



DESIGN, AUTOMATION & TEST IN EUROPE

14 – 15 March 2022 · on-site event  
16 – 23 March 2022 · online event

The European Event for Electronic  
System Design & Test

# MemPool-3D

## Boosting Performance and Efficiency of Shared-L1 Memory Many-Core Clusters with 3D Integration

**Matheus Cavalcante**, Anthony Agnesina, Samuel Riedel, Moritz Brunion,  
Alberto García-Ortiz, Dragomir Milojevic, Francky Catthoor, Sung Kyu Lim, Luca Benini

**ETH** zürich



Universität  
Bremen

imec



ALMA MATER STUDIORUM  
UNIVERSITÀ DI BOLOGNA



**PULP**  
Parallel Ultra Low Power


# MemPool: PULP's scaled-up shared-L1 system

- Shared-L1 memory is a very common architectural pattern
  - Only scaled-up to a few tens of cores
- **MemPool** takes this to the extreme
  - **256 cores**, sharing **1 MiB of L1**, divided into 1024 SRAM banks, in **5 cycles** of latency
  - **500 MHz** (w.c.) at GlobalFoundries **22FDX** technology
  - And **open-source**, as it should be!
    - [pulp-platform/mempool](https://pulp-platform.org/mempool)



# Connecting 256 cores to 1024 SRAM banks

- It is no easy feat to connect 256 cores and 1024 memory banks
- Hierarchical interconnection to avoid major routing congestion
  - And still, congestion is a major issue
- Low latency target is very constraining
  - We need to cross the whole macro with a single pipeline stage at the middle
  - Wire propagation delay limits the operating frequency




4.6mm × 4.6mm 22FDX MemPool cluster

# MemPool: a prime benchmark for F2F 3D ICs


- 3D Integration tackles MemPool's implementation issues
  - MemPool is limited by **wire propagation delay** → **vertical integration reduces the design footprint and wire length**
  - MemPool is **routing congested** → with **Macro-3D** we can **share the BEOs of both dies** to avoid congestion bottlenecks
- *How much PPA can we gain from MemPool-3D?*

# Tile Floorplanning and Partitioning


- MemPool-3D's tile memory die floorplan
  - 16, 32, 64, 128 KiB of L1 SPM



MemPool-3D, 16 KiB  
*Utilization: 51%*



MemPool-3D, 32 KiB  
*Utilization: 65%*




MemPool-3D, 64 KiB  
*Utilization: 89%*



MemPool-3D, 128 KiB  
*Utilization: 100%*


# Group implementations in 2D and 3D

128 KiB of L1 per tile (8 MiB total)




2D

*Footprint  
46% smaller!*




3D

# Groups: MemPool-2D vs. MemPool-3D




# Groups: MemPool-2D vs. MemPool-3D





# Putting it all together: Energy Efficiency

- *MemPool-3D consistently outperforms MemPool-2D*
  - Smaller footprint leading to fewer buffers, shorter wire length, and smaller power consumption
- Larger L1 capacity → decreased energy efficiency
  - Larger SRAM banks, larger leaking





DESIGN, AUTOMATION & TEST IN EUROPE

14 – 15 March 2022 · on-site event  
16 – 23 March 2022 · online event

The European Event for Electronic  
System Design & Test

# MemPool-3D

## Boosting Performance and Efficiency of Shared-L1 Memory Many-Core Clusters with 3D Integration

**Matheus Cavalcante**, Anthony Agnesina, Samuel Riedel, Moritz Brunion,  
Alberto García-Ortiz, Dragomir Milojevic, Francky Catthoor, Sung Kyu Lim, Luca Benini

**ETH** zürich



Universität  
Bremen

imec



ALMA MATER STUDIORUM  
UNIVERSITÀ DI BOLOGNA



**PULP**  
Parallel Ultra Low Power