



CHALLENGES OF HIGH-CAPACITY DRAM STACKS AND POTENTIAL DIRECTIONS


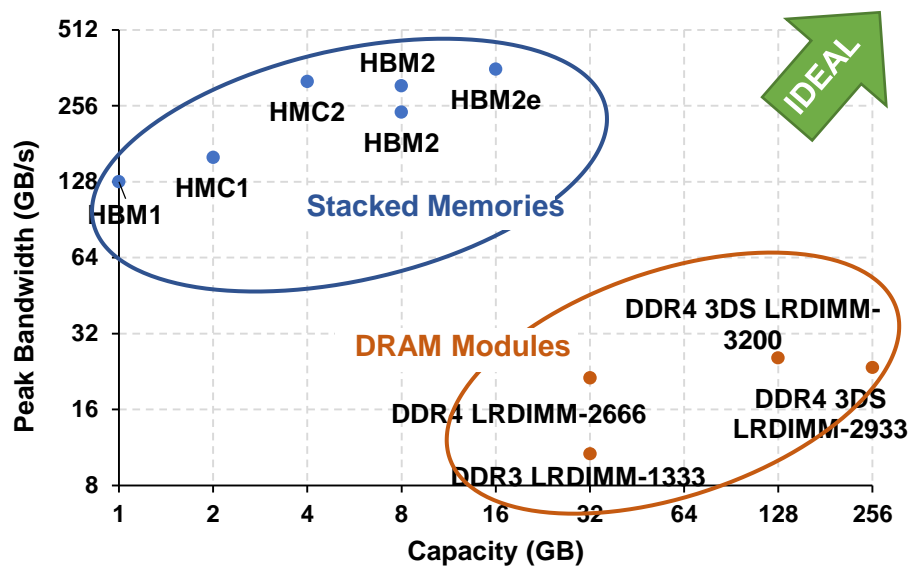
AMIN FARMAHINI-FARAHANI, SUDHANVA GURUMURTHI, GABRIEL LOH,
MIKE IGNATOWSKI (PRESENTER)
AMD ADVANCED RESEARCH, LLC

- ▲ Major challenges of “tall” DRAM stacks and potential directions to overcome/mitigate those challenges
 - TSV speed and density
 - Thermal conductivity
 - Stack height
 - Power delivery
 - Reliability
 - Design and manufacturing cost
- ▲ Key messages
 - Broad range of challenges needs to be addressed and different techniques need to be investigated to understand their tradeoffs
 - Alternative DRAM-to-DRAM bonding techniques could be considered and evaluated as a potential enabler
 - Urge researchers to participate in studying challenges and evaluating techniques for high-capacity, high-bandwidth memory stacks

HIGH-CAPACITY STACKED MEMORY (HBM)



- ▲ In-package stacked memory provides high bandwidth and low power
 - But much lower capacity than DRAM DIMMs
- ▲ High capacity can enable high performance and low energy consumption
 - Fitting larger application's working set, faster and more energy-efficient access to data
 - Less data movement between stacked memory and 2nd-level memory
- ▲ AMD GPUs are used beyond graphics for HPC and ML/AI
 - Drive for more capacity and bandwidth




Stacked Memory

Capacity and bandwidth comparison between a single memory stack and a single DRAM module

IMPROVE CAPACITY OF IN-PACKAGE STACKED MEMORY

- ▲ Higher density DRAM dies (smaller process technology, more capacity per die)
 - Scaling down DRAM cells is challenging
 - Affects yield due to larger DRAM die area
- ▲ Use of NVM dies
 - NVM-only stacks or DRAM-NVM hybrid stacks
 - Requires completely new design
- ▲ Higher number of memory stacks around processor
 - Has cost and packaging implications
 - Requires larger interposer and package substrate
 - Requires more interposer interconnects
- ▲ Higher number of DRAM dies within a stack
 - Going beyond 8 DRAM dies within a stack (taller stacks)
 - Natural evolution of current HBM trend and orthogonal to previous approaches
 - But has several challenges...



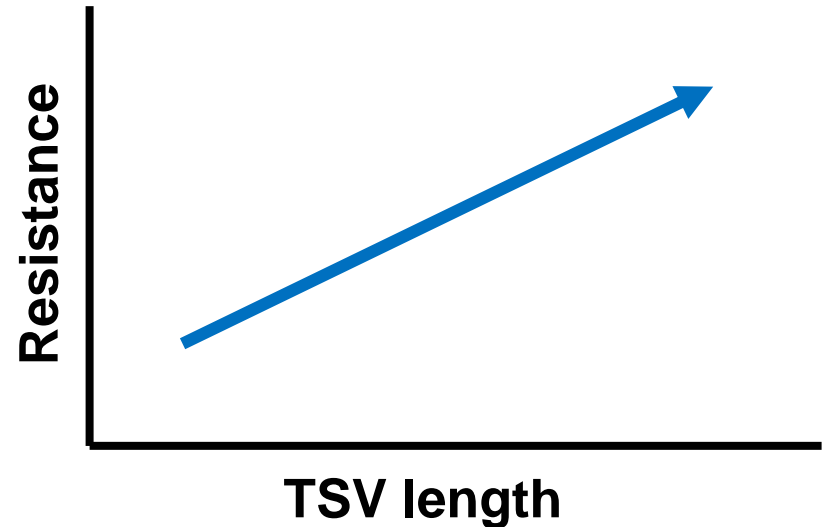
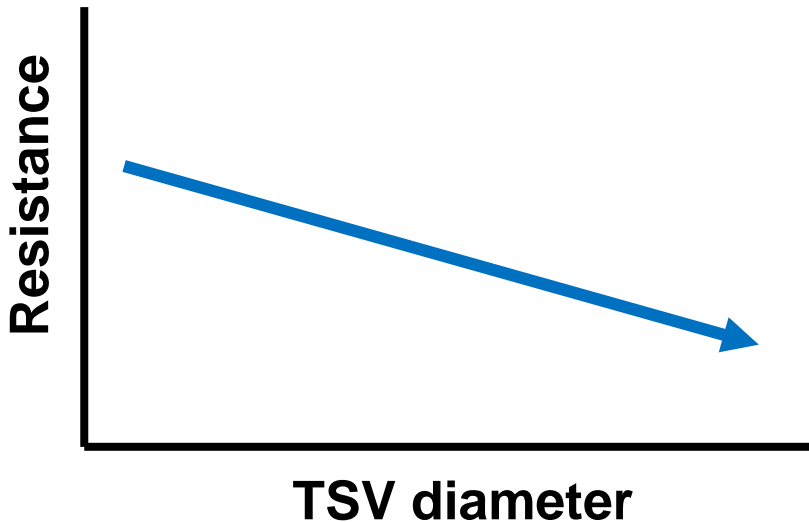

Stacked Memory

Taller stacks to improve capacity

TSV SPEED ISSUES – TSV RESISTANCE




- ▲ TSV RC delay is a function of its resistance R and capacitance C
- ▲ TSV resistance is a function of its material resistance, diameter, and length
 - First order model is the resistance of a cylinder



TSV SPEED ISSUES – TSV CAPACITANCE



- ▲ TSV RC delay is a function of its resistance R and capacitance C
- ▲ TSV capacitance is a function of TSV diameter and length, the thickness of oxide, etc.



- ▲ TSV RC delay characteristics are different from those of metal wires
- ▲ Metal wires have small capacitance to resistance ratio
- ▲ TSVs have very large capacitance to resistance ratio

TSV SPEED IN TALL STACKS



- ▲ Internal stack bandwidth is contingent on
 - TSV data rate (TSV RC delay)
 - TSV density
- ▲ As the number of DRAM layers increases, both TSV resistance and capacitance of top dies increase
 - TSV speeds can degrade with taller stacks (if TSV physical structure remains the same)
 - Capacity increase at the expense of a decrease in bandwidth! Not desirable!
- ▲ Higher TSV data rate is required for bandwidth improvement

Taller stacks would have lower TSV data rates
Must enable high-capacity AND high-bandwidth stacked memory

TSV SPEED – POTENTIAL DIRECTIONS



- ▲ Potential directions to improve internal stack bandwidth
 - Use shorter and thinner TSVs
 - Use larger TSV drivers
 - Insert data buffers to buffer TSV data in intermediate layer locations
 - Use wider TSV buses
 - Possibly allowing to slow down the TSV frequency
- ▲ Investigate potential directions and analyze their overheads and tradeoffs
- ▲ Existing DRAM-to-DRAM HBM bonding techniques may not accommodate higher internal stack bandwidth without negative cost and power implications

Higher TSV speed and density through fundamental change in TSV physical structure (shorter, thinner, denser)
=> alternative bonding techniques

- ▲ HBM provides all bits from a single row of a single bank in a single die
 - Good for power
 - Can negatively impact reliability
- ▲ Types of faults in stacked memories
 - DRAM cells
 - DRAM logic
 - TSVs (data TSVs and address TSVs)
- ▲ Faults can cause multi-bit errors
- ▲ As the number of layers within a stack increases, the probability of failure in a stack increases
- ▲ Replacing a processor package with faulty HBM is more costly than replacing a faulty DIMM

Higher fault rates may occur in taller stacks

RELIABILITY – POTENTIAL DIRECTIONS



- ▲ More robust offline and/or online test and repair schemes to detect and repair a variety of faults before causing failure
 - Enhance yield
- ▲ Stronger error detection and correction schemes to detect and potentially correct multi-bit errors
 - Additional ECC bits and pins may be needed
- ▲ Redundancy storage in logic die for new fault-tolerant schemes
 - Probably more TSVs needed
 - Design and manufacturing cost
- ▲ Chip-kill-like schemes
 - Potential impacts on energy efficiency and bandwidth
 - Design cost

Evaluate potential architectural directions to improve reliability
of future stacks

- ▲ In HBM, a large number of TSVs are used for power and ground
- ▲ Voltage drops over long resistive wires and different locations in a die and stack receive different voltages => IR-drop
 - HBM2 typical Vdd of 1.2v and minimum of 1.14v
 - Static and dynamic IR-drop
 - Power delivery issues may prevent meeting the target data rate
- ▲ Outlook
 - More power delivery challenges with an increase in the number of layers and higher bandwidth
 - Lower supply voltage in future generations
 - Smaller TSV diameters cause more resistance in power delivery network

Additional power delivery requirements for taller stacks

POWER DELIVERY – POTENTIAL DIRECTIONS



- ▲ Design a better power delivery network
 - For example through better distribution of TSVs and PG TSVs on the die edge
- ▲ Provide more power and ground TSVs
- ▲ Provide on-package voltage regulators

- ▲ All directions above potentially incur area and cost overheads

Effectiveness and cost of directions need evaluation

VERTICAL STACK HEIGHT



- ▲ Height of an 8-high HBM2 stack is 700-800 μ m
- ▲ Increasing the number of DRAM layers (and thus requiring underfill between DRAM layers) in a stack would add to the height even further
- ▲ High vertical height of future stacked memories could potentially pose packaging and thermal conductivity challenges
 - Height could be limiting for 16-high DRAM stacks
- ▲ Potential directions in existing bonding
 - Die thinning and underfill thinning (but marginal improvement)
- ▲ Alternative bonding and stacking techniques
 - Opportunities to thin DRAM dies, forgo microbumps and underfill, and better thermally conductivity

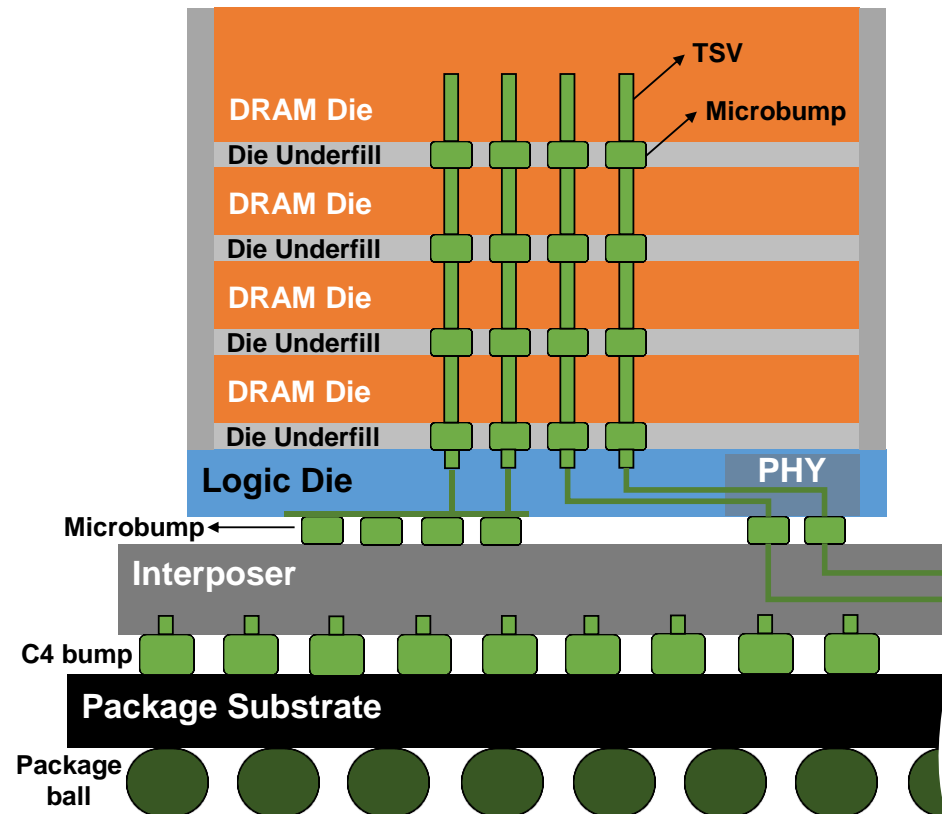
Height could be limiting for taller stacks
Alternative bonding techniques could help

DRAM-TO-DRAM BONDING



PICTURES SHOW CURRENT DRAM-TO-DRAM BONDING IN HBM

- ▲ Bonding is the process of attaching dies to one another or to a substrate to provide electrical and physical connectivity between dies



DRAM-TO-DRAM BONDING



- ▲ Existing microbump bonding
 - Large microbump sizes degrade TSV density
 - Thick die and underfill increase the stack height and degrade thermal conductivity
- ▲ Potential alternative bonding
 - Hybrid bonding: Simultaneous metal-to-metal bonding and oxide-to-oxide bonding
 - Direct oxide bonding: Oxide-to-oxide bonding using a low-temperature process

ALTERNATIVE DRAM-TO-DRAM BONDING



- ▲ Alternative bonding can potentially
 - Improve TSV density and speed
 - Improve thermal conductivity
 - Reduce stack height
- ▲ High-volume manufacturing feasibility and cost of alternative bonding needs evaluation
 - As well as reliability, yield, power delivery
 - Research opportunity, thus academia and industry can step in
- ▲ Alternative bonding can potentially enable taller, higher bandwidth stacks

Mass production of DRAM stacks using alternative bondings are unlikely in the next few years

Academia in partnership with industry can help with early research phases

CONCLUSIONS



- ▲ Presented challenges and potential directions to enhance in-package memory capacity, bandwidth, latency, reliability, and cost
- ▲ Main obstacle is stacking a high number of DRAM dies to provide the required capacity while achieving high memory bandwidth
 - As well as high thermal conductivity and high-volume, high-yield production process
- ▲ Alternative bonding techniques can potentially break the obstacle
 - More capacity, more bandwidth, better energy efficiency
- ▲ More research needed
 - Intriguing challenges and thus research opportunities
 - We set the stage and presented challenges and some potential directions
 - We recommend researchers from different domains such as packaging, reliability, and design architecture to participate

THANK YOU

DISCLAIMER & ATTRIBUTION



The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION

© 2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. SPEC is a registered trademark of the Standard Performance Evaluation Corporation (SPEC). Other names are for informational purposes only and may be trademarks of their respective owners.