


# Xylem: Enhancing Vertical Thermal Conduction in 3D Processor-Memory Stacks

Aditya Agrawal, Josep Torrellas and Sachin Idgunji




University of Illinois at Urbana Champaign and Nvidia Corporation

<http://iacoma.cs.uiuc.edu>

MICRO, October 2017


# Xylem



# Motivation: Thermal Issues in 3D Stacking


## Processor-Memory Stacks:

- Reduced interconnect length and power
  - Higher memory bandwidth
  - Smaller form factors
  - Heterogeneous integration
- 
- 2.5D processor-memory stacks exist
  - 3D processor-memory stacking is the future



Major challenge: Thermals

# 3D Stacking Technologies




Face-to-back (f2b) die interface (not to scale)

# Stack Organization: Memory on Top



- ✓ Processor power and I/O signals do not traverse TSVs
- ✓ Processor IR drop similar to current designs
- ✓ DRAM and processor die floorplans are independent because TSV count and location are governed by stacked DRAM standards

✗ Thermal challenges



# Contributions

- Identify the thermal bottleneck in 3D stacks: **Die-to-Die(D2D) layers**
- Improve vertical conduction through the D2D layer:
  - **Align and short dummy  $\mu$ bumps with Thermal TSVs**
  - **Generic and custom TTSV placement schemes**
- Use the resulting thermal headroom:
  - Boost processor frequency (400-720 MHz) & performance (11-18%)
- **Exploit thermal heterogeneity**: cores closer to TTSVs conduct heat better
  - Conductivity-aware thread placement and migration, and frequency boosting

# Thermal Resistance in the Stack

- Thermal resistance per unit area:

Layer	$R_{th}$ (mm <sup>2</sup> -K/W)
Bulk Silicon	0.83
Proc. Metal	1.00
D2D	13.33

D2D layer is 13-16x more resistive than bulk silicon or metal layers


# Shortcomings of Prior Work

- Underestimated the thermal resistance of D2D layer by assuming:
  - High conductivity
  - Small thickness
- Focused on increasing the conductivity of the bulk silicon using TTSVs
- Concluded that TTSVs alone are effective

Our approach: Combine TTSVs with a mechanism to reduce D2D resistance



# Propose: Dummy $\mu$ bump-TTSV Alignment & Shorting



# TTSV Placement: Constraints

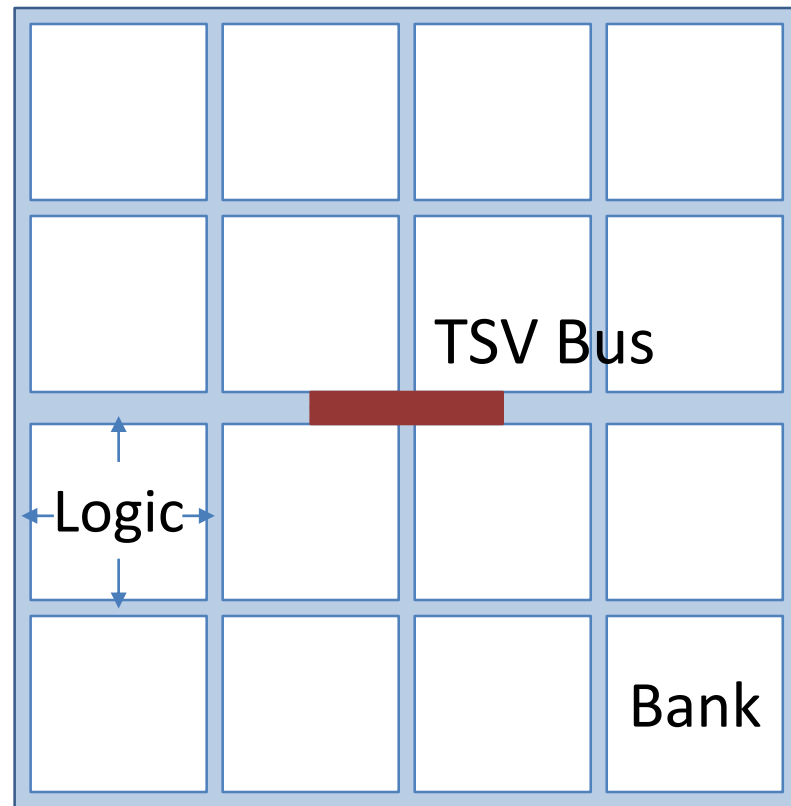
## TTSVs:

- Cannot disrupt regular DRAM arrays: Place in the DRAM peripheral logic
- Distribute TTSVs and avoid TTSV farms
- Maintain Keep Out Zone (KOZ) around each TTSV

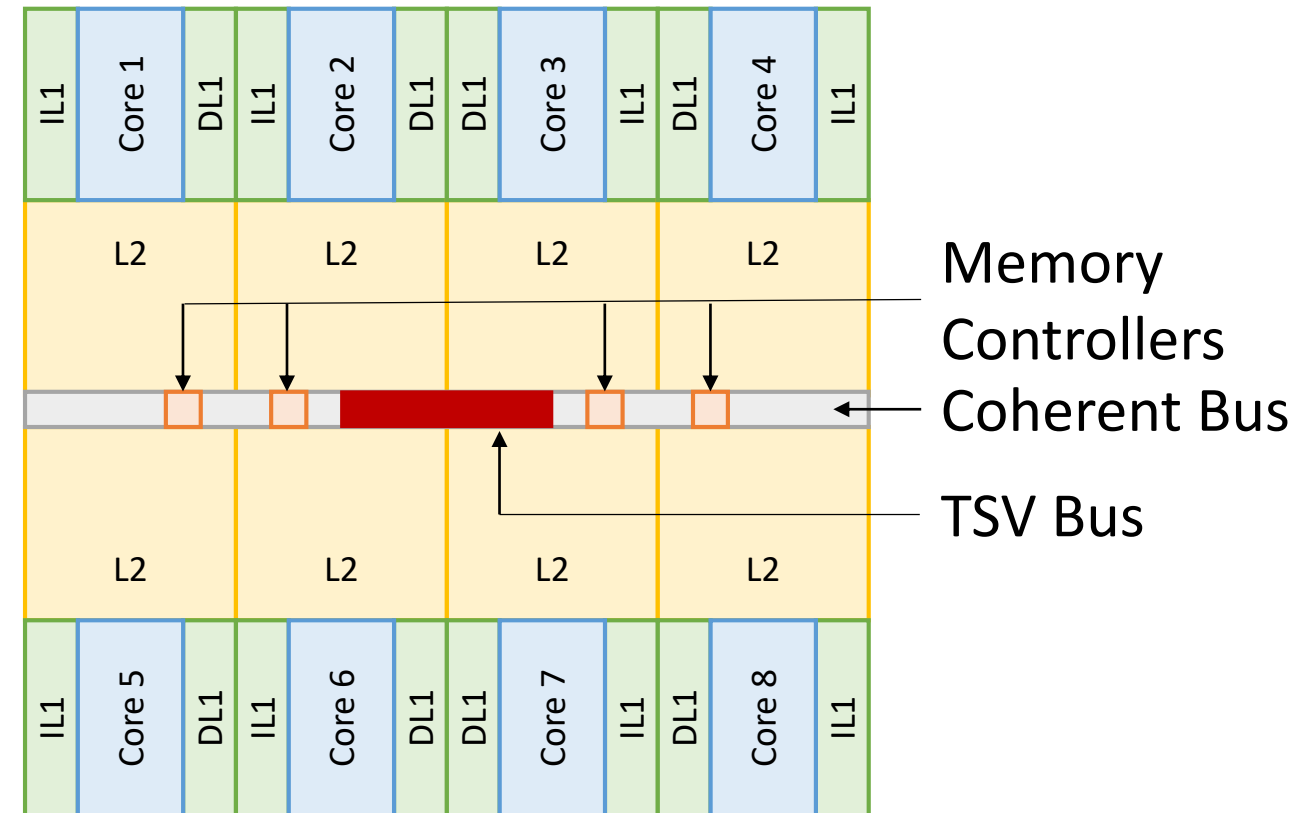
## Dummy $\mu$ bumps:

- Anywhere in the D2D layer except the electrical  $\mu$ bump locations

# DRAM and Processor Baseline Floorplans




DRAM (Wide IO) die floorplan

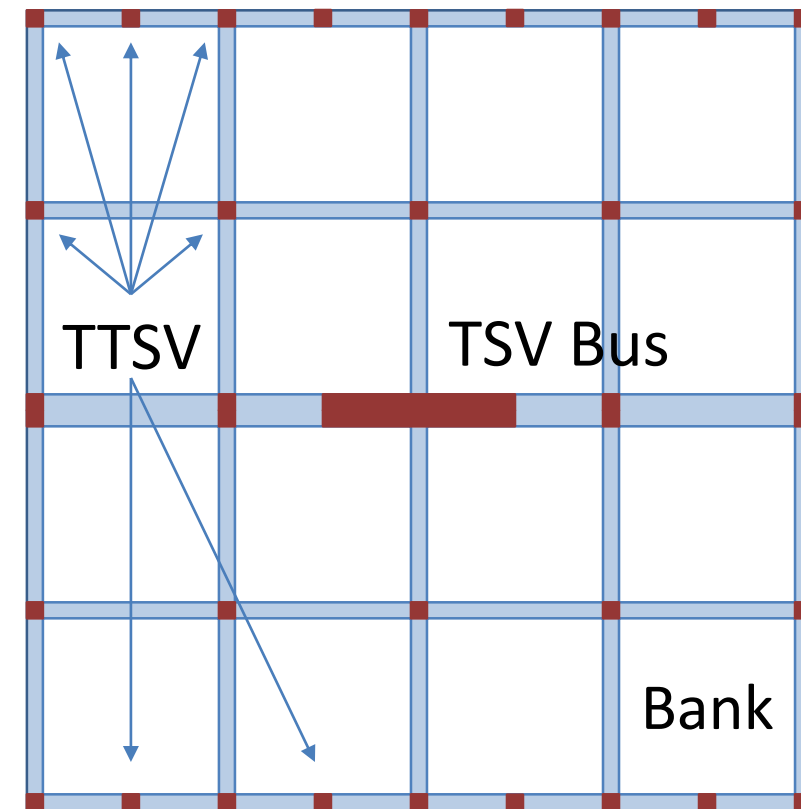


Processor die floorplan

# Proposal: TTSV Placement Schemes



Generic (oblivious to hotspots)



Custom (aligned with hotspots)

# Proposal: Frequency Boosting


- TTSV placement & TTSV- $\mu$ bump alignment and shorting:
  - Increases thermal conduction from the processor die to the heat sink
  - Reduces the temperature of the processor die
- Proposal: Increase processor frequency to consume the thermal headroom
  - Increase application performance

# Proposal: Conductivity ( $\lambda$ ) Aware Techniques

- TTSV- $\mu$ bump alignment and shorting creates high conductivity paths
  - Areas closer to TTSVs dissipate heat more easily
  - Result is thermal spatial heterogeneity in the stack
- Proposal: Three  $\lambda$ -aware optimizations to further improve performance
  - $\lambda$ -aware thread placement
  - $\lambda$ -aware frequency boosting
  - $\lambda$ -aware thread migration

# Evaluation Setup

- 8-core OoO processor die @ 2.4 GHz
- 8 high Wide IO memory on top
- Processor timing and power: SESC & McPAT
- DRAM timing and power: DRAMSim2
- Thermal analysis: 3D HotSpot
- Applications: SPLASH-2, PARSEC & NAS



Memory-on-top configuration

# Result Summary

TTSV Placement	Generic	Custom
Area Overhead	0.63%	0.81%
Proc. Temp. Reduction	5.0 °C	8.4 °C
Avg. Frequency Boost	400 MHz	720 MHz
Avg. Performance Gain	11%	18%

$\lambda$ -aware techniques enable further 100-200 MHz improvements



# Conclusion

- Identified that **D2D layer is the thermal bottleneck** in 3D stacks
- Improved vertical conduction through the D2D layer:
  - **Align and short dummy  $\mu$ bumps with TTSVs**
  - **Generic and custom TTSV placement schemes**
- Used the resulting thermal headroom to
  - Boost processor frequency (400-720 MHz) & performance (11-18%)
- **Exploited thermal heterogeneity**: cores closer to TTSVs conduct heat better
  - **Conductivity-aware** thread placement and migration, and frequency boosting
  - Enable further 100-200 MHz improvements

# Xylem: Enhancing Vertical Thermal Conduction in 3D Processor-Memory Stacks

Aditya Agrawal, Josep Torrellas and Sachin Idgunji

University of Illinois at Urbana Champaign and Nvidia Corporation

<http://iacoma.cs.uiuc.edu>