

# Chipelets Myths and Opportunities

*What, why & how to create a new economic model for HPC*

**John Shalf**

Department Head for Computer Science  
Lawrence Berkeley National Laboratory

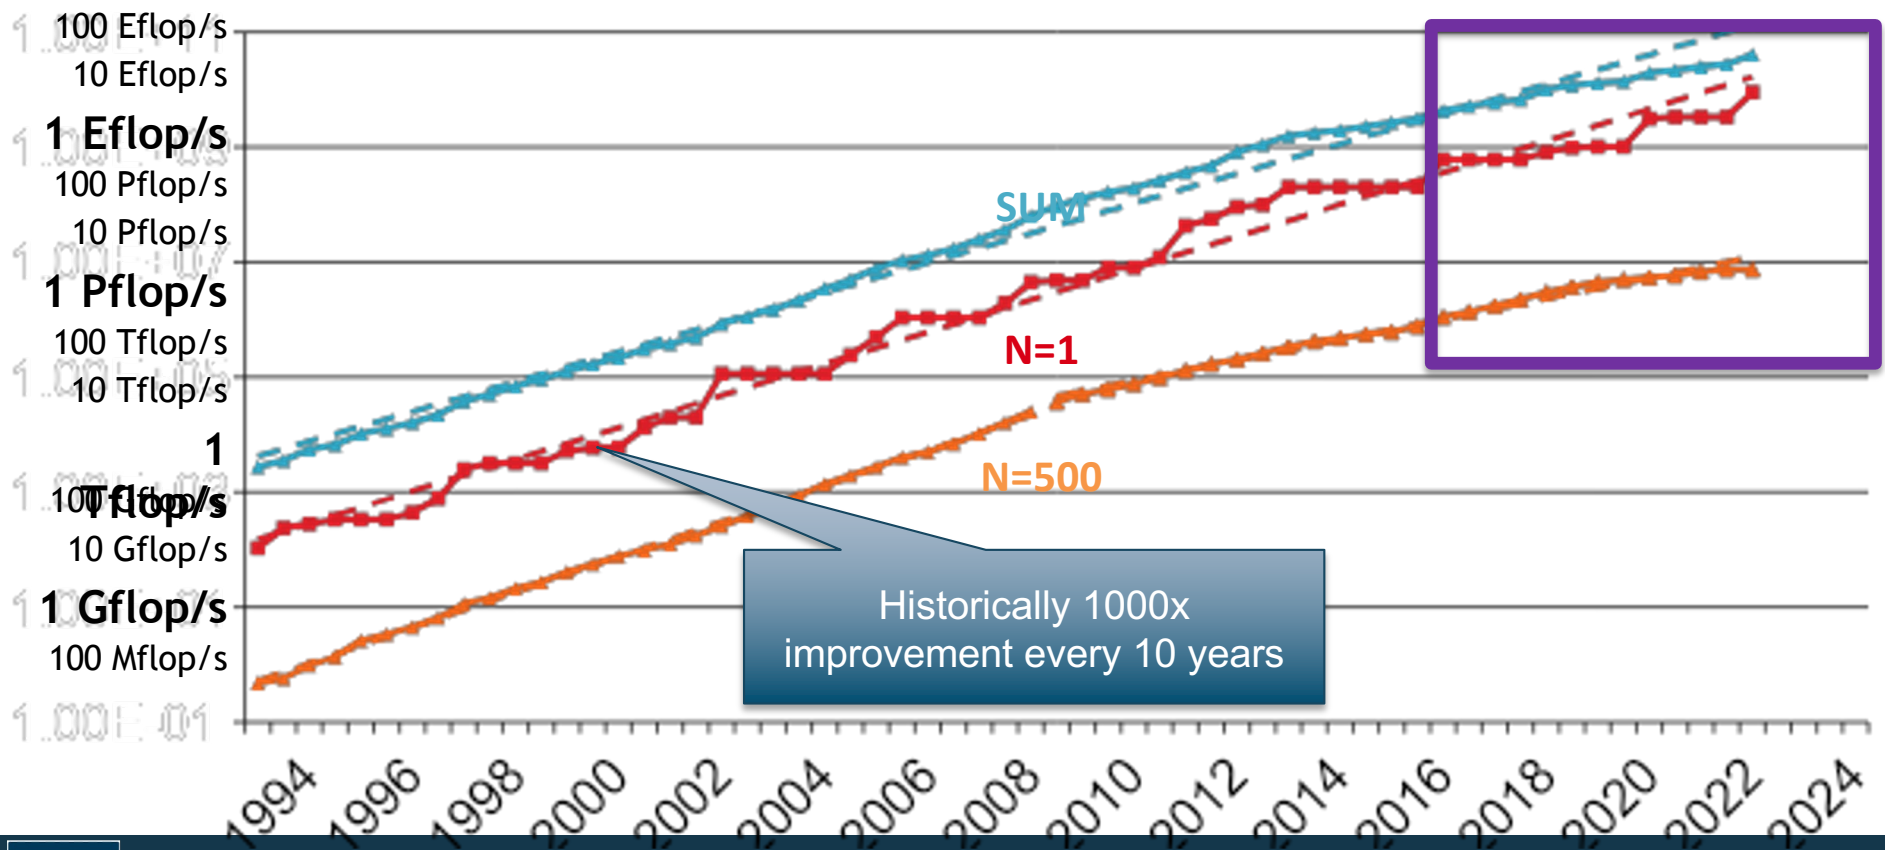
April 26, 2023

Salishan Workshop on High Speed Computing




[jshalf@lbl.gov](mailto:jshalf@lbl.gov)

# Projected Performance Development



# PROJECTED PERFORMANCE DEVELOPMENT



# Specialization:


Natures way of Extracting More Performance in Resource Limited Environment

Powerful General Purpose



Xeon, Power

Many Lighter Weight  
(post-Dennard scarcity)



KNL, AMD, Cavium/Marvell, GPU

Many Different Specialized  
(Post-Moore Scarcity)



Apple, Google, Amazon

# Industry: *Heterogeneous Integration Roadmap*

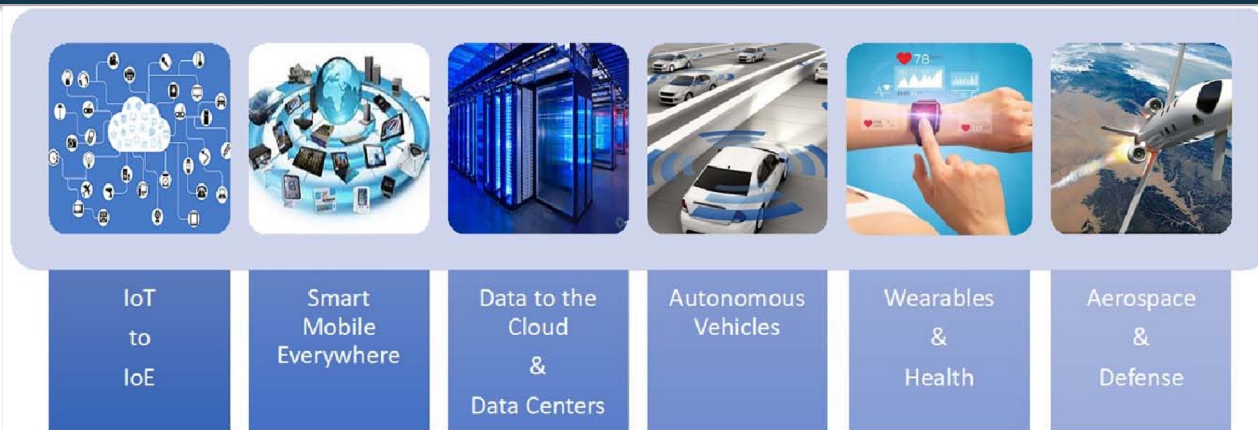


## HETEROGENEOUS INTEGRATION ROADMAP

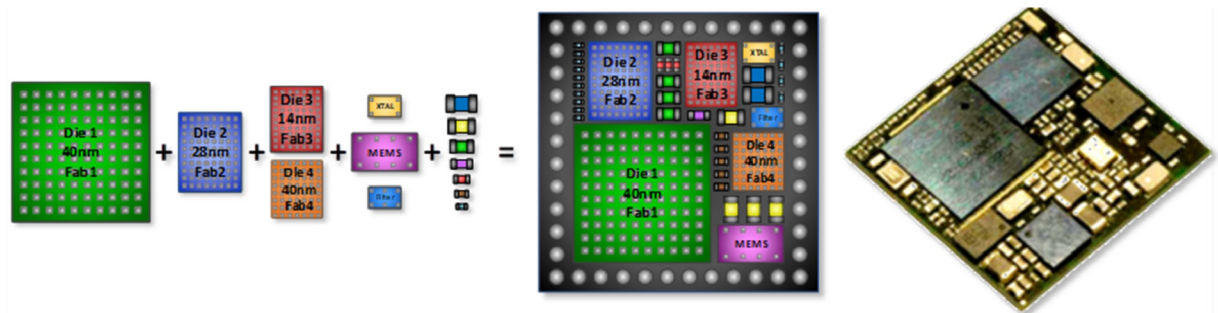
2019 Edition

<http://eps.ieee.org/hir>

HPC and Megadatecenters is 2<sup>nd</sup> chapter









All future applications will be further transformed through the power of AI, VR, and AR.



Die + Heterogeneous

System in Package (SiP)

# What is a Chiplet?



# Heterogeneous Integration Roadmap (HIR)

## 5 generations of bandwidth density doubling

June, 2019

Chapter 22: Interconnects for 2D and 3D Architectures

Generations <sup>6</sup>		1	2	3	4	5
Raw Bandwidth Density (GBps/mm) <sup>7,8,9</sup>		125	250	500	1000	2000
Minimum Bump Pitch ( $\mu\text{m}$ )		55	50	40	35	30
Package Technology	IO/mm <sup>10</sup>	500	667	1000	1500	2000
	/mm <sup>2</sup>	331	400	625	816	1111
Signaling Speed (Gbps) <sup>11</sup>		2	3	4	5.33	8

Table 4: Interconnect Roadmap for Solder-based Interconnects


Raw Bandwidth Density (Gbps/mm <sup>2</sup> )	5	10	20	40	80
Package Technology					
Signaling Speed (Gbps)					

For 100mm<sup>2</sup> chiplet (GPU is 700mm<sup>2</sup>)  
 Generation 1: 5 Terabytes/second D2D bandwidth  
 Generation 5: 80 Terabytes/second D2D bandwidth

*Using modest 8Gbit/sec clock rate*

Table 5: Interconnect Pitch Roadmap. Interconnects will transition away from solder gradually as pitches scale to less than 30 $\mu\text{m}$ .


# AMDs Chiplets Strategy



Monolithic 32-core Chip  
777mm<sup>2</sup> total area  
1.0x Cost



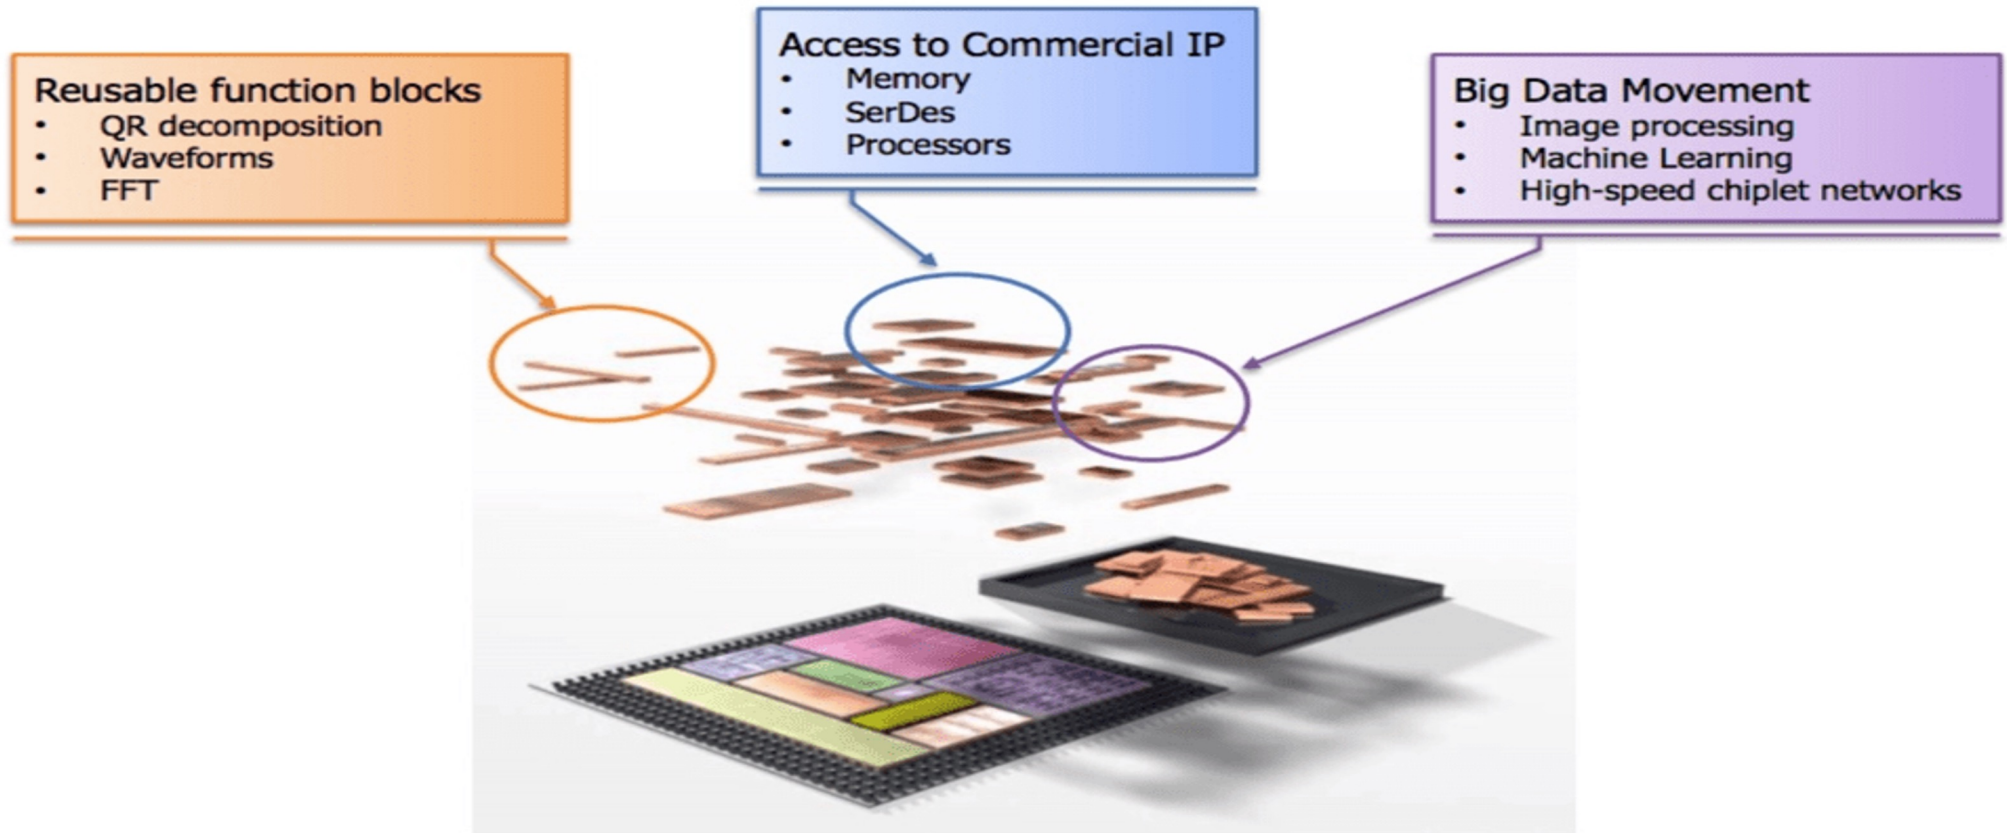
4 x 8-core Chiplet, 213mm<sup>2</sup> per chiplet  
852mm<sup>2</sup> total area (+9.7%)  
0.59x Cost



Naffziger, S., Noah Beck, T. Burd, K. Lepak, Gabriel H. Loh, M. Subramony and Sean White. "Pioneering Chiplet Technology and Design for the AMD EPYC™ and Ryzen™ Processor Families : Industrial Product." 2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA) (2021): 57-70.




## How do chiplets enable domain specialization?



CHIPS modularity targets the enabling of a wide range of custom solutions

- Home
- Topics
- Sectors
- Exascale
- Specials
- Resource Library
- Podcast**
- Events
- Solution Channels
- Job Bank
- About
- Subscribe



June 15, 2022

AMD is getting personal with chips as it sets sail to make products more to the liking of its customers.

The chipmaker detailed a modular chip future in which customers can mix and match non-AMD processors in a custom chip package.

“We are focused on making it easier to implement chips with more flexibility,” said Mark Papermaster, chief technology officer at AMD during the analyst day meeting late last week.

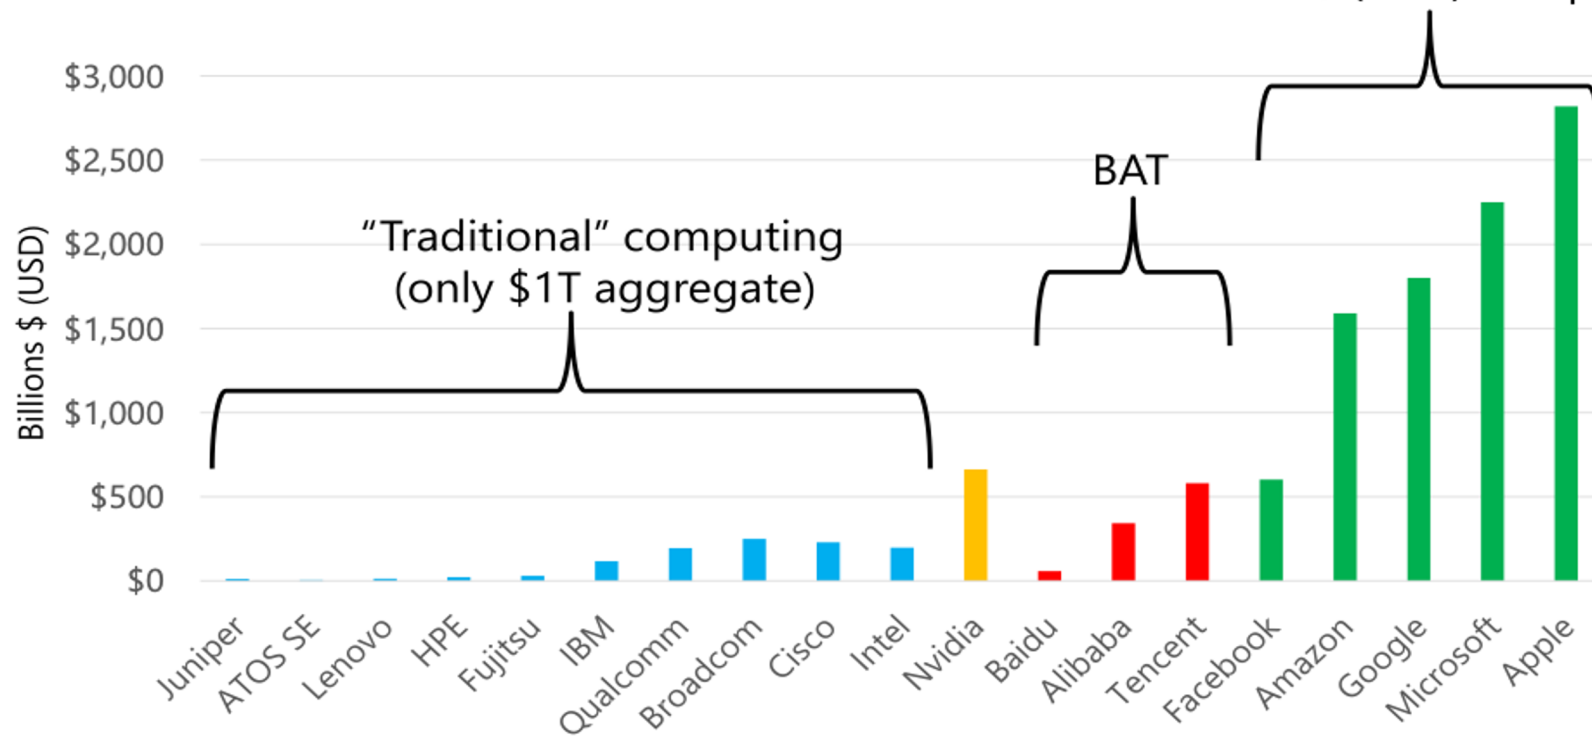


AMD will allow customers to implement multiple dies — also called chiplets or compute tiles — in a tight chip package. AMD already uses tiles, but is now welcoming third parties to make accelerators or other chips to be included in 2D or 3D packages alongside its x86 CPUs and GPUs.

# Why? It is not good enough anymore to understand the technology Now we must also understand the market context

Dan Reed, Jack Dongarra, Dennis Gannon, 2022  
<https://arxiv.org/pdf/2203.02544.pdf>


Control of the computing ecosystem  
Trillion+ \$ (USD) companies



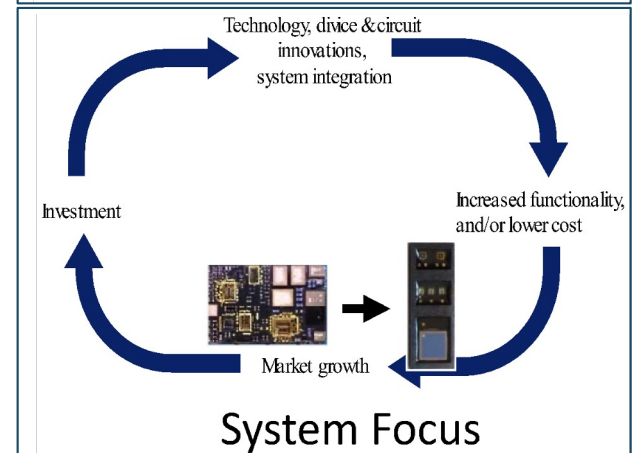

# Why? Domain specific Architectures driven by hyperscalers

*in response to slowing of Moore's Law (switch to systems focus for future scaling)*

Dharmesh Jani, Facebook –  
ODSA Workshop, Regional Summit, Amsterdam, Sep. 2019




AI/ML/data workload explosion needs DSAs




# Who? :ODSA: Open Domain Specific Architecture

## Creating an Open Chiplet Marketplace for Hyperscale Datacenters


**Open D2D Interface**  
Reduce barrier to interoperation



**Reference Designs**  
Starting point for new designs



**Reference Workflows**  
Reusable, open practices






- Package/Integration Partner
- RF Product Partner
- Foundry
- ASIC Provider
- ASIC High-Level Design Partner

**ODSA Activities**



**Chiplet Marketplace**

Integrate best-in-class chiplets from multiple vendors through open interfaces



OCP modular form factors

# Benefits of Chiplets

- 1. Not everything benefits from being on a leading edge design node**
  - Analog electronics, I/O (PCIe), and memory controllers get worse & crazy expensive (can do it in the appropriate design node)
- 2. Don't need to respin an entire chip just to change one component**
- 3. Modularity (Modular HPC effort with ODSA, Intel, Meta, Google, and DOE)**
- 4. Better yield (known good die)**
- 5. Break reticle limit**
- 6. Can co-integrate diverse accelerators with high bandwidth**
- 7. Concentrate DOE/HPC effort on the part of the chip that really matters to our applications (80:20 rule)**
- 8. Leverage the Hyperscale economic force / supply chain (OCP)**
- 9. Enables better IP reuse (*combined with the licensable IP market*) and economies of scale for the broadly reusable chiplets**

# Chipelets Summit: The HPC edition!



January 23, 24, 2023  
San Jose, California

- **Organizing “chipelets for HPC” for DOE and IC**  
*(with Jeff Vetter@ORNL and Bapi Vinnakota@OCP/ODSA/LBL, and others!!)*
  - Targeting summer 2023 (TBD)
  - Co-sponsored by OCP/ODSA
  - Bay area because will have hyperscale system architects participating as well
- **First meeting more of a level-set**
  - Follow-up will be to give community time to develop ideas and associated whitepapers
  - ~Fall 2023

