Investigation of Cost-Optimal Network-on-Chip for Passive and Active Interposer Systems

Dylan Stow
University of California, Santa Barbara
Santa Barbara, California
dstow@ucsb.edu

Itir Akgun
University of California, Santa Barbara
Santa Barbara, California

Yuan Xie
University of California, Santa Barbara
Santa Barbara, California
yuanxie@ucsb.edu

ABSTRACT

Interposer-based packaging is becoming a widespread methodology for tightly integrating multiple heterogeneous dies into a single package, with the potential to improve manufacturing yield and build larger-than-reticle-sized systems. However, interposer integration also introduces possible communication bottlenecks and cost overheads that can outweigh these benefits. To avoid these drawbacks, the abundant interposer interconnect can be leveraged as network-on-chip interconnection fabric to provide high-bandwidth, low-latency communication between chiplets and memory stacks. This work investigates this new interposer design space of passive and active interposer technologies, network-on-chip topologies, and clocking schemes to determine the cost-optimal interposer architectures for a range of performance requirements.

CCS CONCEPTS

• **Hardware → 3D integrated circuits**; **Economics of chip design and manufacturing**.

KEYWORDS

Active Interposer, Cost and Yield Modeling, Network-on-Chip

1 INTRODUCTION

As Moore's Law slows and cost per semiconductor area increases [23], computer architects will need to utilize new architectures and packaging technologies to achieve the system improvements traditionally realized through transistor process scaling. Recently, the concept of multi-die integration has received renewed attention as a promising solution to these challenges. Instead of the monolithic fabrication of modern Systems-on-Chips (SoC) onto a single large die, die-level integration technologies integrate multiple semiconductor dies, each fabricated individually, into a single package. This concept, proposed as far back as the original Moore's Law paper [15], has long been considered as a method to improve semiconductor yields through the use of Known Good Die (KGD) validation, ensuring the functionality of each die before integration. While a single critical defect can disable the functionality of an entire monolithic system, pre-bond validation of each smaller die reduces the loss per defect and results in a greater number of fully-functional systems [18]. Historically, multi-chip module (MCM) packages have been used to integrate multiple dies onto a single substrate, and recently MCMs have again been deployed to improve yield and provide product scaling [3].

Although MCMs can provide a platform for die integration, the coarse-pitch substrate interconnect can only provide limited bandwidth, with reduced efficiency and increased latency, compared to on-chip interconnect [2]. However, these limitations can be solved by instead using fine-pitched silicon interposers, which are already employed in commercial products to integrate 3D High Bandwidth Memory [12]. These interposers utilize standard semiconductor interconnect technology, such as that in the $65nm$ process node, and bond each die using fine pitched microbumps, with current bump pitch of $55\mu m$ and future pitches of $20\mu m$ [22]. With interposer integration technology, systems can realize the yield and flexibility benefits of multi-die integration while maintaining the high performance Network-on-Chip (NoC) fabric utilized to connect modules in modern SoCs.

While recent research has demonstrated the cost and performance benefits of interposer-based system integration [4][10][17][18][19], the details of the interposer technology and design are still an area of development and debate. Commercial interposers are today manufactured as *passive* silicon interposers [12], which contain metal interconnect but do not have active CMOS transistors. The simple nature of passive interposers greatly reduces wafer costs, but without transistors the interposer can only provide non-repeated point-to-point connections between chiplets. An alternative to passive interposers is *active* interposer technology, in which the interposer is instead manufactured from a standard CMOS process (with the addition of die thinning and Through-Silicon Via [TSV] insertion). Active interposers have been utilized in several recent studies [8][10][19][22] to provide high-speed repeated links and to move NoC routers onto the interposers, thereby providing more network bandwidth than available on monolithic SoCs. Unfortunately, the active interposer can become a significant cost overhead, especially if it utilizes advanced process technologies. To minimize active interposer cost, previous work has limited the amount of active logic area on the interposer to improve yield [10].

Prior studies have demonstrated the relative cost and frequency differences of passive and active interposers [19], but no work has conclusively demonstrated the architecture-level impact of this interposer technology selection: should the interposer for a given system be passive or active, and if active which process node should be utilized? This work explores the complex interactions between the interposer technology, the Network-on-Chip topology, and the physical link and router implementation, which then determine the interposer active area and yield, the chiplet area and yield and cost, the network link frequency, and ultimately the system cost and the NoC bandwidth and latency. While it is not possible to make conclusions for all possible multi-die systems, traffic patterns,

and performance requirements, this work aims for generalizable conclusions by sweeping a range of bisection bandwidth targets, representing different memory and coherence requirements, for a realistic multi-core system with integrated die-stacked memories and a synchronous NoC that is compatible with standard design automation flows. Further, NoC latency trends are analyzed to study the interaction between interposer technology and network topology. This work concludes that when only considering network bandwidth, passive interposer integration is almost always cost-optimal when compared to active interposers, even when considering the router area overhead impact on chiplet yield. However, due to longer link latencies and clock-crossing overheads in passive interposers, active interposers demonstrate a significant advantage for latency-sensitive systems. Additionally, active interposers in mature process technologies can achieve high bandwidth and lower latency at system costs within 30% of passive interposer systems.

2 CHIPLET INTEGRATION WITH INTERPOSERS

The yield of large semiconductor systems can be improved by splitting the monolithic chip into multiple smaller chiplets that are reintegrated onto an interposer substrate [17]. However, the interposer needs to provide sufficient bandwidth at low-enough latency to avoid degrading the system performance. Due to the plentiful fine-pitch interconnect available in interposers, prior research has demonstrated that interposers can actually improve system performance over monolithic fabrics, which may have limited interconnect budgets, by implementing high-performance Networks-on-Chip with wide links or high radix routers [1][4][8][10]. Some interposer-based architectures have leveraged active interposers with in-interposer routers and repeated links [10], while others have constrained the systems to metal-only passive interposers in order to limit cost overheads [4]. In this section, the economic, architectural, and circuit-level details of these different interposer technologies are elaborated and compared.

2.1 Cost Implications of Die Integration

When splitting a system into multiple smaller chiplets, the total system yield can be improved by employing KGD testing of each chiplet before bonding. The manufacturing yield Y_{die} of unstructured semiconductor logic can be estimated based on die active area A_{die} using the Negative Binomial model given in Equation 1:

$$Y_{die} = \left(1 + \frac{A_{die} D_0}{\alpha}\right)^{-\alpha} \quad (1)$$

where D_0 is the defect density and α is the clustering coefficient (and constant wafer yield loss is ignored). As demonstrated by the curve in Figure 1 (a) with $D_0 = 2000/m^2$ and $\alpha = 3$, a larger die size reduces the odds of defect-free manufacturing. For example, a monolithic system with $A = 336mm^2$ would have a defect-free yield of about 55%, while disaggregating the same system into four chiplets each with $A = 84mm^2$ increases the yield of each chiplet to about 85%. A small scale example is visualized in Figure 1 (b) and (c). If the chiplets can first be validated with sufficiently high coverage and can then bonded with sufficiently high success rate (commercial HBM integration has achieved greater than 99% bond yield [12]), then total manufacturing yield can significantly improve.

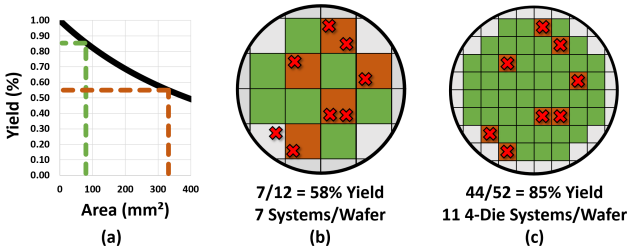

Figure 1: (a) Yield versus area with defect density $D_0 = 2000/m^2$ and $\alpha = 3$, with examples design points highlighted at $A = 336mm^2$ ($Y = 0.55$) and $A = 84mm^2$ ($Y = 0.85$). Scaled visual examples with eight defects per wafer are shown for (b) a monolithic system and (c) an equivalent 4-chiplet system.

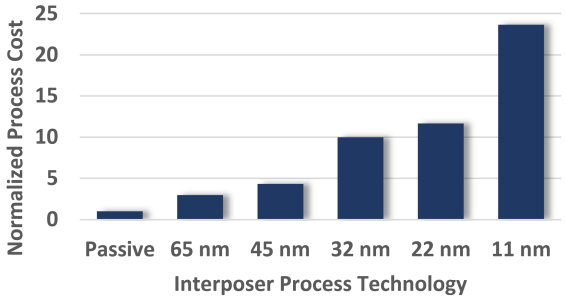

Figure 2: Normalized wafer costs for passive and active interposer.

An additional benefit, also visualized in Figure 1 (b) and (c), is that reduction of the die size can increase the number of rectangular dies that can be fit into the circular wafer, further increasing the effective yield.

However, employing an interposer also introduces cost overheads, including the manufacturing of an additional large silicon die. To minimize this overhead, passive interposers only include interconnect layers, and no active transistors, and therefore have high yields and much lower wafer costs. Active interposers, reliant on relatively expensive transistor fabrication, may need to limit the amount of active area devoted to routers and other logic in order to minimize yield loss. Active interposer process technology selection determines the wafer cost (advanced processes are much more expensive for equal area) and the active area (advanced processes can implement routers in less area, thus improving yield given the same defect density). Figure 2 demonstrates the relative process costs for different possible interposer technologies, based on TSV integration and process cost data from *IC Knowledge* [7]. An additional integration cost overhead is the cost of KGD validation of each chiplet before bonding. Although validation time (and therefore cost) primarily scale with design size and complexity, constant per die overheads result in greater validation costs for multiple small dies than for a single monolithic die [7]. Despite these overheads, prior work has demonstrated that both passive and active interposers can indeed be more cost effective than monolithic systems at standard design sizes [4][19]. An example is shown in Figure 3

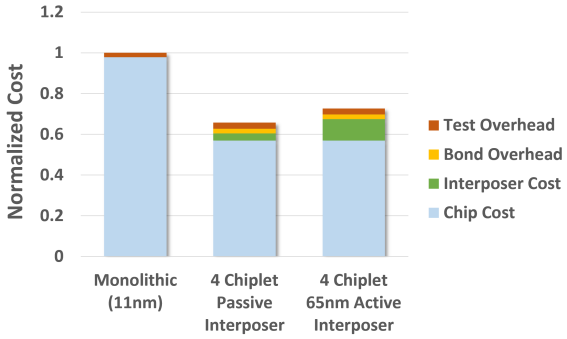

Figure 3: Manufacturing cost example for $336mm^2$ chip in $11nm$ process, with equivalent four-chiplet systems on $448mm^2$ passive and $65nm$ active interposers.

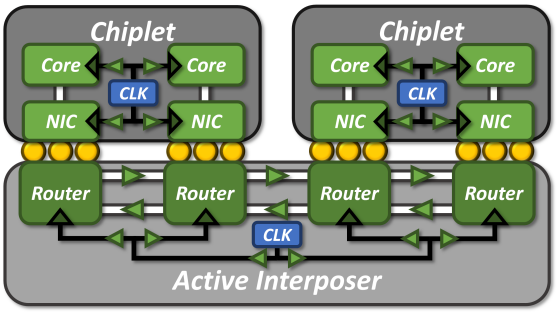

Figure 4: Logical diagram of Network-on-Active-Interposer, with on-interposer routers, repeated interconnect, and synchronous interposer clock.

for a $336mm^2$ monolithic system, compared against four-chiplet systems on $448mm^2$ passive and $65nm$ active interposers (assuming a conservative $1mm$ spacing between and around the chiplets), with $D_0 = 2000/m^2$, $\alpha = 3$, and 99% bond yield.

Accordingly, the cost of each monolithic chip, individual chiplet, or passive or active interposer can be estimated using Equation 2:

$$C_{die} = \frac{\left(\frac{C_{wafer}}{N_{die}} + C_{test}\right)}{Y_{die}} \tag{2}$$

where C_{wafer} is the process-dependent wafer cost, C_{test} is the cost to validate the die after manufacturing, and N_{die} is the number of dies per wafer, which can be estimated with Equation 3:

$$N_{die} = \frac{\pi \times (\phi_{wafer}/2)^2}{A_{die}} - \frac{\pi \times \phi_{wafer}}{\sqrt{2 \times A_{die}}} \tag{3}$$

where ϕ_{wafer} is the wafer diameter. The total manufacturing cost of an interposer system with n chiplets can then be estimated using Equation 4:

$$C_{2.5D} = \frac{C_{int} + \sum_{i=1}^{n}(C_i + C_{bond_i})}{(Y_{bond})^n} \tag{4}$$

2.2 Active Interposer Integration

To provide high-bandwidth, low-latency, scalable connectivity between modules within and between chiplets, Network-on-Chip interconnection architectures can be deployed that span the chiplets and the interposer. These interconnection networks, composed of router modules and connecting links, pass message packets between the logic module and memory controller nodes on the chiplets. While many different network topologies are possible, the selection of either active or passive interposer technologies influences many logical and physical aspects of the resulting network.

Active interposers are manufactured to include active transistor devices within the interposer, providing several unique capabilities not possible with passive interposers:

- **Router Location:** Router modules can be moved from the chiplets to the interposer, reducing area in the chiplets and therefore improving chiplet yield and cost. Further, by moving routers into the interposer, links do not need to pass through the chiplet-interposer interface, thus saving microbump resources for the power delivery network.
- **Link Transmission:** Links on active interposers can utilize repeaters to reduce transmission delay, reducing the delay of long routes from a quadratic to linear relation with distance. Reduced link delays allow for faster NoC frequencies and further single-cycle link distances. Further, with routers in the interposer, delay is reduced by avoiding microbump capacitance [6] and the capacitance of Electrostatic Discharge (ESD) protection circuits, which are required to protect the chiplets interface during bonding [21] (and contribute significantly greater capacitance than a microbump [19]).
- **NoC Clocking:** A primary challenge for SoC-scale synchronous NoCs is the distribution of a low-skew, low-jitter clock to all router modules [11]. Alternative clock schemes, like mesochronous or asynchronous NoC, may incur area or latency overheads and may be difficult to design using standard EDA flows. However, by utilizing the buffers and interconnect available in the active interposer, H-tree clock networks can be generated with sufficiently low jitter (less than $150ps$) to allow for GHz-frequency synchronous NoC in the active interposer [13]. The interposer can even supply a clock signal to the chiplets to avoid interposer-chiplet synchronization latency [14], but this work assumes that a synchronization latency overhead is incurred between the interposer and chiplets to enable flexible DVFS of each core.

This active interposer NoC architecture is depicted in Figure 4.

An active interposer can be implemented using existing CMOS technologies. Because global interconnect performance has remained relatively constant throughout recent process technologies, the link delay has low sensitivity to the interposer process selection. However, NoC routers can consume significant area when implemented with many ports and wide flits, especially in older, larger process technologies. Advanced process technologies can implement complex routers with much smaller areas to improve yield, but wafer costs are significantly higher than for older processes. Relative process technology costs are shown in Figure 2 and router area is addressed in Section 3.

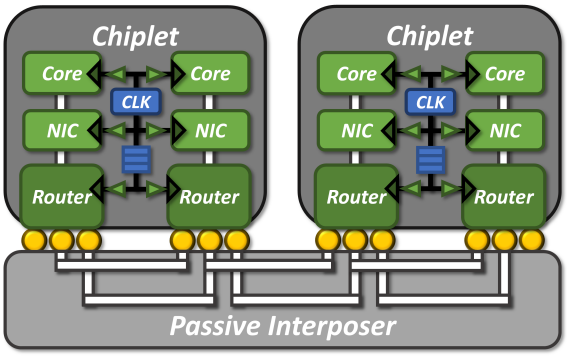

Figure 5: Logical diagram of Network-on-Passive-Interposer, with on-chiplet routers and non-repeated interconnect.

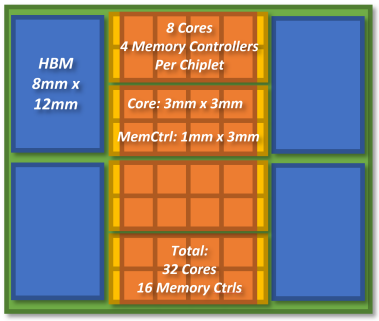

Figure 6: Base system under evaluation, with thirty-two $3mm \times 3mm$ cores, sixteen $1mm \times 3mm$ memory controllers, and four $8mm \times 12mm$ stacks of High Bandwidth Memory, all integrated with an interposer. The system is demonstrated with four eight-core chiplets in this example.

2.3 Passive Interposer Integration

Unlike active interposers, passive interposers only contain metal interconnect, so they cannot include active logic like routers, repeaters, or FIFO queues in the interposer. This leads to several challenges when constructing an NoC across the interposer:

- **Router Location:** Router modules must be placed on the chiplets, contributing to chiplet area and degrading yield and cost. Additionally, links that pass through the chiplet-interposer interface consume microbump resources.
- **Link Transmission:** Links on passive interposers cannot utilize repeaters, so long routes exhibit a quadratic delay with distance, resulting in longer link latencies or slower networks. Links must additionally pass through two chiplet-interposer interfaces, which introduce additional capacitive loads from the microbumps and ESD protection circuits.
- **NoC Clocking:** The challenge of low-skew, low-jitter clock distribution to all routers is complicated by the fact that the passive interposer cannot generate a clock on the interposer and cannot include buffers to drive the low-jitter clock network. To complicate matters further, any clock network that spans and utilizes the chiplets must contend with the potentially large inter-chiplet process variation and power noise. Accordingly, a synchronous NoC with GHz frequency is not likely to be feasible with a passive interposer. Instead, a globally asynchronous locally synchronous (GALS) paradigm is more likely. In such a scheme, routers on each chiplet are synchronous, but synchronization must occur when transmitting between chiplets, incurring a synchronization delay of several cycles [10]. To allow for flexible DVFS of each core without constraining the NoC frequency, additional synchronization occurs between the routers and Network Interface Controller (NIC) within the same chiplet.

This passive interposer NoC architecture is depicted in Figure 5.

3 INTERPOSER COMPARISON METHODOLOGY

To compare the cost and performance trade-offs of active and passive interposers, this work studies a base system floorplan based on recent commercial systems [3]. Figure 6 shows the base system with thirty-two generic cores implemented in $11nm$ process technology. The system also contains four stacks of in-package memory, similar to the JEDEC High Bandwidth Memory standard [16], with eight memory channels per stack. Sixteen memory controllers are distributed between the cores and placed on the chiplet periphery, with two memory channels managed by each memory controller. The entire system is integrated using either an active interposer or a passive interposer, which provides connectivity between nodes to route all inter-core message traffic. NIC terminals are located at each core and memory controller. Figure 6 demonstrates a configuration with four chiplets, but systems from one to eight chiplets are considered. The network is clocked at $2\ GHz$.

The network topologies visualized in Figure 7 and detailed in Table 1 are assigned to router nodes connected to these terminals. Misaligned topology alternatives are also included [10].

The selection of active or passive interposer technology determines the relations between network performance and system cost:

- **Router Area:** With an active interposer, the routers are placed in the interposer, increasing the interposer active area used in yield calculation but reducing the chiplet area. With a passive interposer, all routers contribute to the chiplet area and are implemented in the $11nm$ process technology.
- **Link Frequency:** Links on the active interposer are repeated with optimally-sized repeaters in the selected process technology to minimize link delay. For passive interposers, link delay is calculated for the non-repeated route with capacitance overhead for microbumps ($15fF$) and ESD protection ($200fF$) [19]. Delays are computed from SPICE using pi-model interconnect segments. At the $2\ GHz$ network frequency, the active interposer is able to transmit most link distances in a single cycle. Only the longest links (the longest vertical links in the Folded Torus and the longest diagonals in the Double Butterfly and ButterDonut) require two cycles on the active interposer. The passive interposer, however, requires at least two cycles for all but the shortest links (as in the mesh or concentrating routers), and the longest diagonal links take eight cycles. Link latencies for each distance in the topologies are given in Table 2. Multi-cycle links, depending

Table 1: Network-on-Interposer Topologies

Topology	Nodes	Links	Diameter	Avg. Hops	Bisection Links
Mesh	32 (4×8)	52	10	4.9	8
Concentrated Mesh	16 (4×4)	18	6	3.5	5
Double Butterfly	16 (4×4)	24	3	3.1	8
Folded Torus	16 (4×4)	32	4	3.0	8
ButterDonut	16 (4×4)	28	3	3.0	12
Misaligned Concentrated Mesh	12 (3×4)	17	5	3.1	4
Misaligned Double Butterfly	12 (3×4)	16	2	3.1	8
Misaligned Folded Torus	12 (3×4)	24	3	2.7	8
Misaligned ButterDonut	12 (3×4)	20	2	2.8	12

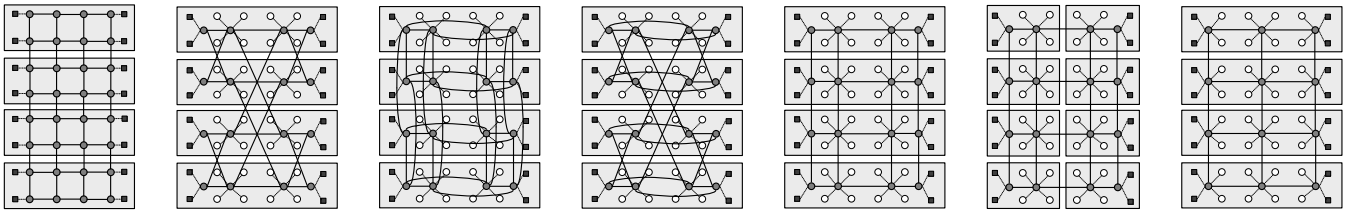

(a) Mesh: 4 Chiplets | **(b) D. Butterfly: 4 Chiplets** | **(c) Folded Torus: 4 Chiplets** | **(d) ButterDonut: 4 Chiplets** | **(e) CMesh: 4 Chiplets** | **(f) CMesh: 8 Chiplets** | **(g) Mis. CMesh: 4 Chiplets**

Figure 7: Network-on-Interposer four-chiplet topologies. Eight-chiplet and misaligned examples demonstrated for CMesh

Distance (*mm*)	Active (cycles)	Passive (cycles)
3.5	1	1
6.5	1	2
10	1	3
13	2	4
19.5	2	8

Table 2: Link latencies for active and passive interposers

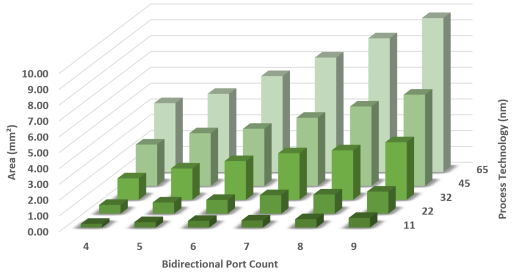

Figure 8: Router area with 512-bit flit width versus bidirectional port count and process technology.

on EDA methodology, may require flip-flops in the active interposer or chiplets [4] to meet synchronous requirements.

- Due to clock domain division between each chiplet terminal and the interposer routers, which enables core DVFS, the active interposer incurs a synchronization overhead only on the initial and final transitions from chiplet to interposer. The passive interposer must additionally synchronize when transmitting between chiplets. This work assumes a three-cycle synchronization latency at each clock crossing.

To compare each combination of interposer technology and network topology, it is necessary to first estimate the area of network routers and links and to then compute the resulting yields and costs. Router areas are estimated using the DSENT Network-on-Chip modeling tool [20], with process technologies from 11*nm* to 65*nm*. Each three-staged pipelined router has 16 virtual channels and eight buffers per virtual channel. Routers area is given for port count and process technology with 512-bit flit width in Figure 8. Flit width is also utilized as a parameter. Each link is composed of global-level interposer interconnect with wire width and spacing of 350*nm*. Based on the system floorplan, the router and link areas are calculated for each topology. Figure 9 shows the average and range of router utilization across topologies for each active interposer technology and flit width. Figure 10 shows the average and range of interposer interconnect utilization under the chiplets for flit widths from 32 to 1024 bits, assuming the availability of two X-Y routing layer pairs. Assuming a microbump pitch of 20μm, each chiplet has sufficient microbump resources to connect all topologies with at least 1024 bit flit width, even when allocating half of microbump resources to the power delivery network.

After calculating the areas of each interposer's network components, the cost of the interposer and system can be computed by estimating the interposer yield. The cost and yield estimation methodology from Section 2.1 is utilized to compute the interposer and chiplet active area yield, with defect density $D_0^{Logic} = 2000/m^2$, as well as the global interconnect yield with lower defect density [5]

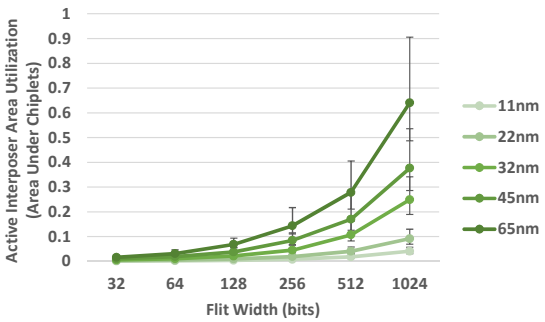

Figure 9: Average router utilization percentage, across network topologies, of active interposer area under the chiplets. Error bars indicate the range between topologies.

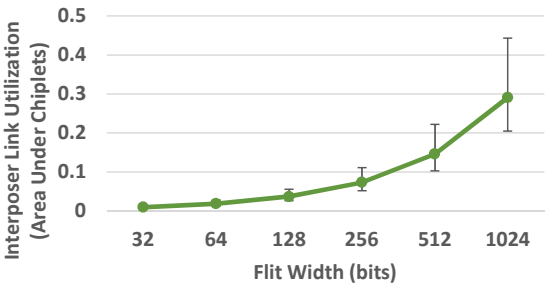

Figure 10: Average link utilization percentage, across network topologies, of interposer interconnect under the chiplets. Error bars indicate the range between topologies.

$D_0^{Link} = 500/m^2$. The manufacturing cost of each interposer and chiplet is then determined by calculating the number of dies per 300mm wafer and incorporating the process-dependent wafer cost, shown in Figure 2, and the yield loss from active area and link defects. Each chiplet is bonded to the interposer with a bond yield of 99% [12]. For all cost results, only interposer and chiplet costs (and not the memory stack costs) are included, but the interposer footprint includes the area required for the memory stacks.

The resulting normalized interposer costs for each process technology are demonstrated in Figure 11 for the ButterDonut topology. When only considering the interposer cost, the passive interposer is significantly less costly than any of the active interposer configurations due to a lower process cost and a lack of active area. For the active interposer technology options, the mature 65nm process is consistently the lowest cost, even despite the large active area utilization at the largest flit width of 1024 bits.

The full system costs for the ButterDonut topology with 512-bit flit width are shown in Figure 12. All active interposer systems have eight chiplets but vary by interposer technology. The advanced process technologies incur significantly higher interposer overheads due to higher wafer costs, while the router yield is less significant even for the 65nm process with high utilization. The passive interposer configurations are shown with a range of chiplet sizes,

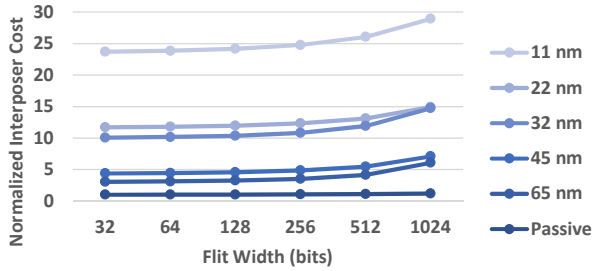

Figure 11: Normalized interposer cost of ButterDonut topology, for each process technology, versus flit width.

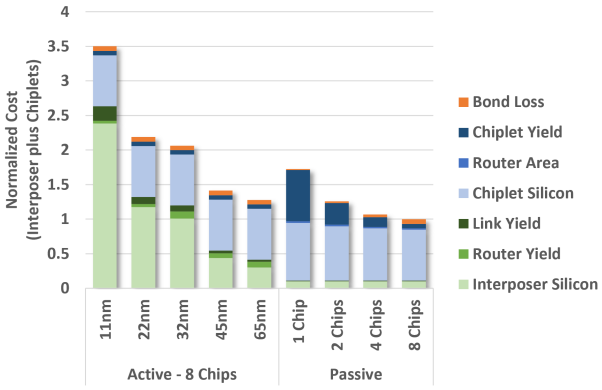

Figure 12: Normalized cost profile of interposer and chiplets across interposer technology options for the 512-bit flit width ButterDonut topology, with eight chiplets for active interposer configurations. Active interposer process cost dominates in advanced processes. Router yield (active) and router area (passive) are small contributions to cost.

demonstrating the improvement in chiplet yield with increased chiplet count, as well as an increase in bond loss overhead.

4 BISECTION BANDWIDTH OF NETWORK-ON-INTERPOSER

To compare the network performance results of a given interposer technology and network topology, the bisection bandwidth metric can be selected to measure the topology-level peak communication between cores and memory on opposite sides of the system. The bisection bandwidth for a topology be calculated with the number of bisection links, given in Table 1, the flit width, and the network frequency of 2GHz. The bisection bandwidths can then be compared against the system costs to determine the cost-optimal topology.

Figure 13 shows the normalized cost versus bisection bandwidth for each active interposer topology at each flit width. The bandwidth value is the average of the bisection bandwidths in the x and y directions. Each active interposer configuration utilizes the 65nm process and has eight chiplets. As demonstrated in the figure, the high-connectivity ButterDonut topology is consistently

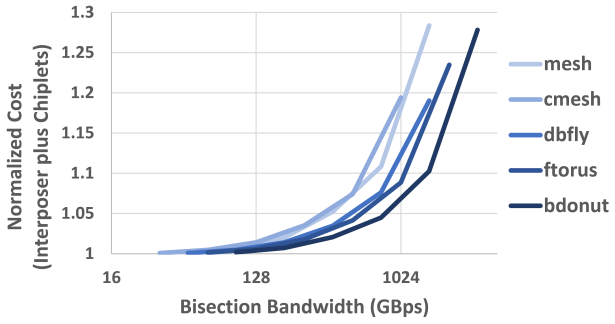

Figure 13: Normalized cost of interposer and chiplets versus bisection bandwidth for each topology on a $65nm$ active interposer with eight chiplets.

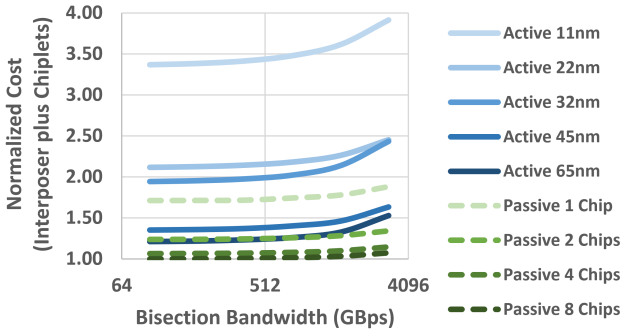

Figure 14: Normalized cost of interposer and chiplets versus bisection bandwidth across interposer technology options, with eight chiplets for all active interposer configurations.

cost-optimal when only considering cost and bisection bandwidth, despite greater router and link area.

Figure 14 compares total system cost and bisection bandwidth for each interposer technology option, with eight chiplets for each active interposer system. As bisection bandwidth is sensitive only to topology, flit width, and frequency, the passive interposer achieves the same bandwidth, but at significantly lower cost. As demonstrated earlier in Figure 12, the passive interposer has high yield and pays little overhead for including the router area in the chiplets. Overall, investigations of network bisection bandwidth demonstrate a system affinity for small chiplet size, to increase yield, and high-radix networks, to improve bandwidth. Router size and link width have less impact on interposer and system cost than the initial selection of interposer technology, and mature processes with lower cost are advantageous for active interposers even at high utilization and lower yield. However, bisection bandwidth is only one of multiple network metrics that should be considered.

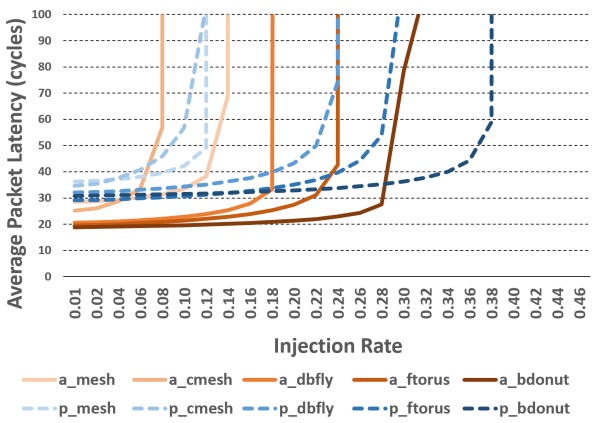

Figure 15: Network latency versus injection rate for active and passive interposers.

5 LATENCY EVALUATION OF NETWORK-ON-INTERPOSER

Although passive interposers are cost-optimal when only considering bandwidth, latency is also a critical metric for many networks-on-interposer systems. In this section, the topologies listed in Table 1 are mapped to passive and active interposers and compared to determine the impact on network latency.

Chiplets on passive interposers, as explained in Section 2.3, are clocked independently. Therefore, a three-cycle clock synchronization overhead is added to the latency of inter-chiplet links when using passive interposers. Routers within an active interposer, as explained in Section 2.2, are synchronous and therefore do not have this link synchronization overhead. However, the connections between the NIC terminals and routers does require a synchronization overhead for both the passive and active interposers in order to allow for independent core DVFS. Link latencies between routers for passive and active interposer implementations also depend on the distance, as previously stated in Table 2.

Methodology: Booksim [9] is used to evaluate the performance of network-on-interposer topologies listed in Table 1. For the following evaluations, the bisection bandwidth of the topologies is fixed by balancing the bisection link count with link width. The network frequency is 2 GHz, as explained in Section 3. The network is evaluated on uniform random synthetic traffic, sweeping over injection rate to observe saturation throughput as well as latency.

Active vs. Passive Interposer on Latency: Figure 15 shows average packet latency for networks on passive and active interposer. Two main observations can be made from this study. First, the active interposer realizes lower average latency than the passive interposer. This is mainly due to lower active interposer link latencies over the synchronous network. Second, topology decisions also impact average latency as a lower average hop count associates with lower network latency, as expected.

Aligned vs. Misaligned Topologies on Latency: Figure 16 shows average packet latency for aligned and misaligned topologies for active interposers. Misaligned topologies, as in Figure 7g, have

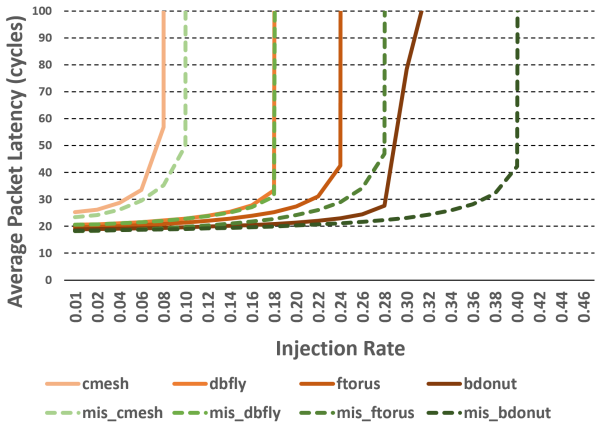

Figure 16: Impact of misaligned topologies on active interposer average packet latency.

fewer routers shared between the cores and memory controller nodes on chiplets. Therefore, misaligned topologies result in lower average latencies compared to the aligned topologies shown in Figure 16 due to lower average hops for both passive and active interposer cases. Additionally, the misaligned topologies result in higher saturation throughput and can tolerate heavier traffic loads.

Chiplet Size on Latency: Figure 17 shows average packet latency for eight-core (eg. Figure 7e) and four-core chiplet sizes (eg. Figure 7f). This evaluation only considers the passive interposer case, since the active case does not incur an off-chiplet synchronization overhead and thus perform similarly on different chiplet sizes. As shown in the figure, smaller chiplet sizes result in higher average network latencies due to the increased frequency of synchronization overhead between chiplets.

To summarize the latency results: 1) Active interposers reduce latency versus passive interposers, 2) Misaligned topologies reduce latency and improve saturation throughput due to lower diameter, 3) Smaller chiplets can increase latency in passive interposers, but they have no effect on latency for active interposers.

6 CONCLUSION

By examining the interaction of interposer technology and network topology, this work concludes that both passive and active interposers may be cost-effective platforms for chiplet integration, depending on system requirements. From a yield and cost perspective, active interposers should generally be implemented using mature process technologies with lower wafer cost, as any yield benefit from smaller routers is overshadowed by the high wafer cost of advanced processes. When only considering bisection bandwidth, passive interposers achieve the same performance at lower cost than active interposers (given the same chiplet size). However, the long multi-cycle links and frequent clock-domain crossings in passive interposers introduce additional latency, and passive interposer systems may sometimes benefit from larger chiplets to reduce this latency overhead.

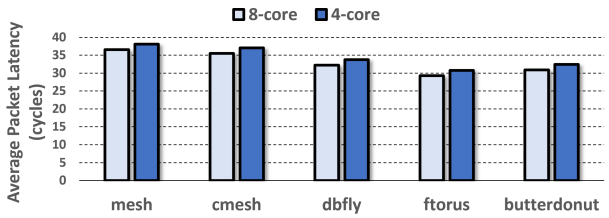

Figure 17: Impact of eight-core vs. four-core chiplet selection on passive interposer average packet latency.

REFERENCES

[1] Itir Akgun et al. 2016. Scalable Memory Fabric for Silicon Interposer-Based Multi-Core Systems. In *ICCD 2016*. 33–40.
[2] Akhil Arunkumar et al. 2017. MCM-GPU: Multi-Chip-Module GPUs for Continued Performance Scalability. In *ISCA 2017*. 320–332.
[3] Noah Beck et al. 2018. 'Zeppelin': An SoC for multichip architectures. In *ISSCC 2018*. 40–42.
[4] Ayse Coskun et al. 2018. A Cross-Layer Methodology for Design and Optimization of Networks in 2.5D Systems. In *ICCAD 2018*. 1–8.
[5] J. A. Cunningham. 1990. The use and evaluation of yield models in integrated circuit manufacturing. *IEEE Trans. Semicond. Manuf.* 3, 2 (May 1990), 60–71.
[6] Pete Ehrett et al. 2017. *Analysis of Microbump Overheads for 2.5D Disintegrated Design*. Report. University of Michigan Ann Arbor.
[7] IC Knowledge LLC 2016. *IC Cost and Price Model, 2016 Revision 05*. IC Knowledge LLC.
[8] Natalie Enright Jerger et al. 2014. NoC Architectures for Silicon Interposer Systems: Why Pay for more Wires when you Can Get them (from your interposer) for Free?. In *MICRO 2014*. 458–470.
[9] Nan Jiang et al. 2013. A Detailed and Flexible Cycle-Accurate Network-on-Chip Simulator. In *ISPASS 2013*. 86–96.
[10] Ajaykumar Kannan et al. 2015. Enabling Interposer-Based Disintegration of Multi-Core Processors. In *MICRO 2015*. 546–558.
[11] Shashi Kumar et al. 2002. A Network on Chip Architecture and Design Methodology. In *ISVLSI 2002*. 117–124.
[12] Chang-Chi Lee et al. 2016. An Overview of the Development of a GPU with Integrated HBM on Silicon Interposer. In *ECTC 2016*. 1439–1444.
[13] Ayan Mandal et al. 2011. An Automated Approach for Minimum Jitter Buffered H-Tree Construction. In *VLSI Design 2011*. 76–81.
[14] Ayan Mandal et al. 2013. A Source-synchronous Htree-based Network-on-chip. In *GLSVLSI 2013*. 161–166.
[15] Gordon E. Moore. 1998. Cramming More Components Onto Integrated Circuits. *Proc. IEEE* 86, 1 (Jan 1998), 82–85.
[16] JEDEC Standard. 2015. High Bandwidth Memory (HBM) DRAM. *JESD235A* (2015).
[17] Dylan Stow et al. 2016. Cost Analysis and Cost-driven IP Reuse Methodology for SoC Design Based on 2.5D/3D Integration. In *ICCAD '16*. 1–6.
[18] Dylan Stow et al. 2016. Cost and Thermal Analysis of High-Performance 2.5D and 3D Integrated Circuit Design Space. In *ISVLSI 2016*. 637–642.
[19] Dylan Stow et al. 2017. Cost-effective Design of Scalable High-performance Systems Using Active and Passive Interposers. In *ICCAD 2017*. 728–735.
[20] Chen Sun et al. 2012. DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling. In *NOCS 2012*. 201–210.
[21] Dimitrios Velenis et al. 2015. Processing active devices on Si interposer and impact on cost. In *3DIC 2015*. TS11.2.1–TS11.2.4.
[22] Pascal Vivet et al. 2015. 3D advanced integration technology for heterogeneous systems. In *3DIC 2015*. FS6.1–FS6.3.
[23] Greg Yeric. 2015. Moore's law at 50: Are we planning for retirement?. In *IEDM 2015*. 1.1.1–1.1.8.