Implementing System-in-Package with Nanophotonic Interconnect

Mark Cianchetti,[*] Nicolás Sherwood-Droz,[†] Christopher Batten[*]
[*]Computer Systems Laboratory   [†]Cornell Nanophotonics Group
School of Electrical and Computer Engineering, Cornell University, Ithaca, NY
{mjc96,nrs35,cbatten}@cornell.edu

There has been significant interest in nanophotonics for global on-chip communication and for inter-socket communication between processors and/or main memory. Our goal in this short paper is to motivate future research on *nanophotonic system-in-package* (NSiP): an integration strategy that uses CMOS-compatible nanophotonic devices to implement an efficient package-level network for high-performance SiPs.

1. Nanophotonic System-in-Package (NSiP)

An NSiP is composed of *nanophotonic chiplets* that communicate using tightly integrated nanophotonic devices. We use the term *nanophotonic chiplet* to emphasize that these components are specifically engineered for an NSiP as opposed to standard electrical chips that might use wire or flip-chip bonding and find applications as discrete parts. Note that a nanophotonic chiplet will also include electrical interfaces (e.g., power, ground, off-NSiP I/O), and an NSiP will likely include both nanophotonic chiplets and standard chips with just electrical interfaces. As with standard SiPs, chiplets can be combined in both 2D or 3D configurations. Figure 1 illustrates two classes of NSiP integration: *monolithic disintegration* uses the same total silicon area as in a monolithically integrated system-on-chip (SoC), but divides the silicon area among multiple small chiplets; *macrochip integration* uses significantly more area than what is possible with monolithic integration by composing large reticle-sized chiplets.

2. Potential Advantages of NSiP

Electrical SiPs have three advantages over SoCs: enabling systems not possible with an SoC; reducing the non-recurring engineering (NRE) cost by composing off-the-shelf (OTS) chiplets; and mitigating high-marginal cost due to low yield. These advantages must be weighed against reduced performance and efficiency as compared to intra-SoC communication and increased marginal cost due to additional assembly and testing. NSiPs have the same advantages but can potentially provide inter-chiplet latency, energy efficiency, and bandwidth density that is comparable or even better than purely intra-chip communication. NSiPs will still increase the marginal cost, which makes this integration strategy most appropriate when an SoC can not achieve the design goals, or for low- to medium-volume markets. We now discuss each of the three advantages in more detail to motivate our interest in NSiPs.

Enable Systems Not Possible with an SoC – NSiPs using macrochip integration have been previously proposed [4, 5] enabling very large single-package systems that are simply not possible with an SoC. NSiPs also allow mixing chiplets fabricated in a process customized for that chiplet's function. For example, Figure 1(a) illustrates an SoC with embedded DRAM, but the NSiP in Figure 1(b) can potentially achieve much higher DRAM density with similar processor-to-memory performance by using nanophotonics and a DRAM chiplet fabricated in a customized DRAM process.

Reduce NRE vs. SoC – NSiPs allow low-cost system design through OTS chiplet composition. A similar motivation
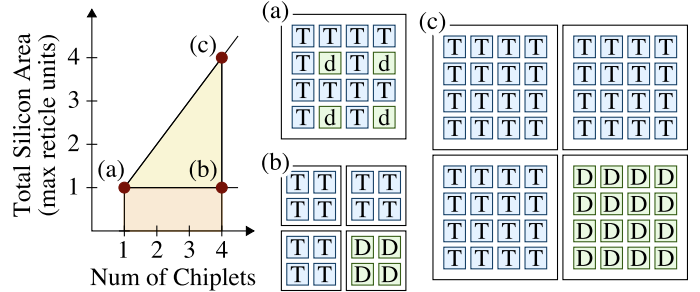

Figure 1: Classes NSiPs Integration – (a) SoC using monolithic integration, (b) NSiP-4 using monolithic disintegration, (c) NSiP-4 using macrochip integration. T = tile w/ processor, SRAM, or accelerator, d = bank of embedded DRAM, D = bank in standard DRAM chip.
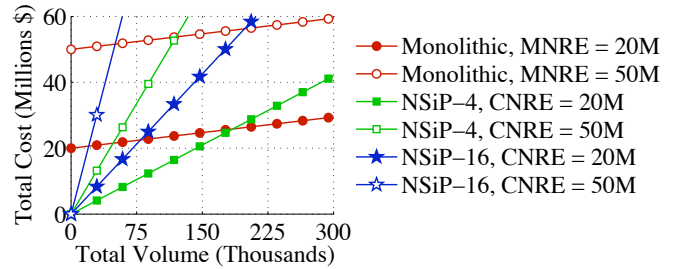

Figure 2: Total Cost vs. Volume – MNRE = NRE for producing monolithic die, CNRE = NRE for producing chiplet, cost model based on [3].
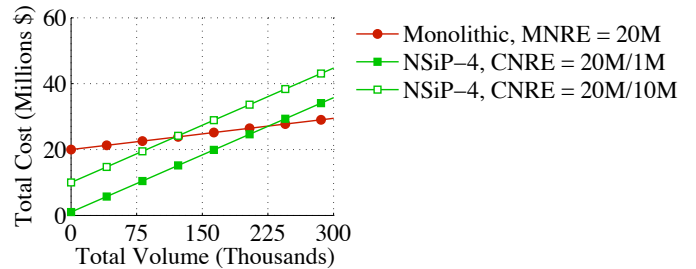

Figure 3: Total Cost vs. Volume w/ Custom Chiplet – MNRE = NRE for producing monolithic die, CNRE = NRE for mass produced chiplet and custom chiplet in cheaper technology, cost model based on [3].

is behind recent advanced electrical SiP architectures [3]. Figure 2 shows the total cost as a function of volume for a 45-nm 200-mm^2 monolithic SoC, a 4-chiplet NSiP, and a 16-chiplet NSiP at two NREs. NSiPs are cost effective for low- to medium-volume markets by amortizing the chiplet NRE over many different products. Figure 3 shows the total cost if three OTS chiplets are composed with one custom chiplet fabricated in an older less-expensive technology just for this NSiP. The impact on the cross-over point is modest, suggesting that limited customization in an older technology could provide an interesting intermediate point between a fully customized SoC and an completely OTS NSiP.

Reduce Marginal Cost vs. SoC – By testing chiplets individually it may be possible to compose only working chiplets. Preliminary analysis for estimated defect densities in modern processes, suggests this may not be a compelling advantage un-

less future processes result in significantly lower yields. There is, however, an opportunity for more flexible system binning by composing chiplets that meet a certain design constraint. For example, the fastest chiplets can be composed to create more high-performance NSiPs than possible with an SoC.

3. NSiP Device-Level Strategy

Previous work on optical interconnect in SiPs has relied on thin-film opto-electrical components integrated in the actual package [1]. Unfortunately, this approach can incur significant overhead, so we are working with device experts on a back-end-of-line (BEOL) nanophotonic technology that uses deposited poly-silicon rings, multi-layer silicon-nitride waveguides, and germanium photodetectors. Unlike a front-end-of-line approach, BEOL devices can be deposited on a wide variety of chips fabricated in different processes. Progress has been made on demonstrating ring modulators using a high-temperature deposition [6], and there is on-going work on fabricating these and other devices within a more reasonable BEOL thermal envelope. If successful, this technology could enable NSiP prototypes to be implemented by depositing optical devices (in an academic research lab) on custom chiplets fabricated through a standard CMOS foundry.

4. NSiP System-Level Strategy

Figure 4 illustrates the 2-fly flattened butterfly topology [2] we are currently investigating as a template for NSiPs. This low-diameter network topology minimizes inter-chiplet latency and enables us to exploit nanophotonics when implementing long global channels. It is also a good match for combining nanophotonic channels with electrical buffering, switching, and arbitration. We explicitly avoid any form of optical switching to simplify the design and reduce the risk associated with more complicated devices. Scaling to larger numbers of tiles is possible by increasing the radix (i.e., integrating more tiles onto each chiplet) or adding a second stage to the flattened butterfly (i.e., two stages of E-O-E conversion). Even larger systems may eventually require some form of optical switching.

Our low-diameter, low-latency topology provides tightly coupled congestion feedback between all routers enabling efficient routing algorithms such as universal globally adaptive load-balanced (UGAL) routing [2]. UGAL routes packets either minimally or non-minimally through a random intermediate router. The choice is based on the number of flits per queue on each output of the local router. The tightly coupled nature of our small network and the extra bandwidth density of nanophotonics makes it feasible for each input terminal to have knowledge of all queues in the network as opposed to a small subset. Thus we are exploring UGAL with global information (UGAL-GI) where credits for all intermediate queues are sent to each router to be factored into the adaptive routing decision.

Figure 5 illustrates a possible implementation of the topology shown in Figure 4. Nanophotonic transmitters and receivers are tightly integrated into each chiplet using the BEOL technology previously described. A centralized hub chip uses purely passive devices to shuffle wavelengths between the chiplets. There are several advantages to using a hub chip as opposed to directly interconnecting the chiplets. A hub chip is thermally isolated, reduces the number of optical couplings, can use a manufacturing process optimized for nanophotonics if desired, and allows mixing a single fabricated chiplet in various configurations.
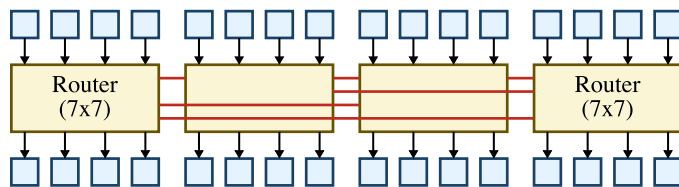

Figure 4: NSiP-4 Flattened Butterfly Topology – Input terminals at top, output terminals at bottom. Inter-router lines represent two independent nanophotonic channels in opposite directions.
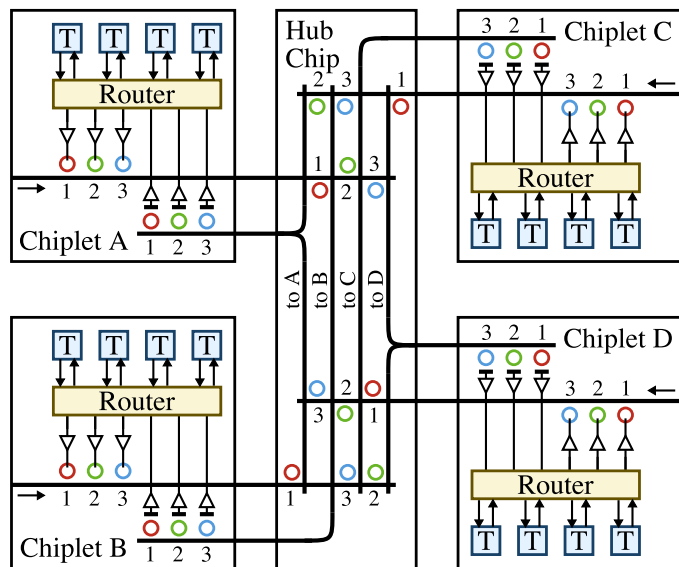

Figure 5: NSiP-4 Microarchitecture and Abstract Layout – T = tile, ring = many parallel rings, ring numbers = specific set of wavelengths.

5. Preliminary Results

We compared a 16-tile, four-chiplet system to two monolithically integrated electrical on-chip networks: a 4×4 mesh and a concentrated ring. First-order optical power calculations suggest very reasonable laser power requirements largely owing to a minimal number of optical couplers and the hub-chip's low-loss passive devices. Preliminary cycle-level simulations on a variety of synthetic traffic patterns suggest that our NSiP approach is able to achieve comparable latency and throughput as the fully electrical on-chip networks. We also compared UGAL to UGAL-GI with varying phit sizes based on possible nanophotonic technology projections, and as expected observed 12-25% higher throughput than UGAL on adversarial traffic and 5-10% higher throughput over all patterns. We are currently working on more detailed power and performance models for larger NSiPs, and we are investigating the possibility of fabricating a small proof-of-concept NSiP prototype.

References

[1] G.-K. Chang et al. Chip-to-Chip Optoelectronics SOP on Organic Boards or Packages. *IEEE Tran. on Advanced Packaging*, 27(2), 2004.

[2] J. Kim et al. Flattened Butterfly Topology for OCNs. *MICRO*, 2007.

[3] M. M. Kim et al. Brick and Mortar Silicon Manufacturing. *ISCA*, 2007.

[4] P. Koka et al. Silicon-Photonic Network Architectures for Scalable, Power-Efficient Multi-Chip Systems. *ISCA*, 2010.

[5] Y. Pan et al. Exploring Benefits and Designs of Optically Connected Disintegrated Processor Architecture. *WINDS*, 2010.

[6] K. Preston et al. Deposited silicon high-speed integrated electro-optic modulator. *Optics Express*, 17(7):5118–5124, 2009.