International Innovative Research Journal of Engineering and Technology
Vol. 8, Issue. 4, June 2023, pp. 17~24
ISSN: 2456-1983, DOI: 10.32595/iirjet.org/v8i4.2023.173

First Demonstration of High-Density Interconnects Integrated with Ultra Density Optical Technology

Raguvaran S[1], N.Gopinath[2]
[1]Department of Computational Intelligence, School of Computing, SRM Institute of Science and Technology, SRM Nagar, Kattankulathur 603203, Chennai, Tanil Nadu, India
[2]Department of Computational Intelligence, School of Computing, SRM Institute of Science and Technology, SRM Nagar, Kattankulathur 603203, Chennai, Tanil Nadu, India

Article Info

Article history:

Received Feb 07, 2023
Revised Apr 21, 2023
Accepted May 10, 2023

Keywords:

high density
holistic optics-electronics
optical transport
electrical interconnects
silicon optical modulators

ABSTRACT

The performance requirements of interconnects to and on silicon chips, both now and in the future is analyzed in this paper. To determine what optoelectronic and optical components are needed if optics is to address the main connection issues for next-generation high-performance silicon processors, we compare optical and electrical interconnects. Potential advantages of optics include timing, energy, and link density. As the information sectors see exponential development in performance, engineers are currently dealing with some technical challenges. Inter-chip connection bottlenecks are among the most critical problems. A novel "Photonics-Electronics Convergence System" idea is put forth by us. This technology allowed for the first demonstration of high-density optical connects integrated with a silicon optical waveguide, germanium photodetectors, silicon optical modulators, and a 13-channel arrayed laser diode on a single silicon substrate. Focusing on the optical transport and switching layer, we discuss elements of large-scale spatial multiplexing, vast optoelectronic arrays, and holistic optics-electronics-DSP integration, as well as optical node topologies for switching and multiplexing of spatial and spectral super-channels.

Corresponding Author:

Raguvaran S,
Department of Computational Intelligence,
School of Computing, SRM Institute of Science and Technology,
SRM Nagar, Kattankulathur 603203, Chennai, Tanil Nadu, India
Email: prof.sraguvaran@gmail.com

1. INTRODUCTION

Over time, communication technologies have developed to the point where they may be used across numerous industries. The primary characteristic of contemporary communication technology, which forms the basis of today's telecommunications services, is its capacity to transfer enormous amounts of data at very high speeds while maintaining excellent quality and low latency. From the ancient smoke-based pre-coded signal transmission systems of the prehistoric era to the smart mobile phones of the twenty-first century, which offer an unprecedented array of features and capabilities, we have witnessed remarkable progress [1]. To transmit pre-defined (coded) messages (such as celebrations, war, etc.) over line-of-sight (LoS) and short coverage areas, prehistoric man used smoke signs, firelights, and some sound sources. Longer transmission ranges of light-based (fire) communications were proved by the terrestrial signal towers found in the Great Wall of China and the coastal signal towers found throughout the Mediterranean region. These systems are still in use today.

The hydraulic semaphore was created in the fourth century BC by the Greeks, who also employed it for communication in the third century BC during the Punic War. Cyrus the Great is credited with

establishing the first pigeon messenger network in Assyria and Persia in the fifth century BC. They were transporting messages to rival factions in Mesopotamia circa 2000 BC. In the sixth century BC, posthouses provided services for horse messenger relays operated by the vast Persian Empire of Cyrus. The Greek writers Xenophon and Herodotus both gave positive descriptions of the method. The Chinese, Romans, and Egyptians brought new systems after these first advancements.

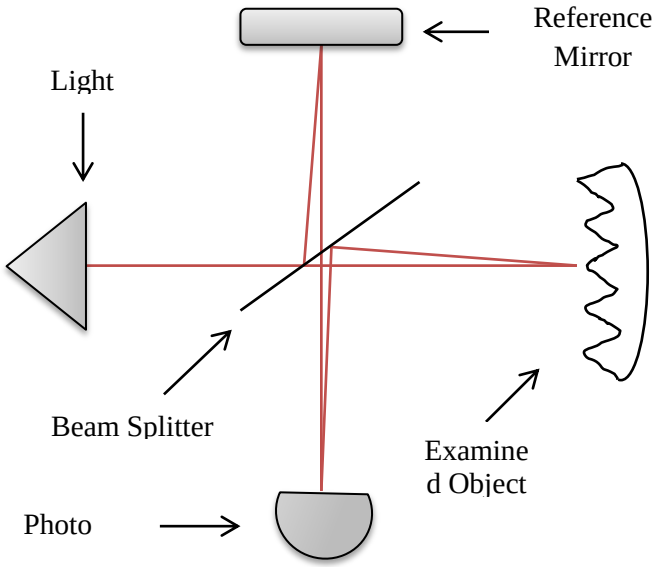

Figure 1. An example of Ultra Density Optical Technology

For more than 40 years, advancements in integrated circuit (IC) density and performance have propelled the semiconductor sector and the Information Revolution that followed. Up to now, evolutionary device scaling and/or a rise in chip size have been used to fulfil the expected entitlement to periodic gains in density and performance. Device performance increases with gate length, junction depth, and gate dielectric thickness scaling for classical transistors. On the other hand, if metal height is not lowered with conductor spacing, scaled chip wiring (interconnect) may experience higher capacitance in addition to increased resistance as a result of a decrease in conductor cross-sectional area. Thus, as feature size scales, RC parasitics contribute more and more to overall chip performance, as Figure 1 illustrates.

The following points to the rest of the document. Section 2 provides an overview of Ultra Density Optical Technology. In Section 3, we detail our suggested architecture for the fundamental Photonics-Electronics Convergence System. The installation of several gates is shown in Section 4, along with the experiment's findings, a comparison of our selected design, and a correlation with the information provided in Section 5.

2. RELATED WORKS

In our work, we demonstrate that every silicon photonic building component performs better than its electrical counterpart, enabling higher functionality with increased bandwidth and energy efficiency. Utilizing the widely available silicon photonic microring resonator, numerous devices that we have investigated are able to operate at high performance levels [4]. The microring resonator cavity improves many of the physical phenomena being studied in silicon photonics, resulting in silicon photonic devices that are far more efficient and compact. Furthermore, the required size of these microring resonators is further decreased by the significant refractive index contrast in these devices.

The chip pads need to be a specific physical size in order for bond wires or solder balls to link them to the chip packaging, which is why the number of pads is mechanically limited. The threshold voltage and supply voltage could not be dropped at the same rate in previous versions of CMOS technology. As a result, the transistor threshold voltage is nearly three times the supply voltage, which is the fundamental CMOS performance limit [5]. Since cooling is a major issue and is reaching its limit in large chip systems, the slight increase in chip electrical consumption is not surprising. The aggregate data rate can be obtained from the port count and the serial backplane data rate.

The performance of supercomputers has witnessed a steady development in the last decade and has generally kept up with the heavy calculation needs of advanced scientic studies. Thus far, scientists working on scientific computing have discovered methods to utilise the processing capacity provided to them in order to enhance the accuracy of their complicated system simulations. It follows that the desire for even more capabilities to propel research forward and/or pave new paths is understandable [6]. As demonstrated by the Top 500 trend lines, the hardware research community and industry have been successful in creating computer architectures and components that can provide a steady increase in performance and satisfy the needs of the majority of researchers.

The growing chip speeds, broader buses, and higher processor counts per system are causing the bandwidth and density requirements for interconnects within high-performance computing systems to expand quickly. The latest iteration of the International Technology Roadmap for Semiconductors highlights a few of these developments [7]. These developments indicate that off-chip or off-module bandwidth will become a more significant limitation for many high-performance devices or modules. Larger caches and software are predicted to provide some respite, but in order to maintain the price-performance trends that computer systems and other high-end systems have demonstrated over time, technologies that offer better chip-to-chip or module-to-module interconnections will become more and more necessary.

Large systems are by definition built from many smaller components, hence one of the most important problems facing any system designer is coming up with cost-effective methods for connectivity and packaging. The integration and interconnection approach of the system designer frequently determines the competitive advantage of a particular system [8]. Consider, for instance, how the quantity of electrical IO signal lines influences packaging decisions made during the design of a large system, such the length scale of a telephone switch's interconnection, or the dominating interconnection technologies in a portion of a large, high-performance system.

Heterogeneous semiconductor technologies are needed for many product applications due to system complexity and manufacturing process integration requirements, which preclude the use of a solution on a single chip. In applications that demand miniaturisation and where extremely high-I/O interconnections are necessary to link these heterogeneous semiconductor technologies, conventional packaging technologies like ceramic or organic chip carriers might not be able to meet the I/O interconnection level requirements [9]. Advanced wire-bond chip stack packages and ceramic or organic chip carriers are examples of traditional interconnection and packaging that may start to be replaced by new packaging techniques like silicon through-via chip stacks and packages.

The LAN industry has created standards for the links that join computers to storage devices and networks. Ethernet has progressed from 10-Mbps electrical specifications to 1-Gbps gigabit Ethernet specifications, encompassing both optical and electrical connections. The 10-Gbit Ethernet working group is already debating a standard, despite the later specification having just been finished [10]. Many suppliers are currently considering offering a variety of options, such as a four-channel, 2.5 GHz per channel “coarse” wavelength-division multiplexed system, four parallel 2.5 GHz channels, and a single-line 10-Gbps serializer/deserializer (SerDes).

3. METHODS AND MATERIALS

3.1 Optical communication performances: 100 gb/s-m distance–bandwidth product

The use of optical data lines in communications systems has grown over the previous 30 years. The commercial penetration of optics is plotted against bandwidth and link distance in Figure 2. Over the past few decades, optical networks have attained bandwidth-reach product performance exceeding 106 Gb/s-m with single-mode solutions and 103 Gb/s-m with multi-mode fibre systems [11]. This approximate trend is observed across five orders of magnitude of bit rates and distance.

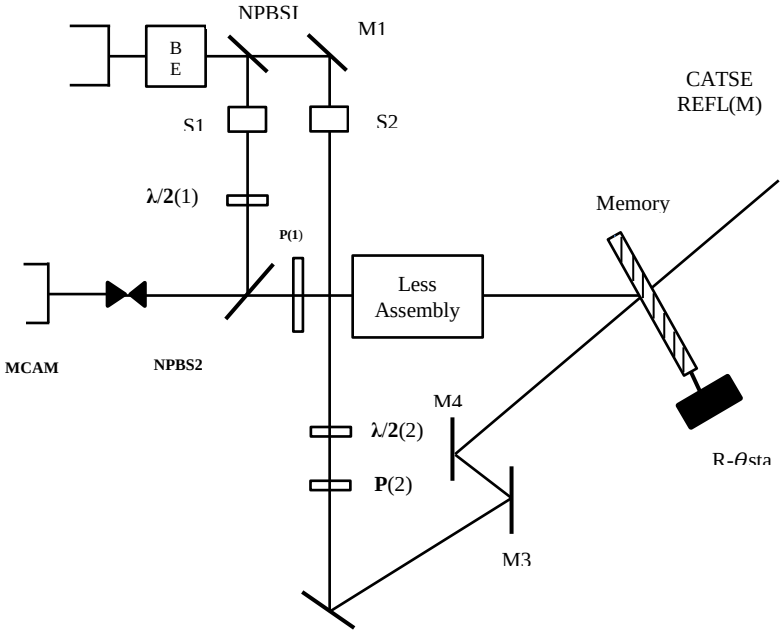

Figure 2. Plotting of the Commercial Penetration of Optics

On the electrical side, 10 Mb/s-m performance trend has mainly been followed by copper-link technologies, such as asymmetric DSL [12], digital subscriber line (DSL), and other twisted pair solutions. This data supports a general conclusion that the bandwidth-distance product of 100 Gb/s-m can be used to indicate the minimum successful technological crossover from electrical to optical lines from a commercial standpoint. A few early long-range electrical coaxial solutions from 100 Kb/s to 1 Mb/s that were implemented in the 1970s are among the rare exceptions to this general tendency. On the optical side, exceptions include applications in automotive and avionics, where adoption may be constrained by factors like susceptibility, electromagnetic interference emissions, or weight of the optic cables.

By extending this trend line towards greater bandwidths and shorter distances, we are getting closer to the conventional sweet spot for serial intersystem electrical links that are commodity-based.

Recently, active optical cables—a new type of transceiver—were released. Because they have original equipment components at both ends, the user can connect an active optical cable to an already-existing socket meant for an electrical transceiver. The connection to the socket supplies power for the optical transceiver. In addition to allowing new optical technologies to compete with current optical and electrical technologies in a way that is transparent to the user, this also removes the need to modify the system for an optical component. Fibre solutions are increasingly integrated into active cables as we get closer to the sub-10 m distance range. This facilitates system adoption because the electrical and optical cables have the same connection interface and signalling system.

If current trends continue, it is anticipated that over the next five years, optical linkages would reach directly into chip-scale packages on printed circuit boards (PCBs), signifying link distances of less than one metre. We shall endorse the idea that the necessary technologies are already well under development, despite the fact that some may perceive this forecast as a drastic divergence from accepted design practices. Optical communications, which have a power dissipation advantage over electrical systems, are crucial to this case. In order to develop a broad guideline for the adoption of optical interconnects, we explore a bit-rate-independent metric in the next section, even though individual technology selections will rely on bit rates, loss metrics [13], and application performance needs.

3.2 Photonics' options in an extremely short-reach interconnects: the metric for bit-energy per unit distance

Electrical links are the most common type of ultrashort reach interconnects. This is mainly because numerous electrical interconnect drivers and receivers can be designed and constructed on a silicon CMOSVLSI chip quite easily. Furthermore, the signal attenuation and other effects of electrical channels that

typically have limited signal bandwidth on PCBs have been overcome with the development of serial-link circuit techniques.

Therefore, as VLSI technologies have continued to scale to larger silicon wafers and finer lithography, the cost per gigabit per second has decreased while the density and bit rates of such high-speed electrical signalling have increased. Power delivery is a new design driver that we must deal with nowadays.

Thermal and cooling considerations, as well as electrical power distribution, were traditionally left until the very end of the system design cycle. Thus, in a highly restricted context, thermal and electrical power engineers had to come up with workable solutions for cooling and power supply, with little chance of changing the logical architecture of the system or having an impact on the design of the microchips or system packaging. This is evolving. With the introduction of 100 W+ super processors, a single-chip socket on a PCB needs to receive more than 100 A of current with minimal resistive losses and inductance. Therefore, it is crucial to incorporate power delivery and thermal considerations early on in the chip design process. Comparably, at the system level, thermal issues become crucial to the system de signer when equipment rack power densities surpass 10kW/ft2. Ultimately, consumers and end users must become aware of the entire system power dissipation as power consumption increases and total computing capacities become power constrained.

Two features of electrical link design are responsible for the majority of the power expenditures when using them at short distances (one metre and less). Initially, many methods are utilized by electrical circuits to surmount channel losses and transmit non uniformities, which deteriorate with connection distance and bitrate. These techniques, which have a significant impact on power dissipation, design complexity, and silicon area, include transmit reemphasis, receiver equalization, and possibly sending multiple bits per hertz. Better dielectrics, blind vias, or better cabling are examples of material advancements that help electrical-link performance, but they also come with unwanted cost increases.

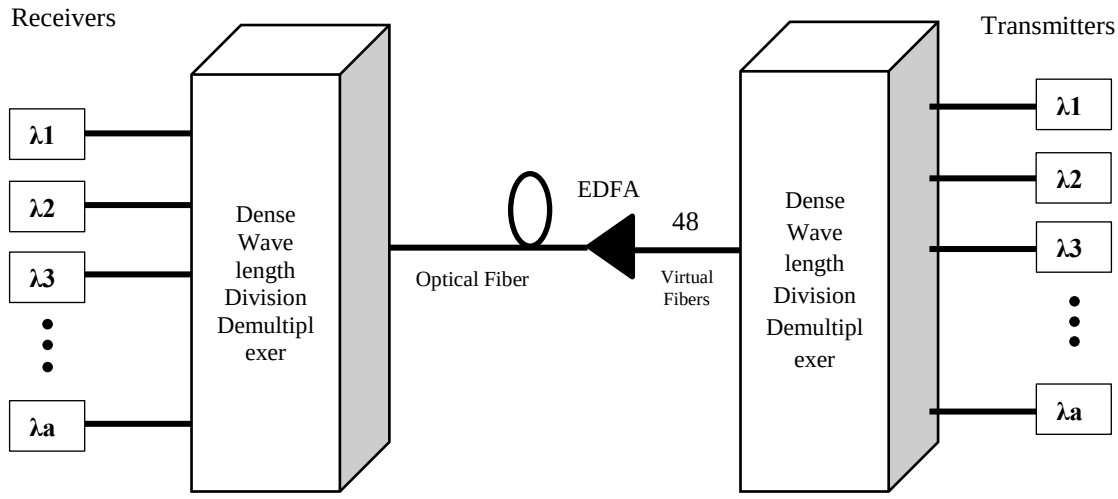

Figure 3. Block Schematic for a DWDM Network

As previously stated, the fundamental carrier in optical networks is Dense Wavelength Multiplexing. Transmitters, receivers, erbium-doped fibre amplifiers, DWDM multiplexors, and DWDM demultiplexors are the essential parts of every DWDM system. The general architecture of a DWDM system is shown in Figure 3 [14].

4. IMPLEMENTATION AND EXPERIMENTAL RESULTS

4.1 Silicon DFB Laser Hybrids

The DFB laser is composed of a quarter wavelength shifted hybrid silicon grating that is 360 μm long, with a grating κ of around 247 cm−1, and a reflectivity peak at approximately 1600 nm. The device layout is displayed in Figure 4. The laser's gain zone is 200 μm long and has a cross section of There are 80 μm-long tapes that connect the passive silicon waveguide to the gain area. According to Figure 4 [15], they are created by linearly shrinking the III–V mesa area above the silicon waveguide. By doing this, the hybrid

waveguide mode is adiabatically transformed into a passive silicon waveguide, allowing for losses of about 1.2 dB per taper and reflections of about 6 × 10−4.

The placement of hybrid silicon light detectors allows for on-chip assessment of the DFB laser performance on both sides of the laser. The two 80-µm-long tapers are part of the 240 µm-long photo detectors. For off-chip spectrum measurements, the detector on the right is positioned 400 µm away to provide space for dicing and polishing. You can get more information. By using integrated photo detectors to gather light from both sides of the laser, the light-current (LI) properties of the DFB laser are monitored on-chip. To estimate the laser power output, the photo detectors' 100% internal quantum efficiency is taken as a conservative estimate of the laser's performance.

That the maximum output power of 4.3 mW and the lasing threshold of 25 mA is both present at 10◦C. This translates to a 1.4 kA/cm2 threshold current density. 50 ◦C is the maximal lasing temperature. The resistance of the laser's device series is 13 Ω. This number is in accordance with the resistance of 4.5 Ω, as measured on FP lasers that are 800 µm in length. By cutting off the appropriate photodetector, polishing, and antireflection coating the silicon wave guide output facet, the lasing spectrum is obtained. Using a lensed fibre, light is captured and fed into a spectrum analyser with a resolution bandwidth of 0.08 nm. The optical spectrum at 30 mA injections current is displayed. The laser features a side-mode suppression ratio of 50 dB and a lasing peak at 1599.3 nm. It is evident that the laser runs in single mode for a span of 100 nm. The delayed self-heterodyne method is used to measure the laser linewidth. With a convoluted Lorentzian linewidth of 7.16 MHz, or 3.6 MHz, which is a typical figure for commercial DFB lasers, a minimum linewidth of 1.8 mW is measured.

4.2 Bragg Reflector Lasers Dispersed

Schematically represented, the distributed Bragg reflector (DBR) laser is comprised of two passive Bragg reflector mirrors spaced 600 µm apart, creating an optical cavity. With an etch depth of 25 nm and a duty cycle of 75%, the gratings have a grating strength of 80 cm−1. With a front mirror length of 100 µm and a back mirror length of 300 µm, the power reflectivity is 44% and 97%, respectively. Two 80-µm-long tapers and a 440-µm-long silicon evanescent gain area are positioned inside the cavity. To reduce absorption, the tapers are electrically pushed in tandem with the gain region. More information is available.

Using an integrating sphere at the laser's front mirror, the output power of a continuous wave laser is determined. The LI attribute of the front mirror output. The apparatus has a differential quantum efficiency of 15% due to its lasing threshold of 65 mA and maximum front mirror output power of 11 mW.

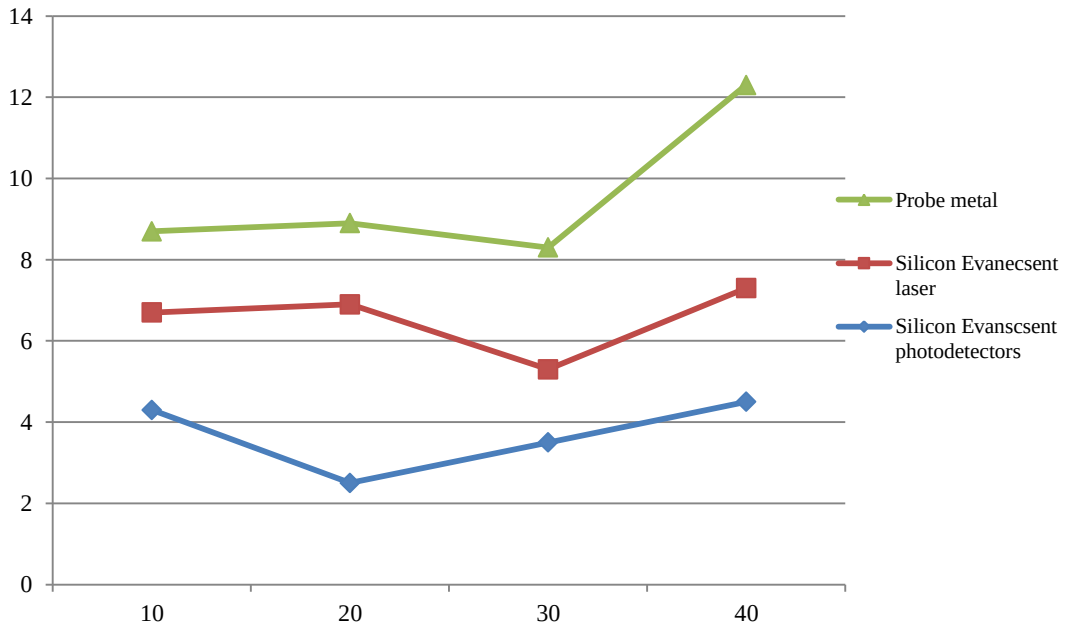

Figure 4. Silicon DFB hybrid device architecture, a microscope picture of the DFB laser that has built-in photodetectors

Numerous crucial laser parameters, including wall plug efficiency and resonance frequency, are impacted by the threshold current, which is mostly dependent on the taper transmission loss.

4.2 Recoveries

Optoelectronic photo receivers with higher performance and greater practicality can be achieved by integrating hybrid silicon amplifiers and photo detectors onto a single substrate. Because III-V QW engineering makes it simple to extend the absorption edge of hybrid silicon photo detectors beyond the 1600 nm range, these devices are intriguing.

Both the amplifier and the detector employ the same III–V epitaxial structure. The detector and amplifier have total lengths of 100 μm and 1240 μm, respectively. The greatest gain for a structure 1.2 mm long at 300 mA is 9.5 dB. The 100-μm detector has 50% quantum efficiency. When the two are combined, the receiver's responsivity rises to 5.7 A/W with preamplification. The gadget exhibits 0.5 dB saturation at 25 mA of photocurrent.

Time-domain impulse response is used to measure the device bandwidth, which is 3 GHz (see Figure 5; frequency is shown in the inset). However, the projected bandwidth that is limited by resistance and capacitance is 7.5 GHz. This suggests that the current III-V layer design is the main factor limiting the device speed. Using QWs with a smaller valence band offset and smaller SCH layer to shorten the whole transit time will result in a wider bandwidth. A nonreturn to zero 2.5 GB/s PRBS with BER measurements reveals a receiver sensitivity of −17.5 dBm.

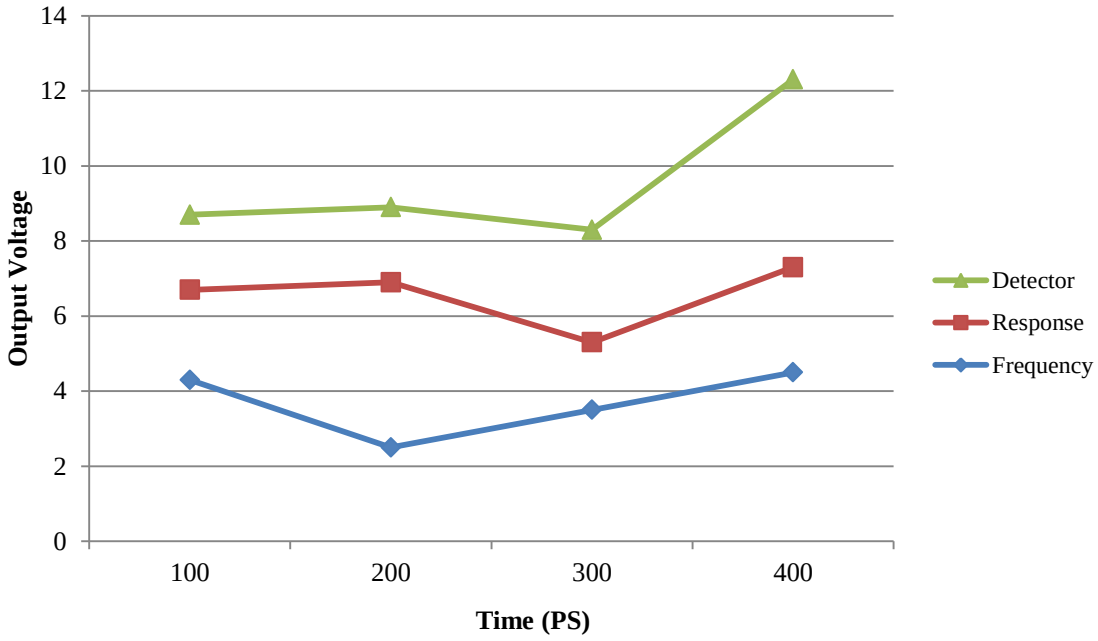

Figure 5. Impulse response of the detector (0.6-ps input pulse). The impulse response's Fourier transform is displayed in the inset

It was recently revealed that an enhanced bandwidth might be achieved with an InGaAs/InP hybrid silicon photo detector. Open eye diagrams were produced up to 12.5 GHz, and the bandwidth was measured at 6 GHz, higher speeds can be attained by carefully planning the photodiode. InP/InGaAs-based wafer-bonded photo detectors with a 20–25 GHz bandwidth, for instance, have been described, demonstrating the potential of the hybrid silicon technique.

4.3 Complete Integration System

The devices displayed here, which include sources, modulators, and receivers, operate around 1.55 μm wavelengths and exhibit distinct III/V QW bandgaps. Nonetheless, a complete transmitter will most likely be made up of a MZM or EAM modulator in addition to a single-frequency source. Additionally, sources, modulators, and detectors must all be integrated on the same device in order to support on-chip interconnects.

Different bandgap QWs can be implemented on the same chip in a variety of ways. Regrowth techniques are one approach; these involve etching away a portion of the QWs and regrowing the portions with a different or absent QW. Wafer bonding techniques required by the hybrid silicon platform most likely do not suit this choice because of the nonflat surface topology at the butt-joint interfaces. Bonding distinct dies with various bandgaps on a single SOI wafer is an additional choice. This technique, however, limits the design flexibility because it requires the placement of components that have to have the same bandgap.

This section presents a method based on QW intermixing (QWI). One possible use of this technology is the integration of a modulator and a source. Displayed is an InGaAsP/InP EAM combined with a sampled-grating (SG) DBR laser. Information about this work is available.

5. CONCLUSION

Our whole hierarchy of interconnects for computing machines for the next ten years is now available thanks to this investigation. Optical interconnects are already present outside of the box and are not going away. Strangely, the most promising technology for connecting boards within the system over the whole board surface is free-space optical interconnects, which are still not commercialized. It has been proved that any available technology can essentially accomplish this. All that is needed is a determined effort on the part of the industry to create products and a standard for their use.

There has been discussion on high-speed modulators, and an effective demonstration of an EAM modulator operating at 25 Gb/s has been provided. On silicon photonic circuits, where modulation speeds are usually constrained, this finding enables high-speed optical communications. While the initial hybrid silicon receiver realizations were limited to 6 GHz in speed, they demonstrated a high sensitivity of −17.5 dBm, which is beneficial for power-efficient communications.

Finally, we have demonstrated that these many components can be merged on a single silicon chip using technologies like QWI. By combining an EAM and an SG-DBR laser to create a complete transmitter that can be modulated at speeds more than 2 Gb/s, QWI has been successfully demonstrated.

In conclusion, hybrid silicon technology provides an entire range of optical connection components that can theoretically all be implemented on a single chip. Furthermore, the fabrication technology works with the CMOS processor that is available today. In summary, this platform shows great promise for on-chip and off-chip interconnects in the future and could offer a more energy-efficient option for high-speed and high-bandwidth systems.

REFERENCES

[1] Raj, A. A. B., Krishnan, P., Darusalam, U., Kaddoum, G., Ghassemlooy, Z., Abadi, M. M., ... & Ijaz, M. (2023). A review–unguided optical communications: Developments, technology evolution, and challenges. *Electronics*, *12*(8), 1922.

[2] Urino, Y., Shimizu, T., Okano, M., Hatori, N., Ishizaka, M., Yamamoto, T., ... & Arakawa, Y. (2011). First demonstration of high density optical interconnects integrated with lasers, optical modulators, and photodetectors on single silicon substrate. Optics express, 19(26), B159-B165.

[3] Meyer, J. U., Stieglitz, T., Scholz, O., Haberer, W., & Beutel, H. (2001). High density interconnects and flexible hybrid assemblies for active biomedical implants. IEEE Transactions on Advanced Packaging, 24(3), 366-374.

[4] Biberman, A., & Bergman, K. (2012). Optical interconnection networks for high-performance computing systems. Reports on Progress in Physics, 75(4), 046402.

[5] Kromer, C., Sialm, G., Berger, C., Morf, T., Schmatz, M. L., Ellinger, F., ... & Jackel, H. (2005). A 100-mW 4/spl times/10 Gb/s transceiver in 80-nm CMOS for high-density optical interconnects. IEEE Journal of Solid-State Circuits, 40(12), 2667-2679.

[6] Rumley, S., Bahadori, M., Polster, R., Hammond, S. D., Calhoun, D. M., Wen, K., ... & Bergman, K. (2017). Optical interconnects for extreme scale computing systems. Parallel Computing, 64, 65-80.

[7] Schares, L., Kash, J. A., Doany, F. E., Schow, C. L., Schuster, C., Kuchta, D. M., ... & Dolfi, D. W. (2006). Terabus: Terabit/second-class card-level optical interconnect technologies. IEEE Journal of Selected Topics in Quantum Electronics, 12(5), 1032-1044.

[8] Levi, A. F. (2000). Optical interconnects in systems. Proceedings of the IEEE, 88(6), 750-757.

[9] Knickerbocker, J. U., Andry, P. S., Buchwalter, L. P., Deutsch, A., Horton, R. R., Jenkins, K. A., ... & Dang, B. (2005). Development of next-generation system-on-package (SOP) technology based on silicon carriers with fine-pitch chip interconnection. IBM Journal of Research and Development, 49(4.5), 725-753.

[10] Lytel, R., Davidson, H. L., Nettleton, N., & Sze, T. (2000). Optical interconnections within modern high-performance computing systems. Proceedings of the IEEE, 88(6), 758-763.

[11] Mekawey, H., Elsayed, M., Ismail, Y., & Swillam, M. A. (2022). Optical interconnects finally seeing the light in silicon photonics: Past the hype. Nanomaterials, 12(3), 485.

[12] Benner, A. F., Ignatowski, M., Kash, J. A., Kuchta, D. M., & Ritter, M. B. (2005). Exploitation of optical interconnects in future server architectures. IBM Journal of research and development, 49(4.5), 755-775.

[13] Thraskias, C. A., Lallas, E. N., Neumann, N., Schares, L., Offrein, B. J., Henker, R., ... & Tomkos, I. (2018). Survey of photonic and plasmonic interconnect technologies for intra-datacenter and high-performance computing communications. IEEE Communications Surveys & Tutorials, 20(4), 2758-2783.

[14] Weninger, D., Serna, S., Jain, A., Kimerling, L., & Agarwal, A. (2023). High density vertical optical interconnects for passive assembly. Optics Express, 31(2), 2816-2832.

[15] Pang, X., Ozolins, O., Lin, R., Zhang, L., Udalcovs, A., Xue, L., ... & Chen, J. (2020). 200 Gbps/lane IM/DD technologies for short reach optical interconnects. Journal of Lightwave Technology, 38(2), 492-503.