) during pulsed wave operation [7]. We compared the ratio of pulsed and CW characterization of wavelength shift, respectively, with the dissipated power (Figure 5b) and substrate temperature (Figure 5c) to determine the device thermal resistance, $R_{\mathrm{th}} = (d\lambda/dP)/\,(d\lambda/dT) = 122$ K/W.

Using the measured thermal resistance of the laser characterized in Figure 5, we calibrated a simple numerical model developed in Comsol Multiphysics to capture electrically induced device heating and thermal conduction effects within the device structure. A Joule-heating model was applied to obtain the electrical power dissipated through the III-V epitaxial structure, whereas heat conduction was numerically simulated over the entire geometry. To capture heating due to non-radiative recombination, we specified a current-dependent resistance in the laser active region based on the latching voltage found from the intercept of the V-I curve threshold for the hybrid laser described above. To drive the model, we specified a current of 60 mA and a voltage ground at anode and cathode, respectively, resulting in a dissipation of $\sim$100 mW over a laser 780 μm in length. The bottom of the substrate boundary temperature was fixed at 300 °K with all other boundaries treated as adiabatic/symmetric. Note that our early models did not consider the electrical and thermal contact resistances between the different layers of material. Electrical contact resistance between the material layers may impact the overall thermal resistance of the device due to distributed heat generation, while thermal contact resistance may cause an additional drop in temperature in regions of localized power dissipation.

Since the thermal conductivity of the active region (k_{MQW}) was least well-known, this value was

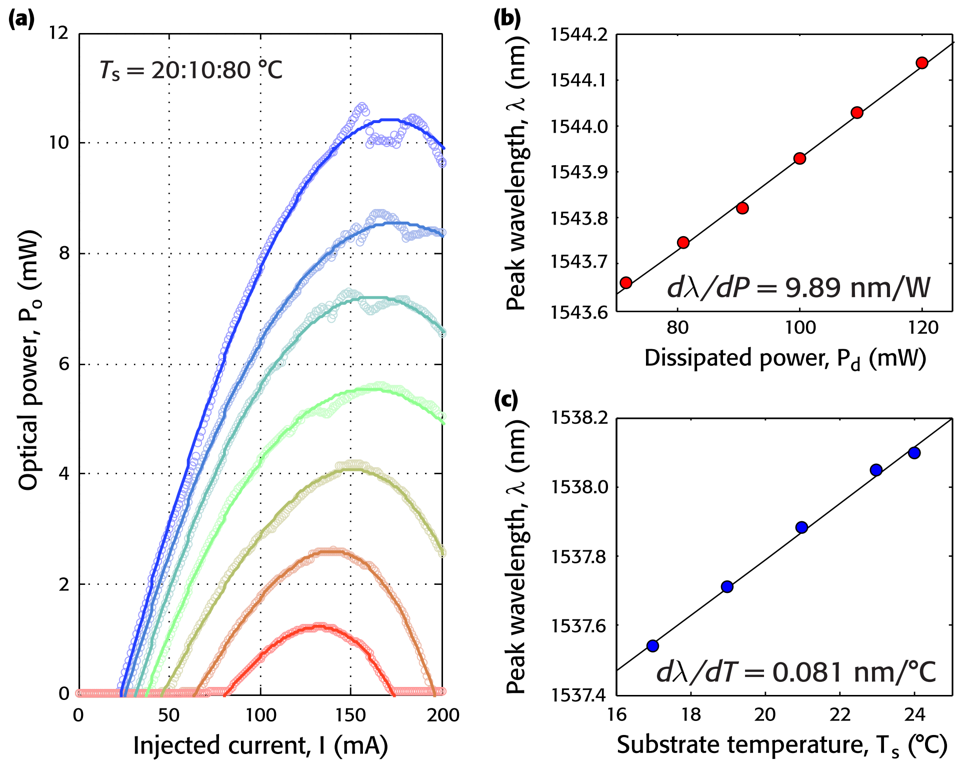

FIGURE 5. **Results of spectroscopic measurements of a hybrid III-V on Si laser emitting in the C-band centered around $\lambda = 1550$ nm. (a) Continuous wave L-I curves of the hybrid laser as a function of substrate temperature, $10^\circ C \leq T_s \leq 80^\circ C$. (b) Continuous wave laser wavelength shift versus dissipated thermal power at $T_s = 20^\circ C$. (c) Pulsed wave laser wavelength shift versus substrate temperature at a given laser output power.**

parameterized and used to find a best fit to the experimental data by varying it over a range of 0.1 W/m.K to 6 W/m.K. The thermal and electrical properties for all other materials were selected based on values referenced in the literature, and we used temperature-dependent properties for Si, InP, and SiO_2 [12, 17, 18]. After calibrating the model, we explored the impact of varying the thermal conductivity of the encapsulation layer around the laser ridge.

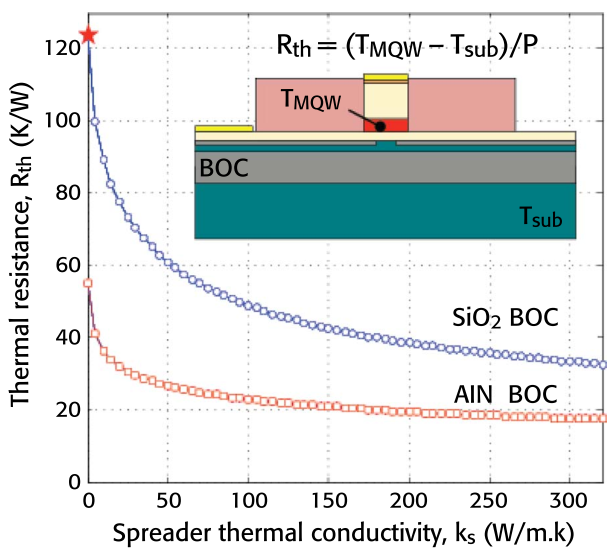

FIGURE 6. **Results of thermal simulations calibrated using the experimental data in Figure 5.**

Figure 6 shows the results of thermal simulations calibrated using the experimental data (indicated by the red star) in Figure 5 by specifying $k_{\mathrm{MQW}} \approx 3.67$ W/m.K, a physically reasonable value considering the InGaAsP composition of the active region [12]. We then simulated the effect of replacing the BCB encapsulation layer with an isotropic material of higher thermal conductivity ($k > 0.3$ W/m.K). Since this layer surrounds the InP-based material, including the heat source region, placing a material of higher thermal conductivity in good thermal contact with it should result in a significantly lower overall device thermal resistance as the heat flow now has a reduced spreading resistance near the source, akin to the improvements in thermal performance that buried heterostructure laser architectures offer over ridge-type architectures in monolithic III-V devices [19]. The results show that the thermal resistance of the hybrid laser could potentially be reduced from 122 K/W to less than 50 K/W by introducing an encapsulant with a conductivity ≥ 130 W/m.K. Combining this with an SOI wafer fabricated with an AlN insulator layer ($k = 136$ W/m.K, 2 μm thickness) instead of SiO_2 further reduced the thermal resistance to $\sim$20 K/W for an encapsulant material with thermal conductivities ≥ 130 W/m.K. The thermal simulation results shown in Figure 6 suggest that using AlN to replace the BCB encapsulate and SiO_2 buried

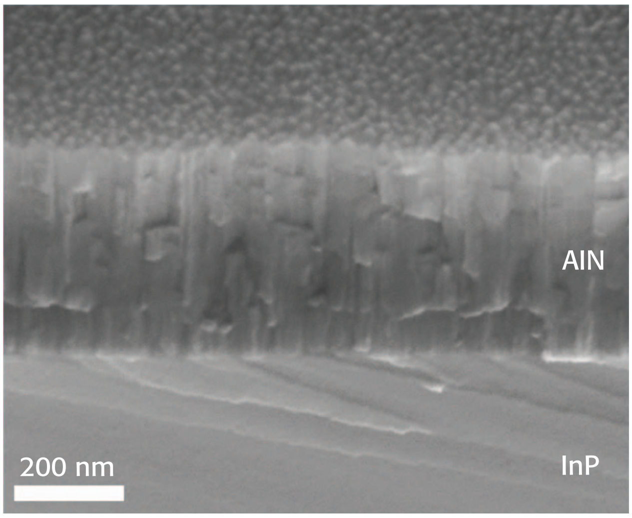

FIGURE 7. Scanning electron microscope image of AlN thin-film deposited by reactive DC magnetron sputtering.

layer may reduce thermal resistance in the hybrid laser design by a scope of 3X to 6X. However, further refinements of the thermal modeling incorporating interfacial effects and deviations from bulk transport behavior based on the structure of the deposited material and the thickness of the device layer are needed to fully understand the extent of the reduction in thermal resistance that may be possible.

Our initial step to explore the integration of AlN into InP-based photonic devices has been to investigate how AlN deposits on single crystal InP. Figure 7 shows a scanning electron microscope (SEM) image of an AlN thin-film deposited by reactive DC magnetron sputtering in a balanced configuration on bare InP wafers without intentional substrate heating. Native oxide was removed from both substrates by chemical etching in an HF solution prior to film deposition. A pure Al target was supplied with a DC power of 2 kW. Pure N_2 was used as the reactive gas entering the chamber at a flow rate of 6 cc/min with the chamber pressure maintained at 6 mTorr. Film thicknesses of 450 nm were deposited at a rate of 18.7 nm/min. Analysis of the deposited film surface grain structure using atomic force microscopy (Asylum Research) indicated a characteristic grain size of $\sim$50 nm suggesting an effective film conductivity of $\leq$ 50 W/m.K [13].

Initial results are promising and elucidate a path to improved laser thermal design. The next steps in this research are to improve the deposition procedure to realize improved film crystallinity (larger grain size) within the device temperature budget ($\sim$300 °C), perform detailed thermal measurements of the deposited AlN films to understand what deposition characteristics provide the highest thermal conductivity, and to investigate the deposition of AlN on non-planar surfaces to understand the film structure formed. Ultimately, we plan to perform full integration in a functioning laser device.

Level 2—Integration of μTECs

A core feature of the TIPS architecture is its efficient approach to temperature control, specifically, the replacement of a single, large, power-hungry macroTEC with μTECs integrated around each heat-generating photonic device, e.g., laser. The use of μTECs promises several advantages including removal of parasitic thermal resistances between the heat source and thermoelectric cooling junction, direct temperature control of individual devices without the need for resistive heaters and their added power demands, reduced thermoelectric material usage, and the potential for increased integration of optoelectronic components with differing operating temperature requirements as in silicon photonics.

Figure 8a provides a schematic of the μTEC concept architecture. The concept involves the serial integration of a μTEC around the hybrid evanescent laser detailed above, assuming an encapsulant layer with a thermal conductivity of 20 W/m.K. In refrigeration mode, heat is pumped from the upper electrode/TE (cold) junction to the lower electrode/TE (hot) junction where it is rejected to the Si substrate resulting in net cooling of the laser active region. The design was driven by two counter-indicative requirements:

1) The introduction of the μTEC should not interfere with the evanescent coupling of light from the InP ridge waveguide to the planar Si waveguide in the SOI substrate or with the optical feedback mechanisms between the laser gain section and a grating etched into the Si waveguide beneath.
2) The laser should be thermally-isolated from the SOI substrate in order to prevent back-diffusion of heat pumped by the μTEC.

The first requirement precluded designs where the device would sit atop the μTEC. However, unlike the uncooled case discussed above where it is desirable to minimize the thermal resistance of all the layers between the laser and the substrate, here the buried SiO_2 layer becomes an advantage by providing a large thermal resistance helping to block back-diffusion of heat from the hot junction to the active region of the laser.

To evaluate the potential of this approach for hybrid laser devices, we developed numerical models in Comsol Multiphysics incorporating the thermoelectric transport governing equations capturing the Joule, Thomson and Seebeck effects [20, 21] to understand the potential performance gains of introducing μTECs. Our implementation of the coupled numerical model capturing electrical, thermal, and thermoelectric effects has been rigorously validated against benchmark cases having exact analytical solutions [22].

The temperature-dependent thermoelectric material properties specified in the simulations correspond to state-of-the-art nanostructured [23, 24] and electrodeposited [25, 26] p-doped $(Bi_{1-x}Sb_x)_2Te_3$ and n-doped $Bi_2(Te_xSe_{1-x})_3$ suitable for room temperature cooling applications. The

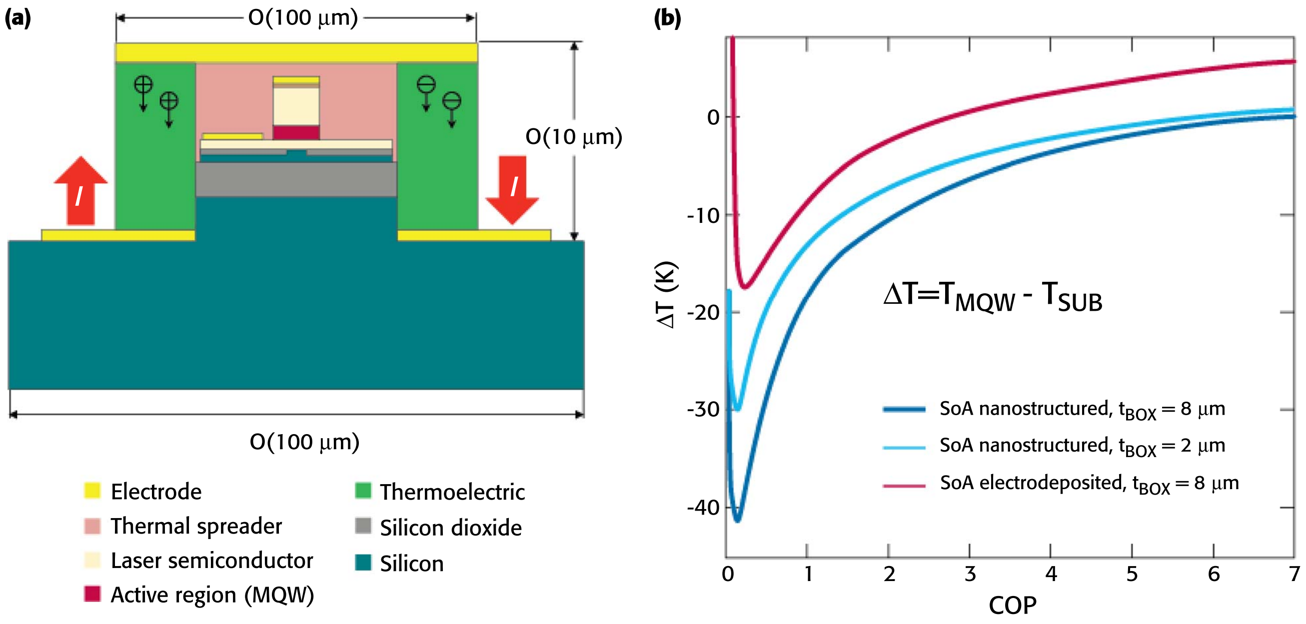

FIGURE 8. **Integrated μTEC temperature control. (a) Schematic of laser device with integrated μTEC (not to scale). (b) Numerically simulated performance of device-integrated μTEC plotting the temperature difference between the active region (T_{MQW}) and the bottom of a 100 μm thick Si substrate ($T_{sub} = 333$ K) as a function of the μTEC COP(= heat pumped/electrical work).**

nanostructured materials, achieved in bulk geometries using a combination of ball milling and hot pressing, represent the best performance measured thus far for thermoelectric materials operating around room temperature, but are aspirational in the context of TIPS since no CMOS-compatible fabrication route has been demonstrated to achieve these levels of performance. On the other hand, though it doesn't achieve the same level of performance as bulk nanostructured materials, electrodeposition shows promise as a back-end process in CMOS fabrication [27, 28]. The specific properties of the electrodeposited materials simulated corresponded to the temperature-dependent power factor of the large diameter nanowires ($\sim$200 nm) studied in [25] and a electrodeposited film thermal conductivity measured in [26] at 300 °K.

Beyond the properties of the thermoelectric material itself, a key performance parameter of μTECs is the electrical contact resistance r_c between the metal electrode and the thermoelectric material that results in parasitic Joule heating and a reduction in device performance. Recent experimental investigations of macroTECs (where contacts are made by soldering) have measured electrical contact resistances of $r_c = 1 \times 10^{-9}\ \Omega.\text{m}^2$ [29]. In μTECs, where the leg height is on the order of 10 μm, these values represent a significant contribution to the overall resistance of the device that can lead to the situation of zero net cooling [30]. However, carefully formed contacts can typically demonstrate better behavior and in our simulations we specified $r_c = 5 \times 10^{-11}\ \Omega.\text{m}^2$ [30–32]. Note, however, that this value is still almost two orders of magnitude larger than the estimated theoretical lower limit based on carrier tunneling transport across the contact [33], suggesting that there is still room for improvement.

Figure 8b shows selected simulation results carried out to characterize the potential performance of the proposed μTEC architecture. The figure shows the temperature difference ΔT between the laser active region and the substrate (which was held at a fixed temperature of 333 °K) plotted against the coefficient of performance (COP = heat pumped/electrical work) for a laser dissipating approximately 0.13 W/mm and a corresponding cold side heat flux of 178 W/cm^2 based on the power dissipated and the area of the thermoelectric elements. By increasing the driving current through the μTEC, we were able to increase the temperature of the active region over that of the substrate. This, however, came at the expense of increasing the COP. Once the point of maximum refrigeration was reached, the temperature differential decreased as Joule heating effects began to dominate μTEC operation. As expected, the nanostructured material outperformed the electrodeposited material due to its superior thermoelectric properties—a maximum $\Delta T \approx 40$ K, compared to 18 K. We also captured the important role of the buried oxide layer, demonstrating that thickening this layer from the standard 2 μm used in III-V lab fabrication to 8 μm had a significant effect on μTEC performance by increasing the maximum ΔT by 33% for the nanostructured thermoelectric material. Again, the latter result highlights the critical need for improved thermal design consideration in the fabrication of laser devices.

Our initial simulations suggest that μTEC integration around active photonic devices has the potential to provide both temperature regulation and appreciable refrigeration. While current state-of-the-art device-integrated

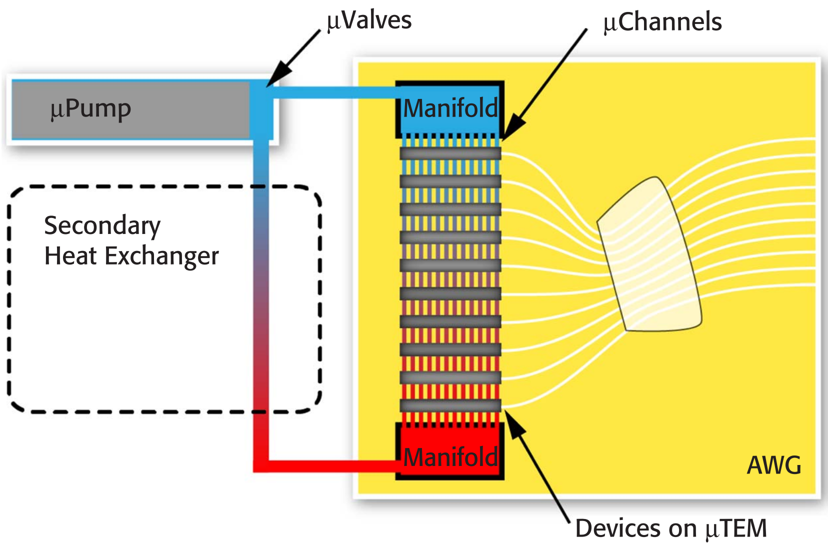

FIGURE 9. Schematic of the μFluidic components including μPumps to pump the fluid, μValves to control the fluid direction, μChannels to carry the fluid to the heat source and heat exchangers to dump the heat load.

μTECs are unlikely to outperform macroTEC modules, thin-film or otherwise, within a thermally-optimized system they promise to free the designer to achieve greater optoelectronic integration. Of course, there are significant challenges to overcome. Design rules need to be formulated, device fabrication strategies need to be further developed and, critically, thermoelectric material and electrical contact processing needs to be optimized to achieve maximum performance.

Level 3—Integration of μFluidics

Microfluidics (μFluidics) is the umbrella term we apply to all of the small scale fluidic-based cooling architectures we need to introduce within the package to remove large heat fluxes in an efficient and scalable way. Figure 9 provides an overview of the μFluidic aspect of the TIPS architecture. A μPump, suitably sized for integration into the package, drives flow (typically water or some other non-electrically conducting fluid with similar performance) around a flow loop and μValves regulate the flow within an array of μChannels that carry the heat from the hot side of the μTECs and the lasers. The μChannels may be etched directly into the device substrate or located in a submount. The coolant passes through the μChannels to remove the heat generated by the laser bar and μTECs and then moves to a secondary heat exchanger rejecting heat to the ambient before repeating the loop. We identified that new approaches in μChannel cooling would provide the biggest gains for the overall TIPS architecture and for that reason, this paper primarily focuses on requirements for μChannel cooling. However, we also point out that the improvements necessary to bring μPump and μValve performance to the level necessary to ensure reliable operation in telecommunications deployment will require considerable research effort. We hope to leverage the considerable knowledge in the biomedical field around the reliable deployment of valves, pumps, and connectors and tailor the designs to meet the requirements of the telecommunications industry [34–36]. However, there are alternative approaches that avoid the need for mechanical pumping including capillary-driven phase-change cooling devices which operate much like a heat pipe. This fluidic-based cooling mechanism has the advantage of being passively driven by the heat source being cooled, but requires more complex fabrication steps for complete integration at chip-level. This approach is being explored in the DARPA-funded ICECool Fundamentals program to develop high performance two-phase thermal management with ultra high heat flux dissipation [37].

Today there are no commercially-available liquid cooled photonic packages; however, liquid cooling has garnered considerable attention within the high end computer/server markets and IBM has played an active role in bringing liquid cooling approaches to market. In particular, they promote the use of integrated water-cooled microchannels around a 3D stacked chip architecture. Here, however, pumping is provided via commercially-available rotary pumps that at $60 \times 70 \times 30$ mm [38] are too large for most telecom applications. Iverson and Garimella reviewed the literature on micro scale pumping technologies [39] and concluded that mechanical displacement and electro-osmotic pumps may offer the best compromise between flow rate and pressure rise. However, both commercially available and research pump designs fall short of telecommunications industry requirements for long-term reliability, small form factor, and energy efficiency. For example, commercial piezoelectric diaphragm pumps—which would provide a rise in pressure adequate to drive water through microchannels—occupy a volume similar in size to a CFP4 pluggable ($\sim$20 cm^3). The primary reason a pump of this size is required is because of the large pressure drop in the microchannels combined with the poor performance of micro pumps. We identified microchannel heat transfer improvements as a key driver for enabling improvements across the entire TIPS architecture, i.e., if we can reduce the drop in microchannel pressure, we can employ much smaller pumps.

From a thermal perspective, μChannels are of particular interest due to the extremely large surface area they provide for heat transfer. In fact, in fully developed laminar flow, the heat transfer coefficient (h) simply scales with the inverse of the channel diameter (D) [40].

$$h \sim \frac{1}{D}.$$

Tuckerman and Pease demonstrated heat dissipation of 790 W/cm^2 with a substrate temperature increase of 71 °C and this result remains one of the highest power densities achieved in single-phase liquid-cooled μChannels [41]. However, a significant rise in pressure ($\sim$200 kPa, or $\sim$2 atm.) was required to drive the flow across a 1 cm long μChannel heat exchanger [41]. This high pressure drop through the channels is required to overcome viscous shear stress which scales with the inverse of the fourth power of the channel diameter [40].

$$\Delta P \propto \frac{1}{D^4}.$$

Therefore, while smaller channels offer better heat transfer performance, the nature of viscous scaling represents a significant challenge in practical systems since much larger and more power-hungry pumps are required. Given that the drop in pressure scales with the channel dimension as discussed above, it is not surprising that recent work has explored the use of liquid metals as an alternative to water in pumping-constrained systems where the increased thermal conductivity of these liquids allows for larger channel dimensions [42–44]. Alternatively, turbulent flow can significantly heighten the transfer of heat due to localized mixing of the fluid layers, which in turn allows more dissipation of heat due to the continuous replacement of the cooling fluid. However, classically turbulent channel flows manifest at prohibitively large Reynolds numbers $(Re \approx 2300)$ for water in micro/minichannel geometries implying extremely large pressure drop penalties. Therefore, to enable our TIPS vision we need to look to different solutions.

Within the TIPS program, we propose to introduce the relatively new μChannel cooling paradigm called *elastic turbulence* [45], where a viscoelastic fluid is used to generate local mixing at the microscale, thereby tricking the flow into thinking it is turbulent at very low Reynolds numbers/flow rates. Mena et al. [46] and Hartnett and Kostic [47] were among the first to demonstrate enhanced heat transfer characteristics of viscoelastic flows in non-circular channels. Heat transfer studies in non-circular channels demonstrated a threefold enhancement of the Nusselt number (*Nu*) without an accompanying drop in pressure for an aqueous solution of 1000 ppm polyacrylamide (Separan AP 273) [47]. The authors attributed the increase in the *Nu* to secondary flows, meaning that *Nu* was dependent on the Weissenberg number,

$$Wi = \frac{N_1}{\tau_{xy}},$$

where N_1 is the first normal stress difference and $\tau_{xy} = \mu \frac{du}{dy}$ is the shear stress.

Much of the experimental and theoretical work on viscoelastic instabilities in fluid flows originated in Bell Labs with several seminal papers [48–50]. Visualization of Taylor-Couette flows by Muller et al. used a 1000 ppm solution of high molecular weight polyisobutylene $(\mathrm{M_w} \sim 4 - 6 \times 10^6)$ dissolved in a low molecular weight polybutene [48]. Mica flakes, which align with the flow to reveal the flow structure, were used to demonstrate how flow is dependent on the Weissenberg number.

More recently, research has demonstrated how a viscous flow between rotating disks can exhibit turbulent behavior with significant amounts of mixing at subcritical Reynolds numbers [45]. As above, this was achieved with the addition of polymers, in this case just 80 ppm of high molecular weight polyacrylamide $(\mathrm{M_w} = 18 \times 10^6)$. Experiments by Burghelea et al. in serpentine microchannels using a polymer solution identical to that in [45] verified that the effect is dependent on the Weissenberg number with a non-linear transition from Newtonian behavior for $Wi > 1.4$ [51]. However, although there has been a wealth of research in the literature on the effect of introducing different polymers in terms of fluid dynamic effects, there has been very little attention paid to understanding the impact they have on the heat transfer/pressure drop tradeoff important to thermal management. In TIPS we are focusing our attention on filling the knowledge gap around how elastic turbulence can be employed in real systems to improve thermal performance.

We are focusing our attention on filling the knowledge gap around how elastic turbulence can be employed in real systems to improve thermal performance.

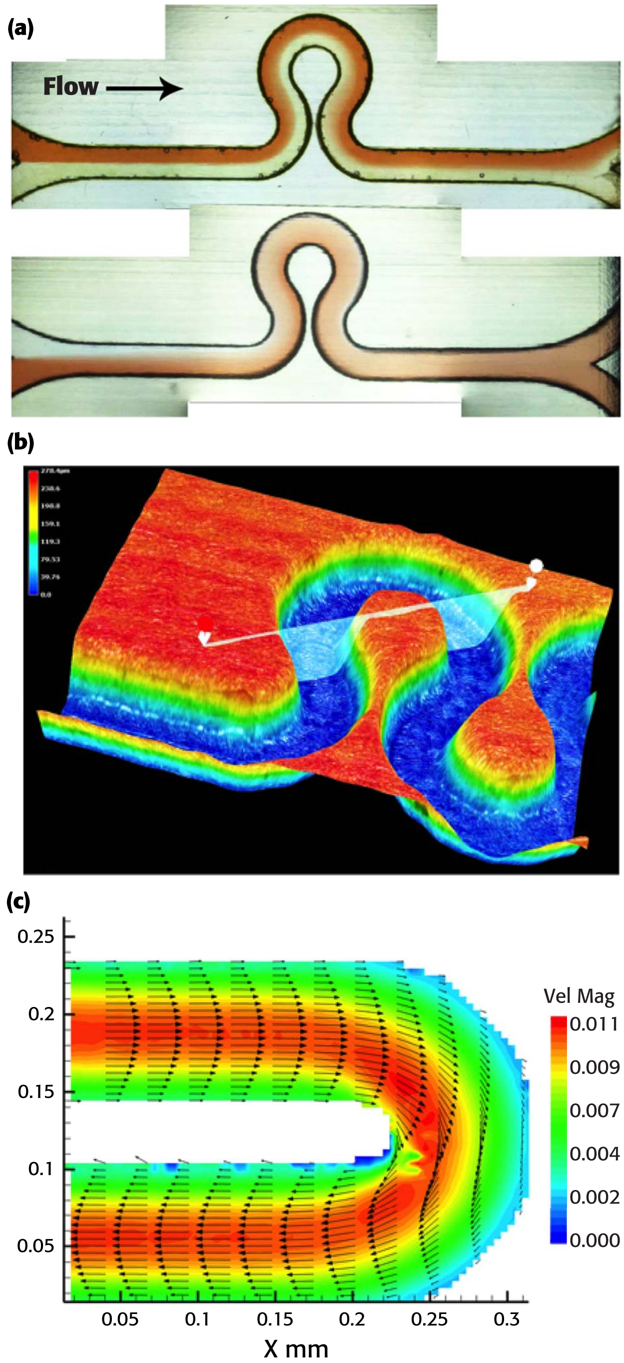

FIGURE 10. **3D illustrations of the μChannel.**

Figure 10 provides a series of three-dimensional views of a representative μChannel design. Figure 10a illustrates the viscoelastic turbulence phenomenon in a 3D printed μChannel with a 400 μm hydraulic diameter, while Figure 10b shows a 3D cross section of a typical μChannel geometry fabricated with high-resolution 3D printing. We employ high resolution 3D printing to de-risk various channel designs prior to investing in silicon fabrication. The channel in Figure 10a has two inlets, two exits, and a single serpentine region around which the flow must pass. The upper image shows deionized water entering from both inlets, one of which is dyed. Because the flow is viscosity-dominated and in the laminar regime $(\mathrm{Re} \approx 100)$, there is no mixing across stream lines as shown in the upper image of Figure 10a. The flow exits much as it enters with only a small amount of diffusion visible across the stream interface. The lower image in Figure 10a shows the result of adding 1000 ppm of polyacrylamide $(\mathrm{M_w} = 40 \times 10^3)$ to the working fluid at a similar Reynolds number. Immediately there appears to be greater diffusion between the parallel fluid streams due to the potential for instability in the shear layer between the inlet streams; however, they still remain distinct until the serpentine region is encountered. By the time the flow has passed the serpentine region, the dye has fully mixed. Thereafter the fluid returns to a laminar state similar to that seen upstream of the serpentine. Figure 10c shows the flow around the bend in a serpentine channel using micro particle image velocimetry (μPIV). We employ μPIV specifically to provide highly-resolved measurements of the micro scale flow field to gain a greater understanding of the underlying flow physics in order to drive improvements for real system deployments.

The mechanism by which the polymer chains affect the flow can be understood by considering that under a typical shear flow condition, polymer chains suspended in the fluid align and stretch with the direction of shear, τ_{xy}. When a disturbance occurs in the flow, such as with an obstacle or a change in flow direction, the polymers release their stored energy and can destabilize the surrounding flow to induce local mixing. This behavior suggests that we can enhance heat transfer locally, via elastic mixing, while not incurring a prohibitive pressure drop penalty over the entire length of the channel as is the case in traditional μChannel cooling. In fact, when employing elastic turbulence, the channel diameter can be substantially increased in size to further minimize the drop in pressure compared to that in traditional μChannels. For example, in our TIPS architecture, it may be possible to replace the μChannel array with a large effective diameter manifold containing localized obstacles (such as posts) that can destabilize the flow to provide targeted and effective heat removal. In this way we circumvent the traditional scaling dynamics of μChannels and provide high heat transfer locally with a minimal pressure drop penalty across the system, thus enabling the use of a much smaller pump.

Note however that practical challenges remain in the implementation of elastic turbulence. In particular, polymer degradation due to mechanical shear and temperature variation is a reliability concern [51]. Li et al. have proposed the use of surfactant micelles, which can self-repair, as a suitable alternative to polymers [52]. Their experiments with a dilute aqueous solution of cetyltrimethyl ammonium chloride (CTAC) exhibited much the same elastic turbulence behavior as previously studied polymer systems, but without the temperature and mechanical degradation effects.

Though a significant amount of work is required for practical implementation, we believe we have identified a path to enable in-package liquid cooling using a new approach that does not require large pumps.

Level 4—Package and System Cooling

As discussed above, efficient cooling within the package represents only one aspect of the overall problem. Densely packing a large number of optical input/output devices on a faceplate within a tightly-confined circuit pack presents additional thermal challenges, as we detail below:

- Industry standards dictate that the case temperature of the optical package must be maintained at 70 °C. In the telecommunications environment, equipment is designed to operate at ambient temperatures as high as 55 °C [53, 54]. In real deployments, there are many optical input/output (I/O) packages placed in line on the faceplate and there may be multiple shelves of systems within a single rack. This means that the ambient air temperature in proximity to the optical I/O packages can be as high as 80 °C due to heating from upstream shelves and upstream components.
- The TIPS project addresses the thermal challenge within the package. The other major problem, however, is to dissipate the heat that is released from the optical I/O package to the surrounding airflow on the circuit board. There are two major barriers: first, because the optical I/Os are pluggable, there is a low contact pressure between the optical I/O package and the heat sink that sits on top of the package. This low contact pressure means that there is a high contact thermal resistance between the optical I/O package and the heat sink. This high resistance limits heat flow. The second major barrier to removing heat from the package is the thermal resistance between the heat sink and the air. Given the current design of I/O package-level thermal pluggable optical interconnects, it will become difficult to dissipate heat fluxes in excess of $\sim$0.31 W/cm^2 for an acceptable increase in temperature [55]. This will significantly reduce the scope for the smaller form factors and increased functionality that the market is driving towards.
- The optical I/Os sit in close proximity to high-power electronic devices that dissipate up to two orders of magnitude more heat than they do and operate at significantly higher case temperatures—between 100 °C for field programmable gate arrays (FPGAs) and 225 °C for power amplifiers. The significant discrepancy in thermal dissipation and operating temperature between the optical I/Os and the remaining components poses further thermal challenges due to thermal cross-talk and thermal pollution issues.
- All of the components on the circuit pack are air-cooled. The components sit on boards that sit within shelves of equipment within a rack. Each shelf has at least one fan tray with multiple axial fans, but the effective use of fans is hampered by strict limitations on noise emissions specified by the standards bodies [54, 56]. A critical area where future research is essential is the reduction of acoustic noise emissions so that higher rates of air flow can be generated.
- Finally, the total allowable power that can be dissipated per shelf is limited to < 30 kW for air-cooled racks. These levels have already been reached in datacom deployments and they are rapidly being approached in telecommunications. Local heat fluxes generated by high-power components are increasing at a rapid pace and air cooling is struggling to remove such high heat fluxes. Liquid cooling has been demonstrated as a viable option in the datacom world and telecommunications must also consider it as a viable solution. Telecommunications applications, however, must adhere to more stringent design regulations as well as expectations for much longer product lifetimes as specified in the NEBS/ETSI design standards. Therefore, further research is required to demonstrate reliable liquid cooling within and external to the photonics package in order to enable the exponential growth in data traffic expected to hit telecom networks in the near future.

Conclusion

Thermal management has traditionally been relegated to the last step in the design process. However, with the exponential growth in data traffic leading to ever-greater levels of component integration and ever-higher levels of energy consumption, thermal management is rapidly becoming one of the most critical areas of research within the ICT industry. Given the vast use of optics for efficient transmission of high-speed data, this paper focuses on a new thermal solution for cooling the components within pluggable optical modules. We call our architecture Thermally Integrated Photonics Systems (TIPS), and it represents a new vision for the thermal building blocks required to enable exponential traffic growth in the global telecommunications network. The thermal blocks are designed such that traditional or emerging photonics architectures can be cooled.

In the TIPS program, we identified that existing thermal solutions cannot scale to meet the needs of exponential growth in data traffic. We identified the main barriers to enabling further growth and developed a research roadmap around a scalable and efficient integrated thermal solution. In particular, we are investigating the effects of replacing inefficient materials and large macroTECs with better thermal spreaders and μTECs. In addition, we are looking to introduce new forms of μChannel cooling into the package to more efficiently remove the heat generated by the lasers and the TECs, leading to future photonic devices that can be deployed in a vastly more dense and integrated manner to address the requirements of future telecommunication networks.

Acknowledgement

Bell Labs Ireland acknowledges the financial support of the Industrial Development Agency (IDA) Ireland.

References

[1] J. Brodkin, "Bandwidth Explosion: As Internet Use Soars, Can Bottlenecks Be Averted?" Arstechnica, May 1, 2014. [Online]. Available: http://arstechnica.com/business/2012/05/bandwidth-explosion-as-internet-use-soars-can-bottlenecks-be-averted/. [Accessed: 14-Jul-2014].

[2] P. Winzer, "Spatial Multiplexing in Fiber Optics: The 10X Scaling of Metro/Core Capacities," Bell Labs Tech J, vol. 19, pp. 22–30, 2014.

[3] R. W. Tkach, "Scaling Optical Communications for the Next Decade and Beyond," Bell Labs Tech. J., vol. 14, no. 4, pp. 3–9, 2010.

[4] R. Soref, "The Past, Present, Future of Silicon Photonics," Sel. Top. Quantum Electron. IEEE J., vol. 12, no. 6, pp. 1678–1687, Dec. 2006.

[5] D. Liang and J. E. Bowers, "Recent Progress in Lasers on Silicon," Nat Photon, vol. 4, no. 8, pp. 511–517, Aug. 2010.

[6] G.-H. Duan, C. Jany, A. Le Liepvre, A. Accard, M. Lamponi, D. Make, P. Kaspar, G. Levaufre, N. Girard, F. Lelarge, J.-M. Fedeli, A. Descos, B. Ben Bakir, S. Messaoudene, D. Bordel, S. Menezo, G. de Valicourt, S. Keyvaninia, G. Roelkens, D. Van Thourhout, D. J. Thomson, F. Y. Gardes, and G. T. Reed, "Hybrid III-V on Silicon Lasers for Photonic Integrated Circuits on Silicon," Sel. Top. Quantum Electron. IEEE J., vol. 20, no. 4, pp. 1–13, Aug. 2014.

[7] M. N. Sysak, D. Liang, R. Jones, G. Kurczveil, M. Piels, M. Fiorentino, R. G. Beausoleil, and J. E. Bowers, "Hybrid Silicon Laser Technology: A Thermal Perspective," Sel. Top. Quantum Electron. IEEE J., vol. 17, no. 6, pp. 1490–1498, Dec. 2011.

[8] D. Liang, M. Fiorentino, S. T. Todd, G. Kurczveil, R. G. Beausoleil, and J. E. Bowers, "Fabrication of Silicon-on-Diamond Substrate and Low-Loss Optical Waveguides," Photonics Technol. Lett. IEEE, vol. 23, no. 10, pp. 657–659, May 2011.

[9] S. Stanković, G. Roelkens, D. Van Thourhout, R. Baets, R. Jones, M. Sysak, and B. Koch, Hybrid III-V/Silicon Laser Based on DVS-BCB Bonding, 2008.

[10] G. A. Slack, "Nonmetallic Crystals With High Thermal Conductivity," J. Phys. Chem. Solids, vol. 34, no. 2, pp. 321–335, 1973.

[11] G. A. Slack, R. A. Tanzilli, R. O. Pohl, and J. W. Vandersande, "The Intrinsic Thermal Conductivity of AIN," J. Phys. Chem. Solids, vol. 48, no. 7, pp. 641–647, 1987.

[12] Ioffe Institute, "Physical Properties of Semiconductors," NSM Archive. [Online]. Available: http://www.ioffe.rssi.ru/SVA/NSM/Semicond/. [Accessed: 10-Jan-2014].

[13] B. E. Belkerk, A. Soussou, M. Carette, M. A. Djouadi, and Y. Scudeller, "Structural-Dependent Thermal Conductivity of Aluminium Nitride Produced by Reactive Direct Current Magnetron Sputtering," Appl. Phys. Lett., vol. 101, no. 15, pp. 151908-1–151908-4, Oct. 2012.

[14] C. Duquenne, M. A. Djouadi, P. Y. Tessier, P. Y. Jouan, M. P. Besland, C. Brylinski, R. Aubry, and S. Delage, "Epitaxial Growth of Aluminum Nitride on AlGaN by Reactive Sputtering at Low Temperature," Appl. Phys. Lett., vol. 93, no. 5, 2008.

[15] C. Duquenne, B. Popescu, P.-Y. Tessier, M.-P. Besland, Y. Scudeller, C. Brylinski, S. Delage, and M.-A. Djouadi, "Magnetron Sputtering of Aluminium Nitride Thin Films for Thermal Management," Plasma Process. Polym., vol. 4, no. S1, pp. S1–S5, 2007.

[16] C. Duquenne, M.-P. Besland, P. Y. Tessier, E. Gautron, Y. Scudeller, and D. Averty, "Thermal Conductivity of Aluminium Nitride Thin Films Prepared by Reactive Magnetron Sputtering," J. Phys. Appl. Phys., vol. 45, no. 1, Jan. 2012.

[17] V. Palankovski, "Simulation of Heterojunction Bipolar Transistors," Ph.D, TU Wien, 2000.

[18] T. P. Pearsall and J. P. Hirtz, "The Carrier Mobilities in Ga0.47-In0.53As Grown by Organo-Mettalic CVD and Liquid-Phase Epitaxy," J. Cryst. Growth, vol. 54, no. 1, pp. 127–131, Jul. 1981.

[19] K. Takemasa, M. Kubota, T. Munakata, and H. Wada, "1.3-μm AlGaInAs Buried-Heterostructure Lasers," Photonics Technol. Lett. IEEE, vol. 11, no. 8, pp. 949–951, Aug. 1999.

[20] D. M. Rowe, Thermoelectrics Handbook: Macro to Nano. CRC Press, 2005.

[21] G. Chen, Nanoscale Energy Transport and Conversion. Oxford, UK: Oxford Univ. Press, 2005.

[22] S. Lei, R. Frizzell, and R. Enright, "Benchmark Numerical Simulations of Solar Thermoelectric Generators," in Proceedings of the 15th International Heat Transfer Conference, Kyoto, Japan, 2014.

[23] B. Poudel, Q. Hao, Y. Ma, Y. Lan, A. Minnich, B. Yu, X. Yan, D. Wang, A. Muto, D. Vashaee, X. Chen, J. Liu, M. S. Dresselhaus, G. Chen, and Z. Ren, "High-Thermoelectric Performance of Nanostructured Bismuth Antimony Telluride Bulk Alloys," Science, vol. 320, no. 5876, pp. 634–638, May 2008.

[24] X. Yan, B. Poudel, Y. Ma, W. S. Liu, G. Joshi, H. Wang, Y. Lan, D. Wang, G. Chen, and Z. F. Ren, "Experimental Studies on Anisotropic Thermoelectric Properties and Structures of n-type Bi2Te2.7Se0.3," Nano Lett., vol. 10, no. 9, pp. 3373–3378, Jul. 2010.

[25] S. Bäßler, T. Böhnert, J. Gooth, C. Schumacher, E. Pippel, and K. Nielsch, "Thermoelectric Power Factor of Ternary Single-Crystalline Sb2Te3- and Bi2Te3-based Nanowires," Nanotechnology, vol. 24, no. 49, pp. 495402, 2013.

[26] C. Schumacher, K. G. Reinsberg, R. Rostek, L. Akinsinde, S. Baessler, S. Zastrow, G. Rampelberg, P. Woias, C. Detavernier, J. A. C. Broekaert, J. Bachmann, and K. Nielsch, "Optimizations of Pulsed Plated p and n-type Bi2Te3-based Ternary Compounds by Annealing in Different Ambient Atmospheres," Adv. Energy Mater., vol. 3, no. 1, pp. 95–104, 2013.

[27] H. Bottner, J. Nurnus, and A. Schubert, "Miniaturized Thermoelectric Converters," in Thermoelectrics Handbook: Macro to Nano. CRC Press, 2006, pp. 46.1–46.18.

[28] G. J. Snyder, J. R. Lim, C.-K. Huang, and J.-P. Fleurial, "Thermoelectric Microdevice Fabricated by a MEMS-Like Electrochemical Process," Nat Mater, vol. 2, no. 8, pp. 528–531, Aug. 2003.

[29] S. R. Annapragada, T. Salamon, P. Kolodner, M. Hodes, and S. V. Garimella, "Determination of Electrical Contact Resistivity in Thermoelectric Modules (TEMs) From Module-Level Measurements," Components Packag. Manuf. Technol. IEEE Trans., vol. 2, no. 4, pp. 668–676, Apr. 2012.

[30] H. Bottner, J. Nurnus, A. Gavrikov, G. Kuhner, M. Jagle, C. Kunzel, D. Eberhard, G. Plescher, A. Schubert, and K.-H. Schlereth, "New Thermoelectric Components Using Microsystem Technologies," Microelectromechanical Syst. J., vol. 13, no. 3, pp. 414–420, Jun. 2004.

[31] S.-P. Feng, Y.-H. Chang, J. Yang, B. Poudel, B. Yu, Z. Ren, and G. Chen, "Reliable Contact Fabrication on Nanostructured Bi2Te3-based Thermoelectric Materials," Phys. Chem. Chem. Phys., vol. 15, no. 18, pp. 6757–6762, 2013.

[32] I. Chowdhury, R. Prasher, K. Lofgreen, G. Chrysler, S. Narasimhan, R. Mahajan, D. Koester, R. Alley, and R. Venkatasubramanian, "On-Chip Cooling by Superlattice-Based Thin-Film Thermoelectrics," Nat. Nanotechnol., vol. 4, no. 4, pp. 235–238, Apr. 2009.

[33] L. W. da Silva and M. Kaviany, "Micro-Thermoelectric Cooler: Interfacial Effects on Thermal and Electrical Transport," Int. J. Heat Mass Transf., vol. 47, no. 10–11, pp. 2417–2435, May 2004.

[34] P. Gravesen, J. Branebjerg, and O. S. Jensen, "Microfluidics: A Review," J. Micromechanics Microengineering, vol. 3, no. 4, pp. 168, 1993.

[35] C.-Y. Lee, C.-L. Chang, Y.-N. Wang, and L.-M. Fu, "Microfluidic Mixing: A Review," Int. J. Mol. Sci., vol. 12, no. 5, pp. 3263–3287, 2011.

[36] E. K. Sackmann, A. L. Fulton, and D. J. Beebe, "The Present and Future Role of Microfluidics in Biomedical Research," Nature, vol. 507, no. 7491, pp. 181–189, Mar. 2014.

[37] D. F. Hanks, Z. Lu, S. Narayanan, K. R. Bagnall, R. Raj, R. Xiao, R. Enright, and E. N. Wang, "Nanoporous Evaporative Device for Advanced Electronics Thermal Management," presented at the Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2014 IEEE Intersociety Conference on, 2014, pp. 290–295.

[38] A. K. Coskun, D. Atienza, T. S. Rosing, T. Brunschwiler, and B. Michel, “Energy-Efficient Variable-Flow Liquid Cooling in 3D Stacked Architectures,” in Design, Automation & Test in Europe Conference & Exhibition, 2010, pp. 111–116.
[39] B. Iverson and S. Garimella, “Recent Advances in Microscale Pumping Technologies: A Review and Evaluation,” Microfluid. Nanofluidics, vol. 5, no. 2, pp. 145–174, Aug. 2008.
[40] S. G. Kandlikar, S. Garimella, D. Li, S. Colin, and M. R. King, Heat Transfer and Fluid Flow in Minichannels and Microchannels. Elsevier, 2005.
[41] D. B. Tuckerman, and R. F. W. Pease, “High-Performance Heat Sinking for VLSI,” Electron Device Lett. IEEE, vol. 2, no. 5, pp. 126–129, May 1981.
[42] M. Hodes, R. Zhang, L. S. Lam, R. Wilcoxon, and N. Lower, “On the Potential of Galinstan-Based Minichannel and Minigap Cooling,” Components Packag. Manuf. Technol. IEEE Trans., vol. 4, no. 1, pp. 46–56, Jan. 2014.
[43] L. S. Lam, M. Hodes, and R. Enright, “Analysis of Galinstan-Based Microgap Cooling Enhancement Using Structured Surfaces,” in ASME 2013 Heat Transfer Summer Conference, Minneapolis, Minnesota, USA, 2013, vol. 3, pp. V003T23A006.
[44] R. Zhang, M. Hodes, N. Lower, and R. Wilcoxon, “High Heat Flux, Single-Phase Microchannel and Minichannel Cooling With Water and Liquid Metal,” Electronics Cooling, 2014. [Online]. Available: http://www.electronics-cooling.com/2014/09/high-heat-flux-single-phase-micro-channel-minichannel-cooling-water-liquid-metal/. [Accessed: 17-Oct-2014].
[45] A. Groisman and V. Steinberg, “Elastic Turbulence in a Polymer Solution Flow,” Nature, vol. 405, no. 6782, pp. 53–55, May 2000.
[46] B. Mena, G. Best, P. Bautista, and T. Sanchez, “Heat Transfer in Non-Newtonian Flow Through Pipes,” Rheol. Acta, vol. 17, no. 4, pp. 454–457, Jul. 1978.
[47] J. P. Hartnett and M. Kostic, “Heat Transfer to a Viscoelastic Fluid in Laminar Flow Through a Rectangular Channel,” Int. J. Heat Mass Transf., vol. 28, no. 6, pp. 1147–1155, Jun. 1985.
[48] S. J. Muller, R. G. Larson, and E. S. G. Shaqfeh, “A Purely Elastic Transition in Taylor-Couette Flow,” Rheol. Acta, vol. 28, no. 6, pp. 499–503, Nov. 1989.
[49] R. G. Larson, E. S. G. Shaqfeh, and S. J. Muller, “A Purely Elastic Instability in Taylor-Couette Flow,” J. Fluid Mech., vol. 218, pp. 573–600, 1990.
[50] E. S. G. Shaqfeh, “Purely Elastic Instabilities in Viscometric Flows,” Annu. Rev. Fluid Mech., vol. 28, no. 1, pp. 129–185, Jan. 1996.
[51] T. Burghelea, E. Segre, and V. Steinberg, “Elastic Turbulence in Von Karman Swirling Flow Between Two Disks,” Phys. Fluids 1994-Present, vol. 19, no. 5, 2007.
[52] F.-C. Li, H. Kinoshita, X.-B. Li, M. Oishi, T. Fujii, and M. Oshima, “Creation of Very-Low-Reynolds-Number Chaotic Fluid Motions in Microchannels Using Viscoelastic Surfactant Solution,” Exp. Therm. Fluid Sci., vol. 34, no. 1, pp. 20–27, Jan. 2010.
[53] W. Ling, S. J. Messana, and P. M. Rominski, “Thermal Design of Central Office Telecommunication Equipment,” Bell Labs Tech. J., vol. 14, no. 3, pp. 21–41, 2009.
[54] Telcordia, Generic Requirements NEBS GR-63-CORE, NEBS Requirements: Physical Protection, Telcordia Technologies, Inc., Piscataway, NJ, 2012.
[55] T. Graham, S. Kipp, and A. Romero, “Thermal Management at the Faceplate White Paper,” in Optical Internetworking Forum, 2012. [Online]. Available: http://www.oiforum.com/public/documents/OIF_PLUG_Thermal-01.0_WP.pdf. [Accessed: 10-Oct-2013].
[56] ETSI, EN 300 753 Equip. Eng.: Acoustic Noise Emitted by Telecommunications Equipment, Euro. Telecom. Stand. Inst., Sophia Antipolis, FR, 2009.

(Manuscript approved July 2014)

Authors

Ryan Enright is a member of technical staff in Bell Labs' Thermal Management Research program in Dublin, Ireland. He received B.Eng. (Hons) and Ph.D. degrees in mechanical engineering from the University of Limerick. Prior to his current posting, he worked as a research assistant at Bell Labs in Murray Hill, New Jersey during his doctoral work, as a CTVR postdoctoral associate at the University of Limerick, and a Marie Curie Fellow at the Massachusetts Institute of Technology. He joined Bell Labs Ireland in 2012 where he is currently the technical lead for the Thermally Integrated Photonics Systems (TIPS) program and a Researcher-in-Residence at the CRANN Institute, Trinity College Dublin. He is author or coauthor of more than 25 journal papers, 50 conference papers, and holds two granted patents with 12 patents pending.

Shenghui Lei is a member of technical staff in Bell Labs' Thermal Management Research program in Dublin, Ireland. He received B.S. and M.S. degrees in thermal engineering from Shandong University, and a Ph.D. degree in mechanical engineering from The University of Manchester. His studies focused on numerical simulations on multiphase reacting flows and radiative heat transfer. After a two-year post-doctoral position at the University of California, he joined Bell Labs Ireland in February 2013 where he has been performing industrially relevant research on thermoelectric cooling, micro-channel flows, and thermal analysis on III-V laser architecture. He has co-authored over ten journal articles

Kevin Nolan is a member of technical staff in Bell Labs' Thermal Management Research program in Dublin, Ireland. He graduated with a B.Eng. in aeronautical engineering and Ph.D. in fundamental fluid mechanics from the University of Limerick, Ireland. After completing a postdoctoral fellowship at the Stokes Institute, he worked on the development of optics and software solutions for next-generation qPCR at biomedical start-up Stokes Bio, where he developed optical tracking systems in microfluidic systems. He then spent two years working on extremely large Direct Numerical Simulation (DNS) databases at Imperial College London as a Marie Curie fellow. His current work focuses on microfluidics, particle image velocimetry, flow visualization techniques, and aero acoustics. He has published extensively in the Journal of Fluid Mechanics on topics related to structure identification in fluid flows from both experimental and numerical data sets.

Ian Mathews is a postdoctoral member of technical staff in Bell Labs' Efficient Energy Transfer Department in Dublin, Ireland. His current research focus is on thermally integrated photonics solutions and solar cells. Dr. Mathews completed his Ph.D. on InP-based multi-junction solar cells at the Tyndall National Institute, University College Cork, Ireland, where he also received his B.Eng. (1st) in civil and environmental engineering.

He has an M.Sc. (with distinction) in renewable energy systems technology from Loughborough University, UK. He has published six journal papers, three conference proceedings, 20 conference presentations, and is the author of one book chapter. He is a member of the IEEE.

Alexandre Shen is a member of technical staff at the III-V Lab in Marcoussis, France a joint laboratory of Alcatel-Lucent Bell Labs, Thales Research & Technologies, and the CEA LETI. He received an engineering degree and a master of science from the Institute of Optics Graduate School, Orsay, and a Ph.D. degree in electrical engineering from Université Paris XI, Orsay, France. He joined Alcatel-Alsthom Recherche in 1998 and led research on InP-based electro-absorption modulators and Mach-Zehnder modulators. He subsequently worked on saturable absorbers and mode-locked laser diodes. Within the III-V Lab and CEA LETI he is now working on Si-based hybrid lasers, other Si-based devices, and photonic integrated circuits. He is author or co-author of more than 13 published patents, 33 peer-reviewed journal articles, and almost 60 international conference papers, including four post-deadline papers.

Guillaume Levaufre is an intern at the III-V Lab in Marcoussis, France, where he is currently working toward his Ph.D. degree. He received a master's degree in engineering in optoelectronics from INSA Rennes, Rennes, France. His early work at the Lab involved inductively coupled plasma etching of InP. His current research is focused on hybrid III-V/Silicon photonic integrated circuits for high-speed telecommunications.

Ronan Frizzell is a member of technical staff at Bell Labs in Dublin, Ireland. He received a B.Eng in aeronautical engineering as well as a Ph.D. from the University of Limerick. The subject of his Ph.D. thesis centered on numerical damage modeling of composite materials and structures. Dr. Frizzell has authored or coauthored nine journal papers, 11 conference papers and has six patents pending. His research interests are in energy harvesting and thermal management, and he has experience in both experimental and numerical analysis techniques, with a particular emphasis on nonlinear finite element modelling.

Guang-Hua Duan is the leader of the Silicon Photonics Research Team within the III-V Lab, a joint laboratory of Alcatel-Lucent Bell Labs, Thales Research and Technology, and CEA LETI. He received a B.E. degree from Xidian University, Xi'an, China, as well as M.E. and Ph.D. degrees in applied physics from the Ecole Nationale Superieure des Telecommunications (Telecom-ParisTech), Paris, France. He came to France to direct research at Université Paris-Sud in 1995. He has authored or served as a coauthor of more than 100 journal papers, 200 conference papers, 25 patents, and as a contributor to three book chapters. He is also a visiting professor at the Ecole Supérieure d'Elecricité, Paris, France, and Ecole Supérieure d'Optique, Palaiseau, France, giving lectures in the fields of electromagnetism, optoelectronics, and laser physics.

Domhnaill Hernon is head of the Efficiency Energy Transfer (ηET) Department at Bell Labs in Dublin, Ireland. He graduated with a B.Eng in aeronautical engineering and a Ph.D. in fundamental fluid mechanics from the University of Limerick, Ireland. Dr. Hernon joined Bell Labs' Thermal Management Research Group in 2006 and became Technical Manager in 2010, leading a research team focused on advanced thermal management and energy harvesting solutions for next-generation telecommunications products. He established the Efficiency Energy Transfer (ηET) Department in 2013 and now manages a team of 16. Dr. Hernon plays an active role in building strong collaborations with multiple university partners to accelerate knowledge transfer across many disciplines ranging from thermal management on multiple scales (photonics, thermoelectrics and reliable active air cooling) to energy harvesting (vibrational energy harvesting for wireless sensors and alternative energy for off-grid power) to energy storage (novel battery architectures). He represented Alcatel-Lucent on the Industry Steering Board of the International Energy Research Centre (IERC) for two years and played a significant role in establishing governance, intellectual property, research, and business strategies for the center. He was named the 2013 Young Leader of the Year by The Irish Laboratory Awards. He has authored over 20 refereed publications and has three patents granted with ten pending. ▼