What Can Chiplets Bring to Multi-Tenant Clouds?

Jiechen Zhao Natalie Enright Jerger Mingyu Gao

UNIVERSITY OF TORONTO

Tsinghua University

Introduction: Chiplets in the Cloud

- Purposes
  - Enabling cost-efficient services
    - performance/TCO
    - 100s millions USD
  - Decarbonizing datacenters
    - J/bit or J/operation
    - Up to ~100MW
- Chiplet-based design philosophy
- Outline of this talk

Introduction: Chiplets in the Cloud

- Purposes
  - Enabling cost-efficient services
  - Decarbonizing datacenters
- Chiplet-based design philosophy
  - Smaller die sizes
  - Chip disaggregation
- Outline of this talk

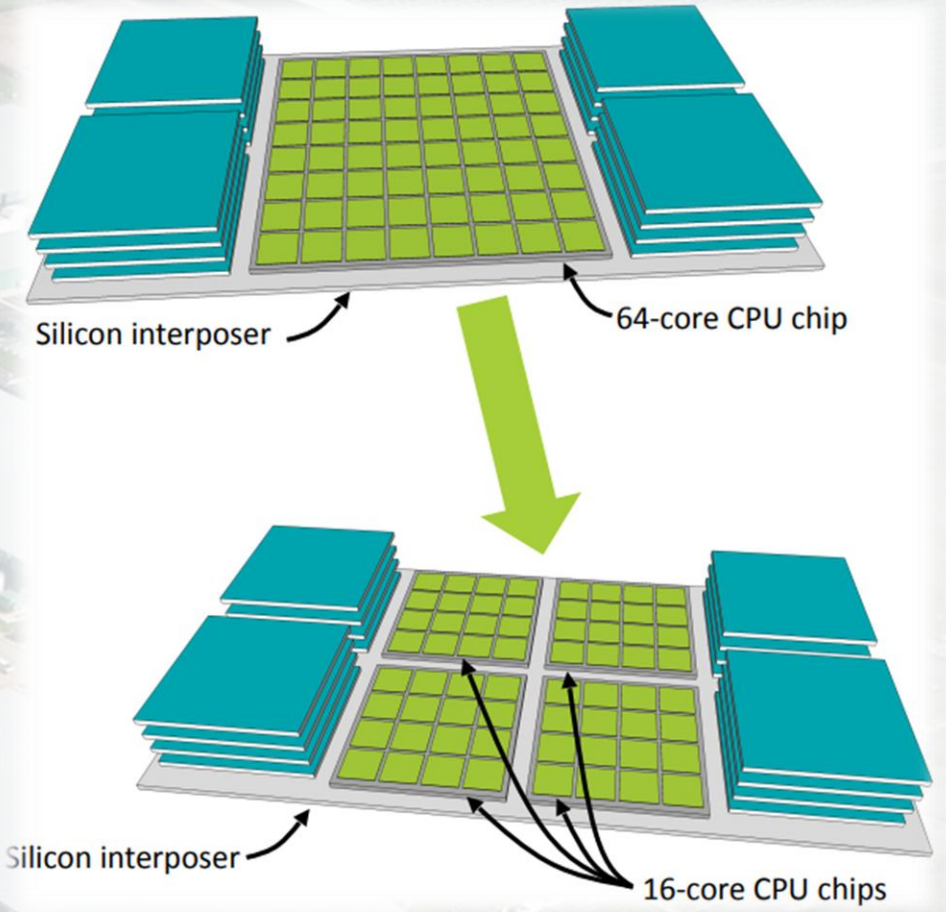

Introduction: Chiplets in the Cloud

- Purposes
  - Enabling cost-efficient services
  - Decarbonizing datacenters
- Chiplet-based design philosophy
  - Smaller die sizes
  - Chip disaggregation
- Outline of this talk
  - Chiplets for cloud hardware
  - Chiplets for multi-tenant clouds
    - Memory, interconnect in server designs
    - Isolation/security management in system designs

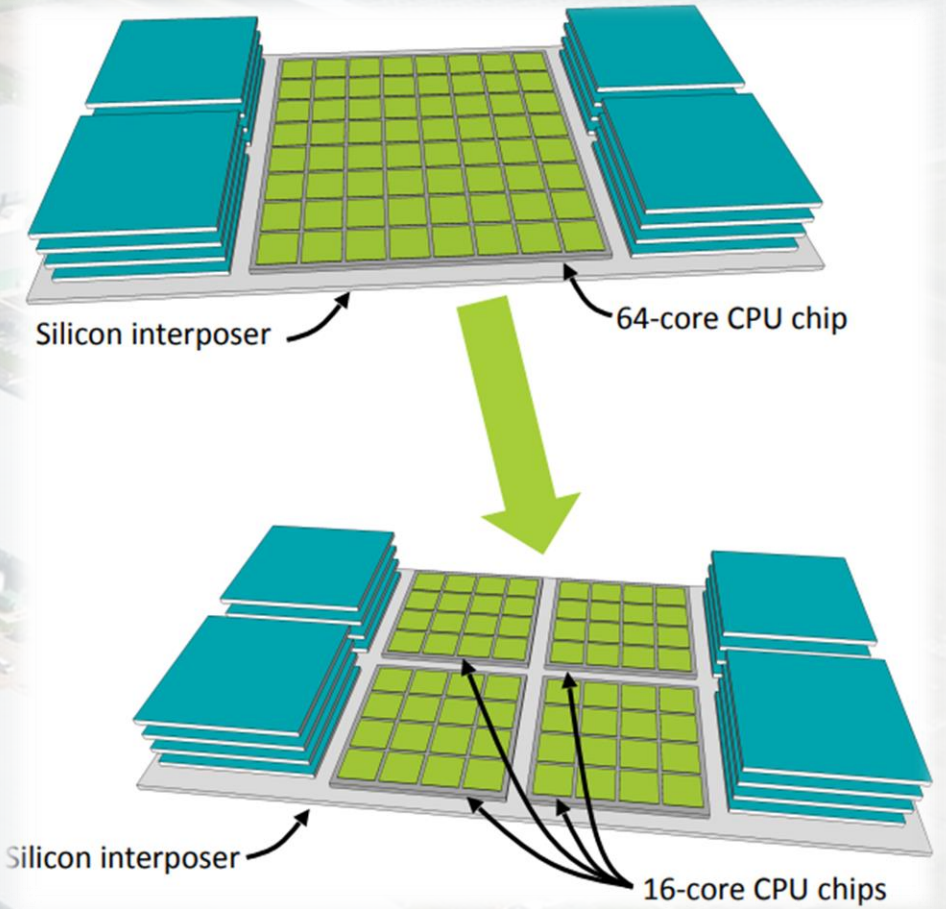

Why Chiplets for the Cloud?

😔 Silicon out of steam

- 15% per year [David Brooks]
- 7nm development prohibitively cost

😄 Smaller dies -> lower manufacturing cost

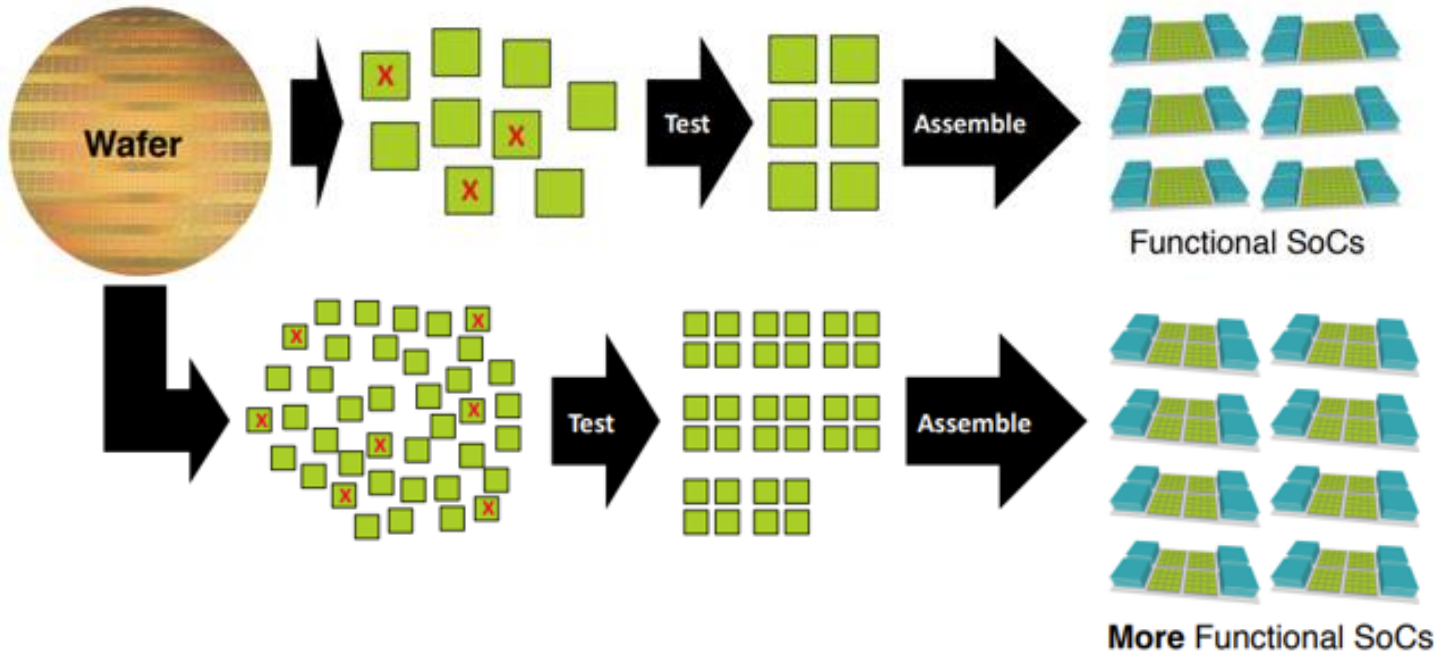

Why Chiplets for the Cloud?

- 😔 Silicon out of steam
  - 😁 Smaller dies -> lower manufacturing cost
- 😔 Challenges in heterogeneous SoCs
  - 😁 Chip disaggregation -> lower design cost

Chiplets allow design reuse and decoupled developments for various IPs!

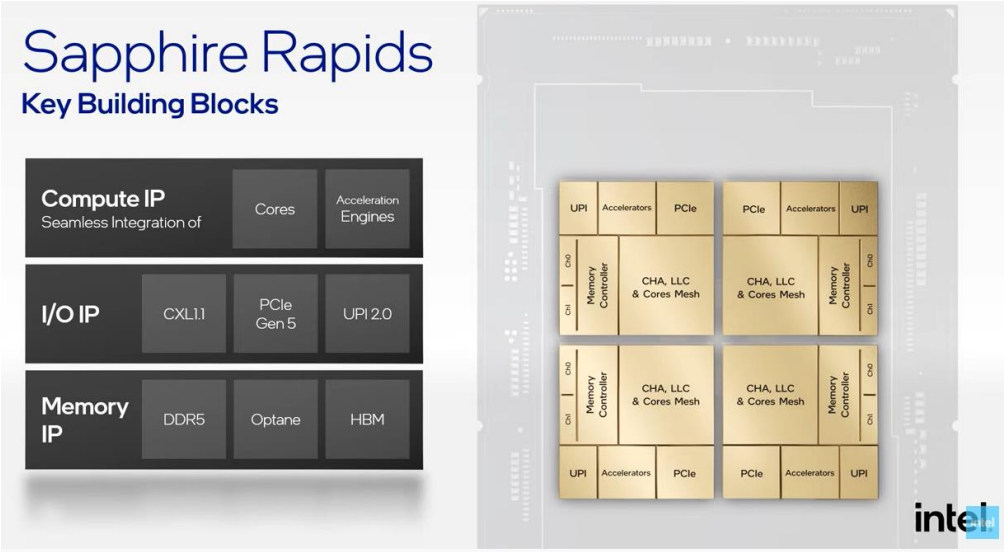

Introduction: Chiplets in the Cloud

- Purposes
  - Enabling cost-efficient services
  - Decarbonizing datacenters
- Chiplet-based design philosophy
  - Smaller die sizes
  - Chip disaggregation
- Outline of this talk
  - ~~Chiplets for cloud hardware~~
  - Chiplets for multi-tenant clouds
    - Memory, interconnect in server designs
    - Isolation/security management in system designs

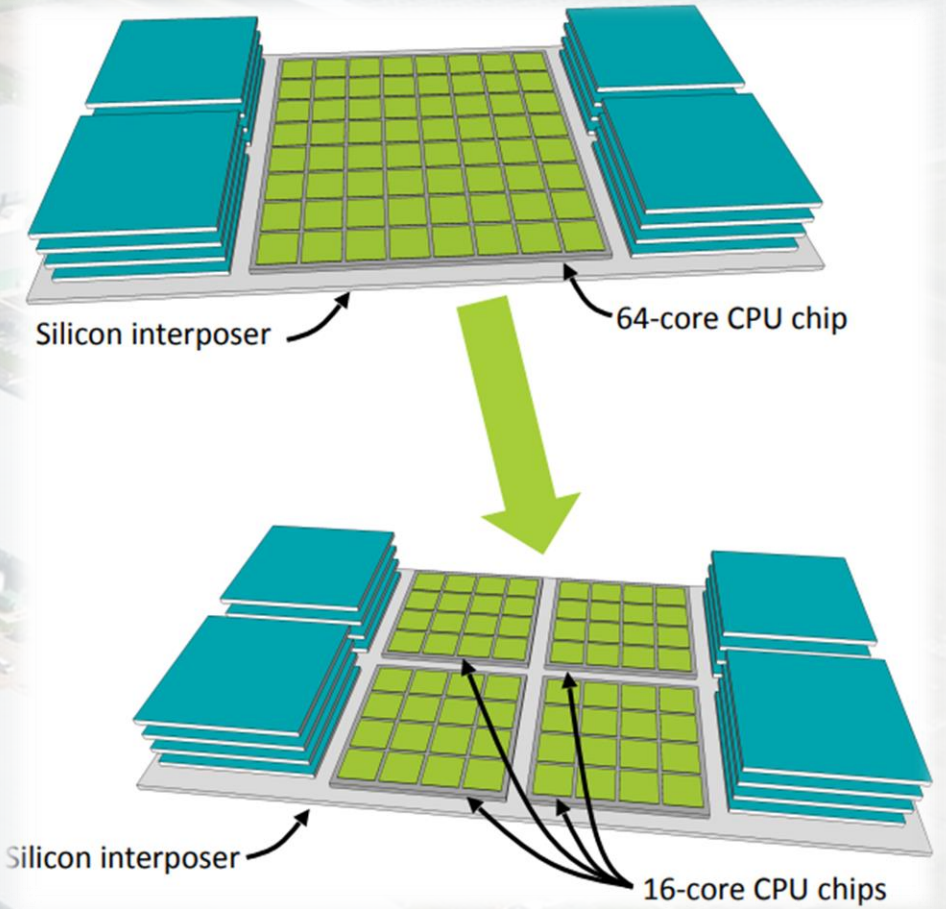

Conclusion: Why Chiplets for the Cloud?

- 😔 Silicon out of steam
  - 😁 Smaller dies -> lower manufacturing cost
- 😔 Challenges in heterogeneous SoCs
  - 😁 Chip disaggregation -> lower design cost
- 😔 Big data amount and costly data movement
  - 😁 Memories closer to compute -> Better use of DRAM -> higher cost/energy efficiency
  - 😁 High-speed and flexible interconnects -> Lower communication cost

Importance of Memory Integration

Experiment settings for tail latency V.S. throughput
- 256 tenants
- Independent Poisson request distribution
- Our configuration: Assume memory bandwidth is infinite

Memory access paths
- Green path: On-board DRAM as a cache -> perf. 😄
- Purple path: Host DRAM access -> perf. 😔

Tenants contend for
- On-board memory (B/W, capacity
- PCIe bandwidth

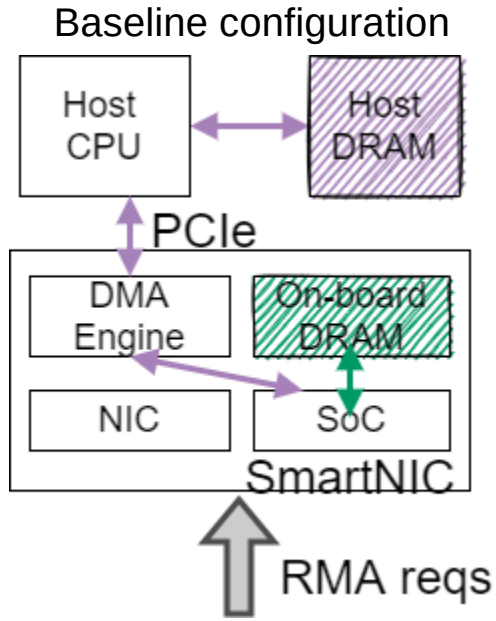

Need: Integrated memories closer to where data is consumed

Throughput (Million Requests per Second)

Memories “Closer” to Heterogeneous Compute

Data access much more expensive
than arithmetic operation

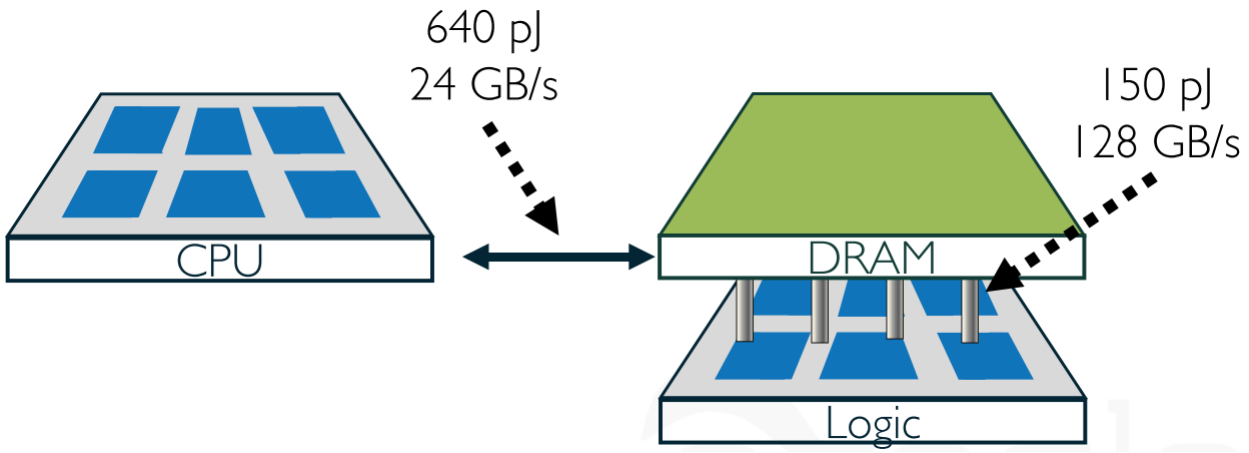

Source: Babak Falsafi

Memories “Closer” to Heterogeneous Compute

Data access much more expensive than arithmetic operation

“Closer” means

- Shorter but wider signals

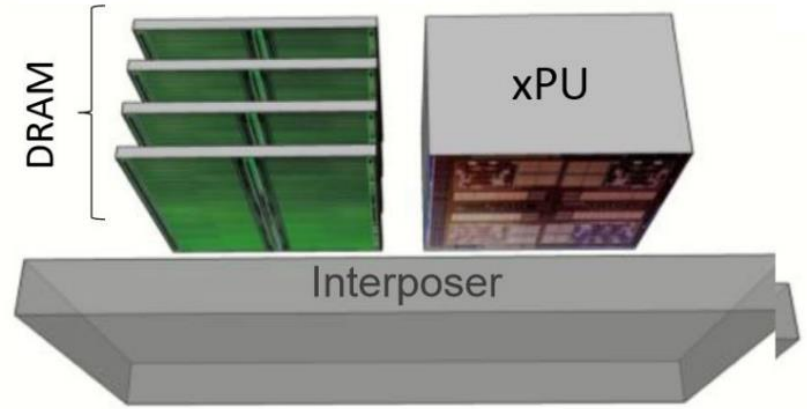

INTERPOSER STACKING (2.5D)

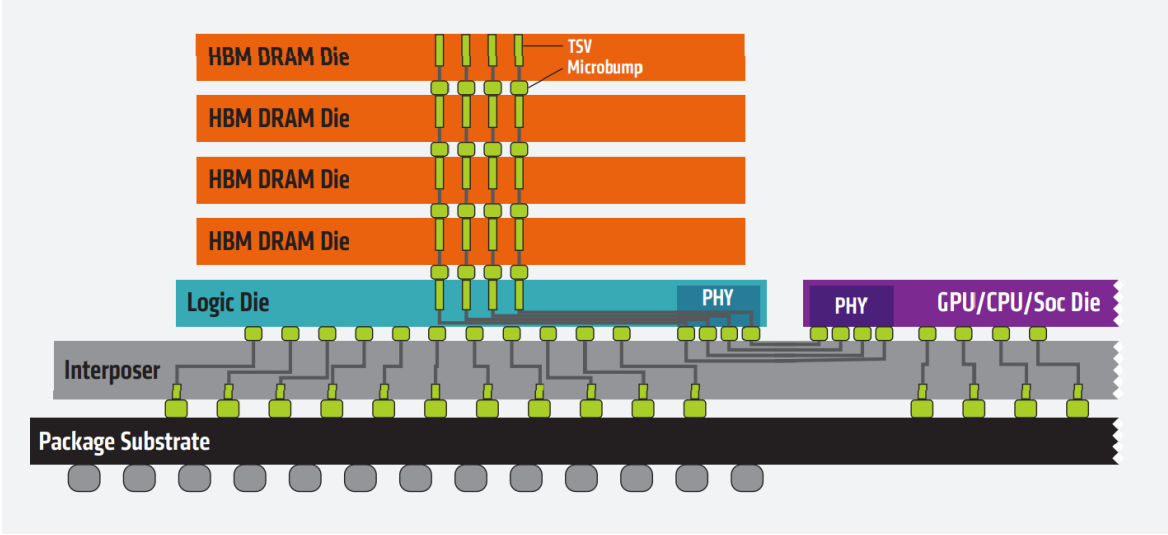

Memories “Closer” to Heterogeneous Compute

Data access much more expensive than arithmetic operation

“Closer” means

- ➢Shorter but wider signals
- ➢Coherent data sharing
  - ➢Memory hierarchy integration

Integrate compute devices in the memory hierarchy is a key

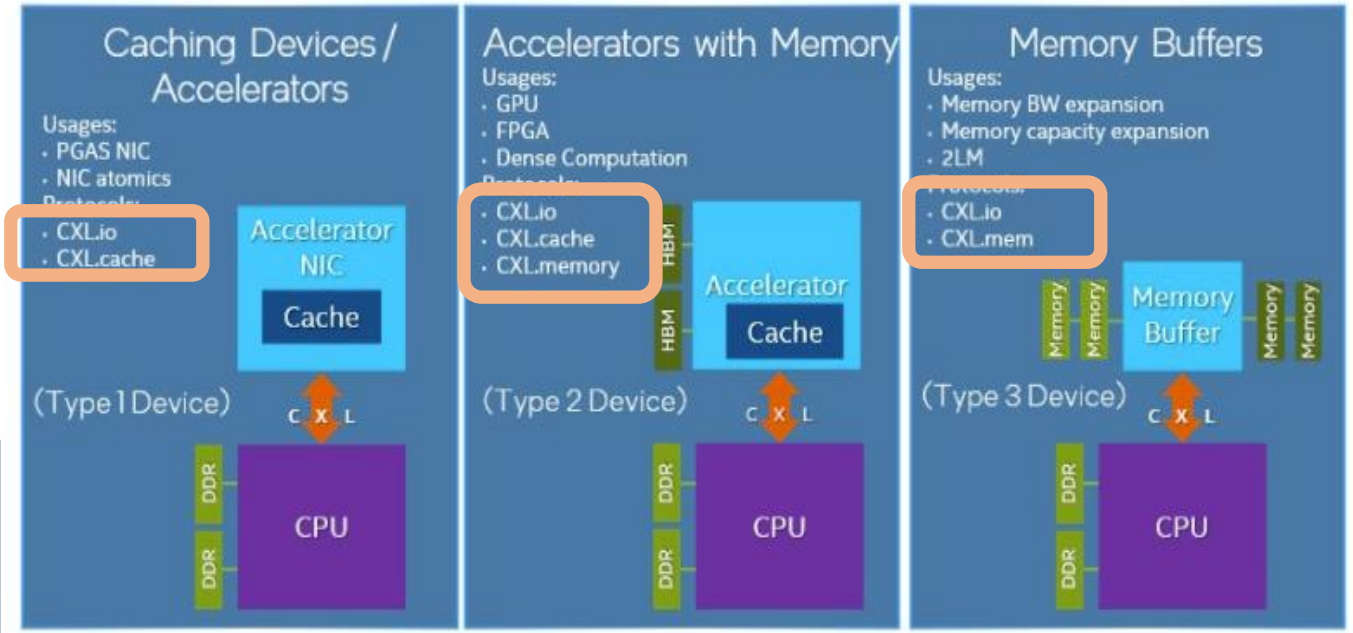

Source: CXL

Flexible and High-Speed Interconnect

Communication patterns in the server is getting more complex

- ➢Localize short communication on-chiplet
- ➢Customize inter-chiplet interconnects

Interposer offers additional routing logics

- ➢Metal layers, passive/active
- ➢[ButterDonut, IEEE Micro'16]: Customize interconnects for high memory bandwidth

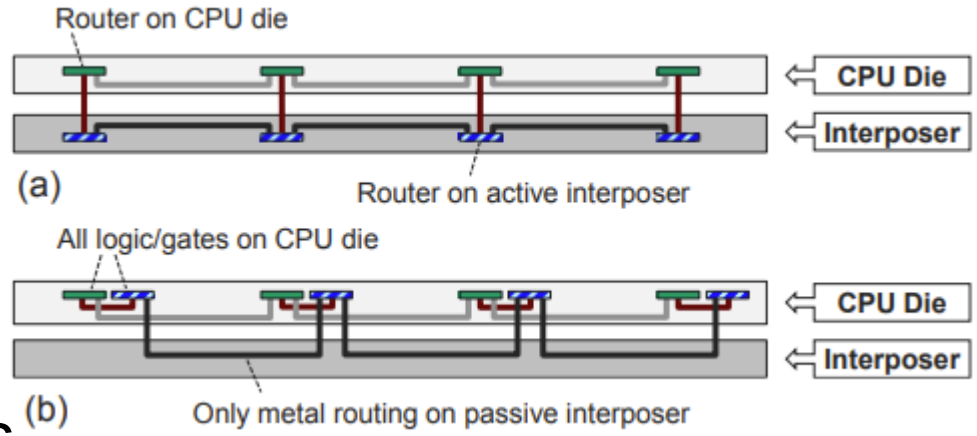

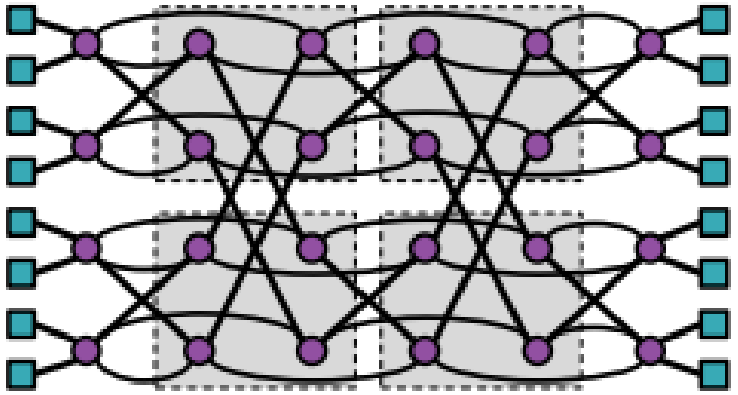

[ButterDonut, IEEE Micro'16]

Flexible and High-Speed Interconnect

Experiment settings for latency V.S. interconnect throughput

- 16 compute dies, each with 16 cores, 1 I/O die connecting Ethernet fabrics
- Each chiplet has an interposer router, 16 interposer routers as a ButterDonut
- 256 tenants
- Independent Poisson request distribution

Bottleneck of ButterDonut for in-memory workloads

- Ingress/egress traffic between the I/O die and compute dies

Aa separate fat-tree interconnect
->1.5x improvement

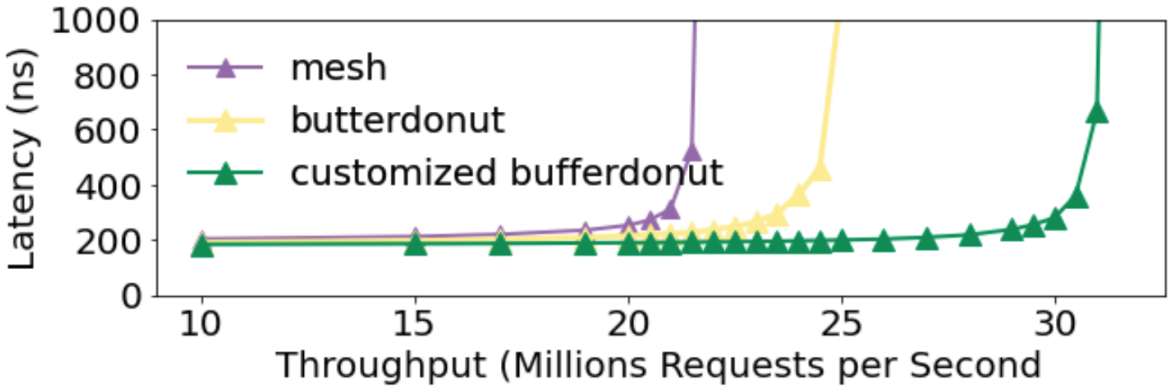

Introduction: Chiplets in the Cloud

- Purposes
  - Enabling cost-efficient services
  - Decarbonizing datacenters
- Chiplet-based design philosophy
  - Smaller die sizes
  - Chip disaggregation
- Outline of this talk
  - ~~Chiplets for cloud hardware~~
  - Chiplets for multi-tenant clouds
    - ~~Memory, interconnect in server designs~~
    - Isolation/security management in system designs

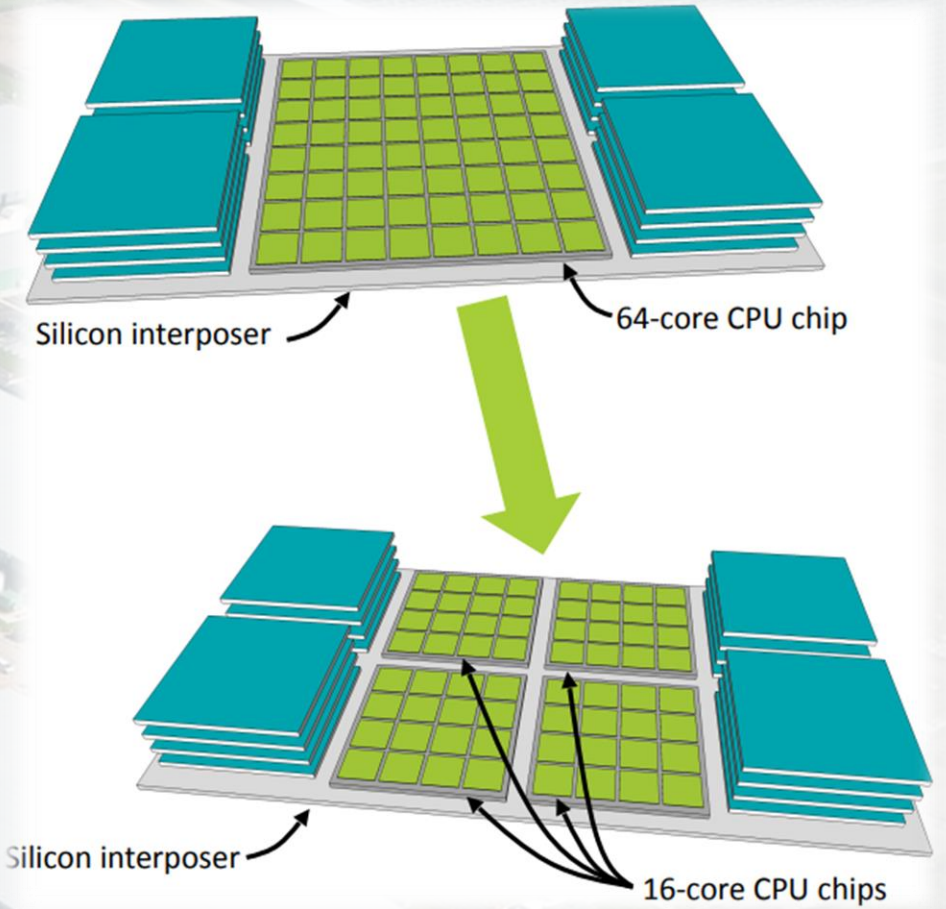

Why Chiplets for the Cloud?

- 😔 Silicon out of steam
  - 😁 Smaller dies -> lower manufacturing cost
- 😔 Challenges in heterogeneous SoCs
  - 😁 Chip disaggregation -> lower design cost
- 😔 Big data amount and costly data movement
  - 😁 Memories closer to compute -> Better use of DRAM -> higher cost/energy efficiency
  - 😁 High-speed and flexible interconnects -> Lower communication cost
- 😔 Multi-tenancy challenges workload isolation and security issues
  - 😁 Chiplets provide natural and physical isolation

Chiplets as A New Dimension of Isolation

First-class constraints in multi-tenancy

- ➢Workload isolation
- ➢Security guarantees

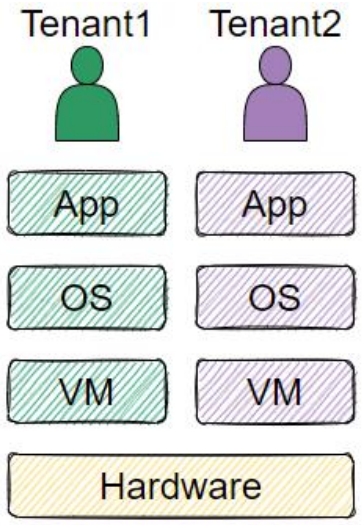

Single-tenant model

😄 Isolation

😔 Elasticity

😔 Transparency

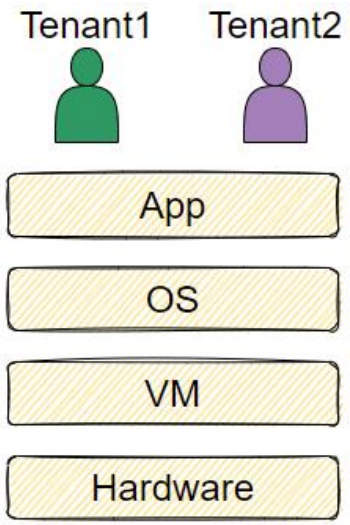

Multi-tenant model

😄 Elasticity

😄 Transparency

😔 Isolation

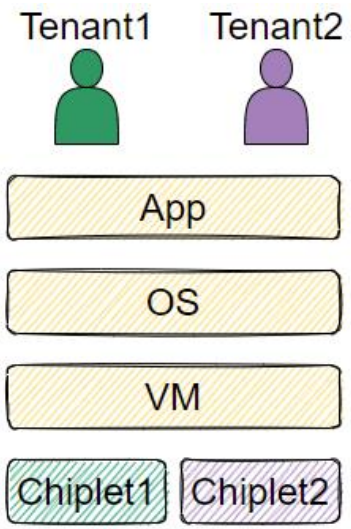

Chiplet-aware Multi-tenant model

😄 Elasticity

😄 Transparency

😄 Isolation

Chiplets as A New Dimension of Isolation

Experiment settings
- Workloads: Silo, Masstree
- Baseline configuration
  - Silo + Masstree co-location
  - 40-core, 28MB LLC
- Chiplet-based configuration
  - Silo on chiplet1, Masstree on chiplet2
  - Each chiplet: 20-core, 14MB LLC

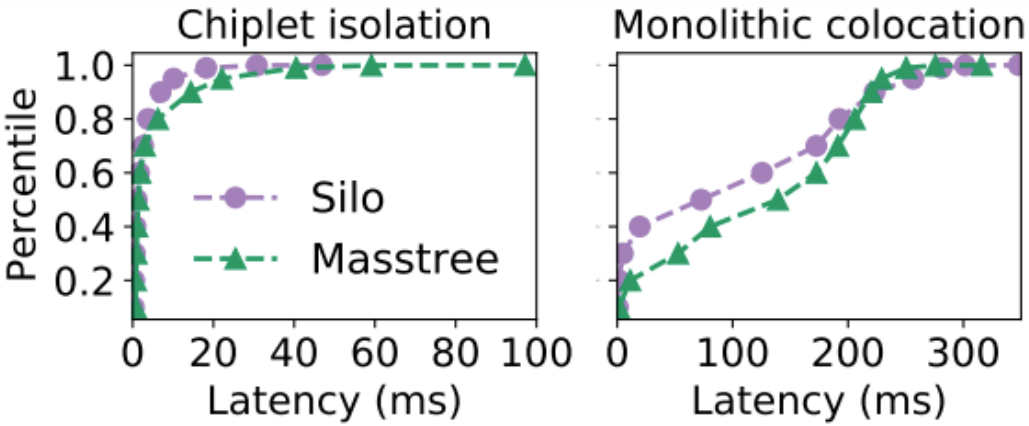

Benefits result from
- Interference removed b
- Cache less overprovisi
- Faster instruction supp

Simple chiplet-workload mapping can compromise management complexities

Blindly increasing one type of resource does not help if interference still exists

Chiplets as A New Dimension of Isolation

To isolate resources for security

- OS- and hardware level mechanisms not sufficient for security [Bolt, IEEE Micro'18]
- Chiplets can isolate previously shared microarchitecture states

Conclusion: Chiplets for the Cloud?

- 😔 Silicon out of steam
  - 😁 Smaller dies -> lower manufacturing cost
- 😔 Challenges in heterogeneous SoCs
  - 😁 Chip disaggregation -> lower design cost
- 😔 Big data amount and costly data movement
  - 😁 Memories closer to compute -> Better use of DRAM -> higher cost/energy efficiency
  - 😁 High-speed and flexible interconnects -> Lower communication cost
- 😔 Multi-tenancy challenges workload isolation and security issues
  - 😁 Chiplets provide natural and physical isolation

Thank you!

Jiechen Zhao Natalie Enright Jerger Mingyu Gao

UNIVERSITY OF TORONTO

Tsinghua University