Silicon Photonics

Keren Bergman, Columbia University

Ray Beausoleil and Dejan Milojicic, Hewlett Packard Labs

Photonics has been used for computer communication at larger distances. Due to performance, power, and packaging, photonics holds promise for moving further to and onto chips. This is called silicon photonics.

DEJAN MILOJICIC: What does silicon photonics (SiPh) mean to you?

KEREN BERGMAN: It's tremendously challenging to integrate photonics on a large scale. Photonic technology primarily deployed in the telecommunications (telecom) industry consists of few integrated components, typically lasers, modulators, and receivers. SiPh is ushering in a new era where we can think of integrating hundreds to thousands of devices into complex photonic-integrated circuits with a design and fabrication scalability similar to CMOS. SiPh is opening up new applications at the intersection of photonics and computing.

RAY BEAUSOLEIL: The push to build bigger, better, and faster computers has accelerated over the past decade due to applications like big data and machine learning. In high-performance computing (HPC), the push to exascale performance has put a strain on existing data communications technologies. This pressure will get worse as the industry implements next-generation heterogeneous architectures for advanced artificial intelligence (AI) workloads. I don't see any route to affordable interconnects for these applications without SiPh.

MILOJICIC: I have heard the phrase from Norm Jouppi (who might have heard it from others) that "ions are best for storing information, electrons for computing, and photons for communication." Photonics in the form of optical networks have been used at larger distances, while electrons in the form of packet-switching interconnects have been resisting the adoption of photonics. Has the time arrived for this to change?

BERGMAN: I agree. Fundamentally, communications with photons are advantageous in terms of bandwidth density: we can pack multiple wavelength channels in a single low-loss waveguide/fiber. The key question

Digital Object Identifier 10.1109/MC.2022.3148491
Date of current version: 8 April 2022

 0018-9162/22©2022IEEE

EDITOR **DEJAN MILOJICIC**
Hewlett Packard Labs;
dejan.milojicic@hpe.com

is, where is the optimal place to put the interface between the electronic and optical domains? SiPh is pushing that interface closer to the chip, and perhaps on chip in the future. Electronic interconnects are superior for ultrashort intrachip connectivity, but because energy costs are too high, they cannot drive that bandwidth off chip. The ultimate goal is for SiPh to deliver "chip-escape" bandwidth and carry the data's movement across the system. Packet switching is another matter. Photonics is essentially a circuit switch technology. This is because some form of buffering/memory is needed to perform packet switching. Optics is not a great technology for buffering or memory because the photons do not like to stay still. Besides, driving optical switching at high speeds is expensive from an energy standpoint. But optical circuit switching can play an important role in system design, together with electronic packet switching. We are developing new photonic architectures that take advantage of this and provide—through optical circuit switched paths—direct, high-bandwidth connectivity across the system. Emerging system architectures with disaggregation can take advantage of flexible optical connectivity.

BEAUSOLEIL: I like Keren's answer so much that instead of answering your question, I'll challenge the claim that electrons are always best for computing. We've published a few papers recently that propose photonic accelerators for some machine learning and optimization applications. I think that this is especially true once the data are already available in an optical format in the interconnect.

MILOJICIC: SiPh has been gaining substantial popularity in the past 10 or more years. What are some of key technological enablers that may influence the adoption of SiPh?

BERGMAN: The relatively high cost of SiPh is dominated by assembly and packaging. Laser sources for multiple wavelength channels are also a key issue. Comb lasers that can generate tens or even hundreds of wavelength channels from a single source are perhaps one of the most promising game changers for enabling SiPh interconnects.

BEAUSOLEIL: Keren's right about cost being the limiting factor at this point in time. Even with new technologies like comb lasers, assembly and packaging won't drop significantly in cost until there's an open ecosystem for SiPh that's similar to that of CMOS. Historically, SiPh has been driven by companies that deploy a vertically integrated business model where they exclusively own much of the intellectual property involved in the design and assembly of SiPh components. With its steep entrance barrier, the vertical model stifles innovation and tends to favor established players in the space. A horizontal business model for SiPh similar to the foundry and outsourced assembly and test model that has been common in the semiconductor industry would make a huge difference.

MILOJICIC: In terms of speeds, feeds, and power, how does SiPh compare to traditional communication based on electrons?

BERGMAN: The key is bandwidth density and the fact that we can place multiple wavelength channels, each modulated with high-speed data, all onto the same "wire" or waveguide. For example, a 64-channel link with each wavelength modulated at 25 Gb/s can carry 1.6 Tb/s. The modulation is performed by electronic drivers and the receiver is electronic, so the energy consumption scales with the data rates per channel. The current research designs are targeting a photonic -link energy consumption of ~1 pJ/b.

BEAUSOLEIL: I think that speeds and feeds, which are specific to the point-to-point connections that interconnects use today, sort of disguise the real potential of multiple-wavelength SiPh. Wavelength is another knob that we could use to change the system-level network architecture for new machine learning workloads. We are researching these now.

MILOJICIC: How would SiPh impact the reliability of components and systems as a whole. Some previous optical solutions required substantial temperature calibration. What is the situation with SiPh?

BERGMAN: With SiPh and CMOS-based wafer-scale manufacturing, the defect rate has been substantially reduced and yields are improving. The laser sources present reliability challenges for co-integration with SiPh and are separately packaged.

BEAUSOLEIL: Keren's right. At this point, SiPh components themselves are comparatively reliable: most of the implementations of SiPh dies use good old-fashioned, 65-nm, design-rule process flows. The laser is where all of the reliability studies are focused now.

MILOJICIC: Security is traditionally linked to communication, especially the encryption of data in flight. Does SiPh open any new opportunities to make systems and data more secure?

BERGMAN: Secure data communication and quantum encryption

demonstrated in fiber optics can definitely be incorporated into SiPh links. I can see some possibilities at the system level but using the wavelength domain to create secure communication subnetworks.

BEAUSOLEIL: As I said before, once the data are optical, then optical computing becomes easier. I have a few crackpot ideas about how to do optical-layer encryption and decryption of the data at the link endpoints that might be worth pursuing someday.

MILOJICIC: Would SiPh affect packaging and production of overall systems, and how?

BERGMAN: Absolutely, packaging is the key challenge to realizing SiPh at the system level. We've solved a great deal of the chip-level fabrication problems, and now it's all about packaging, packaging, packaging. At the moment, it's the "Wild West," with numerous approaches.

BEAUSOLEIL: Well, at the system level, I don't think the impact of packaging (... packaging, packaging ... is there an echo in here?) will be enormous, if we can perfect blind-mate, single-mode, fiber-array connectors at faceplates. But this is a pretty big challenge.

MILOJICIC: How will SiPh affect other components, such as computation or memories? For example, would it affect endpoints such as memories and computational cores (CPUs or accelerators)? Is there an opportunity for influence in another direction, that is, if the design of memories and compute change, could that affect the implementation and deployment of SiPh?

BERGMAN: SiPh can break open the chip input/output (IO) bottleneck. Today we have CPUs and accelerators on an interposer with memory to provide needed high-bandwidth GPU/memory communications. With SiPh, we can bring the same (and higher) high-bandwidth communication beyond the interposer and across the system, with the same low energy consumption. This could mean that instead of fighting for room on the number of high-bandwidth memory modules we can fit on the interposer next to the GPUs, we can imagine a much larger number of GPUs, CPUs, and memory all interconnected using a photonic fabric, and no longer limited by the size of the interposer. The system architecture can have a "flatter" communication because the cost of moving data on the chip is equivalent to moving data anywhere in the system. Of course, latency still plays a role as time of flight will always be limited by the speed of light.

BEAUSOLEIL: I would love to see someone build an on-chip dense wavelength division multiplexing interconnect for many-core CPUs! This is exactly what we proposed at the International Symposium on Computer Architecture back in 2008; we called it *Corona*, after the sun (not the beer). We also invented a cool on-chip computer that could resolve network collisions in approximately eight 3-GHz clock cycles. Perhaps I'm "getting high on my own supply" here, but I think that driving SiPh down to the millimeter scale would have a profound effect on chip design.

MILOJICIC: What are use cases and traffic patterns that can benefit the most from SiPh? Are these related to large bandwidth or short latency, or anywhere in between?

BERGMAN: SiPh interconnects will bring the most benefit to the traffic patterns that require large bandwidths with relatively steady connectivity. The best cases are where we can set up needed connectivity among the computing resources with transparent optical paths through circuit switches and then let the workloads run.

BEAUSOLEIL: Back when I was just a kid (so about 10 years ago), I wrote a paper with Moray McLaren and Jouppi that looked at the impact of SiPh and the Hyper-X network topology on three very different workloads: random-access memory (giga-updates per second), MapReduce, and multidimensional fast Fourier transform. We saw significant improvements in performance across all three.

MILOJICIC: Are there specific topologies that could benefit more from SiPh; for example, all-to-all architectures? If so, would that affect programming models at all, or it would be a completely transparent change. Even if it could be a transparent change, if it is accounted for, could programmers benefit from it?

BERGMAN: The SiPh interfaces do not depend on topology. They will provide 10–100 times the bandwidth to/from the system's endpoints and the electronic switch fabrics. From a programmer perspective, the benefits would come from rethinking the scaling and placement of the workloads. If we think of the available internode and memory bandwidths in terms of bytes/flop, this ratio would increase by 10–100 times. Beyond this, we can think of incorporating optical switching and wavelength-selective functions in the system's interconnection network. This would provide further capabilities, for example, of composable systems within disaggregated architectures, and of using some optical data-movement functionalities like wavelength multicasting. This research is currently being explored together with the co-design of new programming models that can take advantage of these photonic architectures.

BEAUSOLEIL: When we were working on The Machine at Hewlett Packard Labs (a radically memory-driven computer architecture), I spent quite a bit of time talking to programmers about the impact of photonic

interconnects on bandwidth, latency, and so on. It was pretty difficult going at first because their mental models were founded on old-fashioned interconnects. Eventually, I proposed that they pretend the network had infinite bandwidth per channel, zero collisions, and zero latency. Of course, this is ridiculous, but this framework allowed them to re-evaluate their previous assumptions, and we started to make progress. It was a lot of fun, and I learned a lot from them all.

MILOJICIC: Where do you see the largest business impact of SiPh?

BERGMAN: In scaling AI and data analytics systems.

BEAUSOLEIL: At this point in time, I agree.

MILOJICIC: Are there any special industry verticals that could benefit from SiPh, for example, HPC?

BERGMAN: The explosion of machine learning and data analytics is transforming HPC and data centers. This sector is at the forefront of benefiting from SiPh.

BEAUSOLEIL: Anywhere data are communicated over distances of 10 m will ultimately use SiPh. Take your pick: automotive, aerospace, smart buildings, … even telecom. As prices drop, everyone will adopt.

MILOJICIC: Is there a unified approach to SiPh, or are there thousands of flowers blooming, usually resulting in a need to standardize? If so, what are these standards and standards bodies, or is it too early to even think about it? They say that standardization needs to come at just the right time; not too early, not too late?

BERGMAN: There are definitely standards bodies, from both the Ph industry and system vendors (open compute and so on), but it is still early and there are multiple approaches. This is the right time, and it is needed for standardization to evolve further, especially in narrowing down the packaging.

BEAUSOLEIL: Well, my music player is set to repeat here: there won't be any relevant standards bodies in a vertical industry. So far, I think that the attempts to standardize photonics have been thinly veiled to squeeze margins out of suppliers.

MILOJICIC: Open source has taken the whole software community by storm, and as of recently, there is a movement toward open compute. Is there an equivalent open SiPh?

BERGMAN: It is still a siloed industry, with individual vendors developing proprietary technologies.

BEAUSOLEIL: Yes, see my previous comments. We need to push an open ecosystem that supports a horizontal SiPh industry.

MILOJICIC: We are seeing a lot of changes recently with the end of Moore's law, a lot of heterogeneity introduced by specialized accelerators, researchers revisiting quantum and/or quantum-like computing, and so forth. Do you think that SiPh is here to stay, or perhaps a new type of computing (and memory, for that matter) may require even more sophisticated interconnects? In other words, is there anything else, even beyond the horizon of SiPh?

BERGMAN: Yes, absolutely. SiPh breaks open the chip IO bottleneck—enabling ultrahigh-bandwidth connectivity that we now have at the chip—to go systemwide. This is a game-changer, especially for heterogeneous systems. SiPh can enable "deeply disaggregated" systems and an interconnect fabric that provides immense communication scalability.

BEAUSOLEIL: I think Keren is right. SiPh will be critical for enabling disaggregated compute architectures that place memory and memory multiplexing at the center of the universe, instead of CPUs. This also means that board-based hardware accelerators that can be added to racks as needed will become popular in HPC installations that support a wide variety of workloads. I'm still waiting for nonvolatile memory; when we get it, it will be huge and absolutely require SiPh.

MILOJICIC: Any closing thoughts on the future of accelerators?

BERGMAN: We are at an inflection point. SiPh is poised to provide the flexible connectivity and high-bandwidth communications exactly needed for the new generation of heterogeneous systems.

BEAUSOLEIL: I don't think that a general-purpose zetaflop machine based on CMOS will ever be built. Accelerators (optical, neuromorphic, quantum, and so forth) will be the only way to achieve zetascale performance on particularly important (but relatively narrow) problems.

KEREN BERGMAN is a professor at Columbia University, New York, 10027, USA. Contact her at bergman@ee.columbia.edu.

RAY BEAUSOLEIL is a senior fellow at Hewlett Packard Labs, Milpitas, California, 95035, USA. Contact him at ray.beausoleil@hpe.com.

DEJAN MILOJICIC is a distinguished technologist at Hewlett Packard Labs, Palo Alto, California, 94306, USA. Contact him at dejan.milojicic@hpe.com.