

Heterogeneous Integration for HPC

EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA



Heterogeneous Integration for HPC



SERVER

George Michelogiannakis
Research Scientist
Lawrence Berkeley National Laboratory




OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Projected Performance Development



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.


Specialization: More Performance in a 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



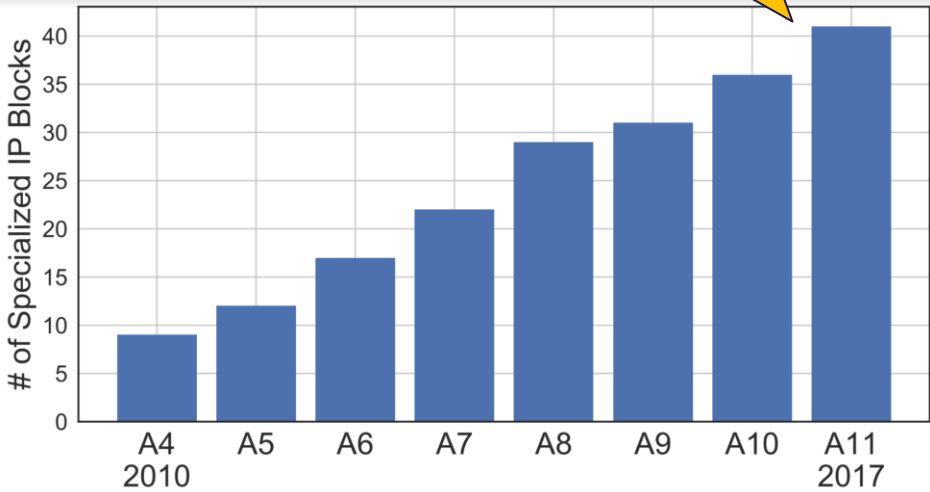
OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Already Happening in Datacenters and Consumer Electronics

40+ different heterogeneous accelerators in Apple A11 (2019)

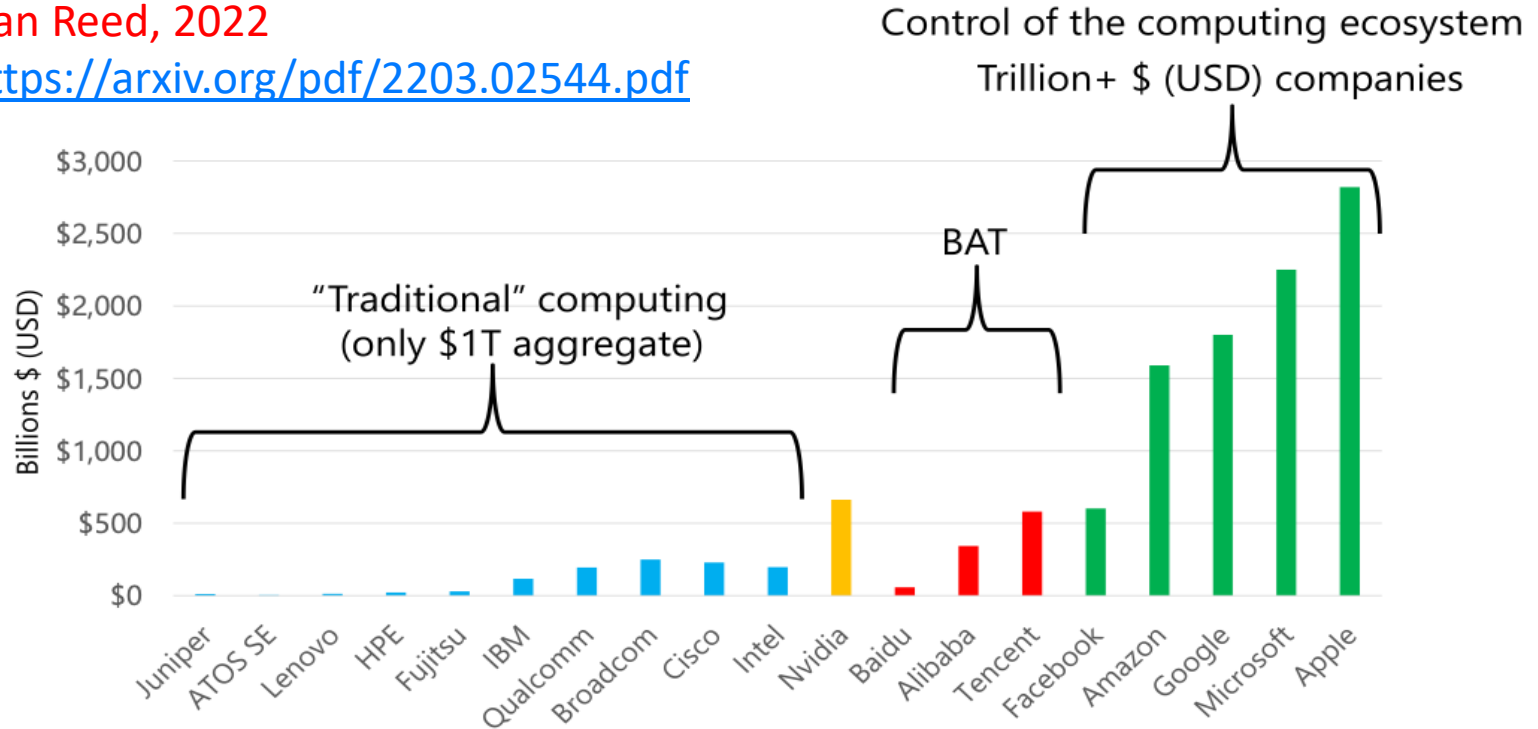
Google TPU:
Deployed in datacenters since 2015



But Those Are Large Markets

Dan Reed, 2022

<https://arxiv.org/pdf/2203.02544.pdf>



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Opportunity for HPC: New Economic Model

Open Chipllets Marketplace is forming (ODSA and UClexpress)

- Licensable IP and assembly by 3rd party lowers that barrier

Leverage this baseline and extend to support HPC

- Smaller incremental cost for HPC to “play”
- *HPC has become “too small to attack the city”*

Leverage commodity ecosystems

- Get **licensed IP** where there is a market to support it
- Use **open-source IP** where the government needs to develop technology to serve its needs
- Partner with chip vendors to realize systems






OCP
GLOBAL
SUMMIT


OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

What is a Chiplet?



How Do Chiplets Help?



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

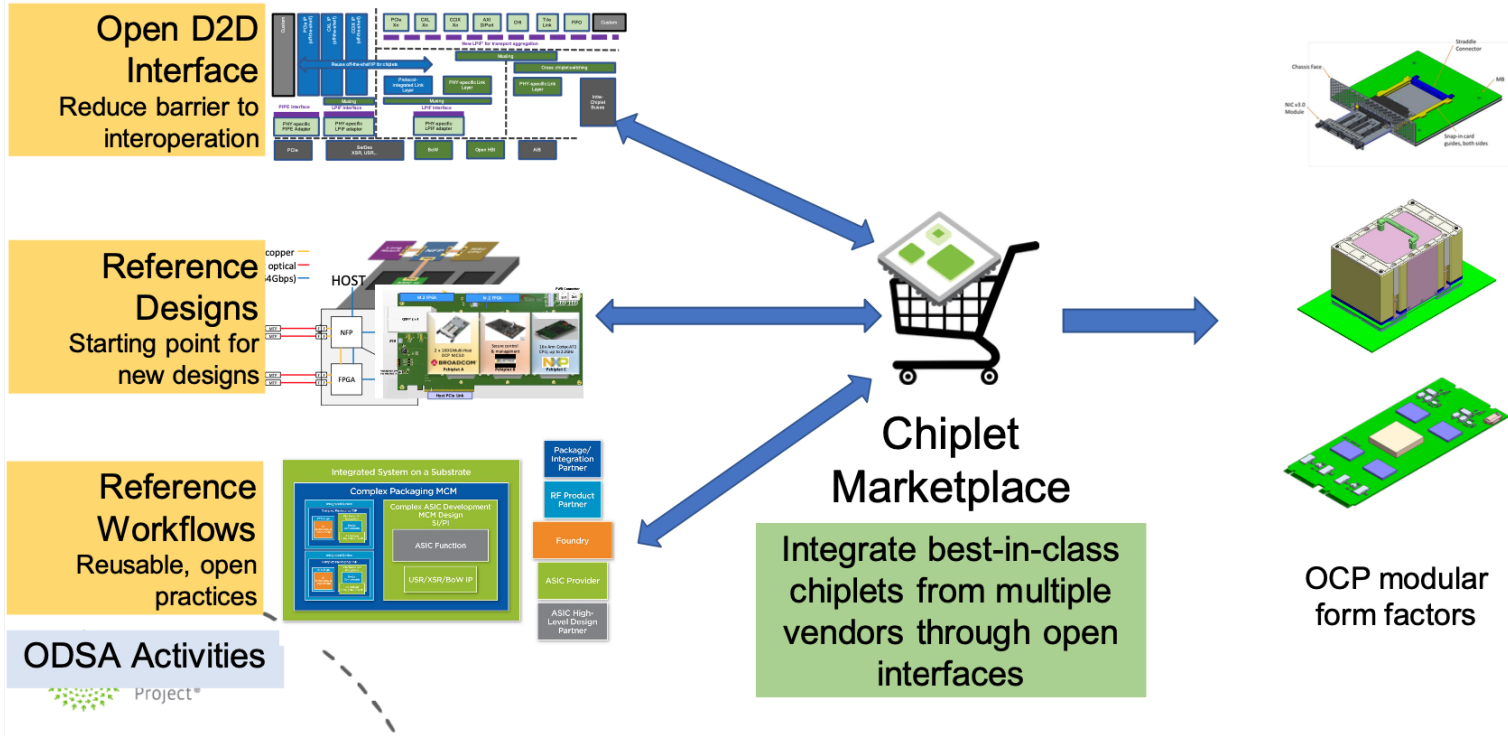


OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

This is Already Happening in ODSA: Open Domain Specific Architecture



[Call to Action]

- Support the development of open-source hardware IP for HPC
- Identify which compute accelerators are HPC-specific and thus should be open source, and which commercialized
- Develop design system-wide management methods to reduce underutilization
- Develop programming abstractions and accelerator usage libraries
- Contact the speaker: mihelog@lbl.gov



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Thank you!

EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

