

A Novel 3D DRAM Memory Cube Architecture for Space Applications

Anthony Agnesina¹, Amanvir Sidana¹, James Yamaguchi², Christian Krutzik², John Carson², Jean Yang-Scharlotta³, and Sung Kyu Lim¹

¹School of ECE, Georgia Institute of Technology, Atlanta, GA

²Irvine Sensors Corporation, Costa Mesa, CA

³NASA Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA
agnesina@gatech.edu

ABSTRACT

The first mainstream products in 3D IC design are memory devices where multiple memory tiers are horizontally integrated to offer manifold improvements compared with their 2D counterparts. Unfortunately, none of these existing 3D memory cubes are ready for harsh space environments. This paper presents a new memory cube architecture for space, based on vertical integration of Commercial-Off-The-Shelf (COTS), 3D stacked, DRAM memory devices with a custom Radiation-Hardened-By-Design (RHBD) controller offering high memory capacity, robust reliability and low latency. Validation and evaluation of the ASIC controller will be conducted prior to tape-out on a custom FPGA-based emulator platform integrating the 3D-stack.

1 INTRODUCTION

The on-board computing capabilities of spacecraft are a major limiting factor for accomplishing many classes of future missions. Effective execution of data-intensive operations such as terrain relative navigation, hazard detection and avoidance, and autonomous planning and scheduling require high-bandwidth and low-latency memory systems to maximize processor usage. Furthermore, the memory system must be capable of providing the necessary operational robustness and fault tolerance required for space applications.

Recently, manufacturers such as Micron, Hynix, Samsung, AMD, and Intel have exploited the 3D stacking technology to create the next generation of 3D stacked memories. This next generation of 3D stacked memory devices employs the use of through-silicon vias (TSVs) for interconnection between dies. This leads to increased bandwidth, lower latency, and lower power consumption, thereby mitigating the “memory wall”.

Unfortunately, these devices are just coming to the commercial market and none of these are ready for space applications. In particular, the behavior and resilience of TSVs in space environment are under investigation. There is also presently a limited memory capacity available for these type of devices (< 8 GB), as well as a limited number of memory dies that can be stacked together (< 8 dies). Manufacturers such as Samsung have stacked NAND flash packages with up to 16 layers using high aspect ratio staircase

wire bond structures. However, these stitch bond interconnection networks will not satisfy the electrical requirements for high-speed DDR operation nor will they offer the selectability required for radiation hardness.

The desire to utilize COTS devices for space has been an ongoing drive by the space community. Extensive testing on COTS devices has been undertaken over recent years to understand the performance of these devices in a space environment. Using COTS devices not only increases the component selection choices but also drives down the procurement time and system cost.

Hence, our work focuses on developing a space-qualified 3D memory cube (3D-M³) utilizing COTS DRAM memory devices supplemented with a custom Radiation-Hardened-By-Design (RHBD) controller. The contributions of this paper are as follows:

- We introduce a new stacking technology for our 3D-M³, which offers a flexible way to achieve high density packages and allows individual die access.
- We design a custom logic controller for the stack, offering a low-latency and reliable interface to the cube as well as advanced error mitigation features.
- We design a FPGA-based emulator platform to validate the combined operation of our 3D-M³ and RTL controller.

2 OVERVIEW OF OUR MEMORY CUBE

Selecting the most radiation tolerant, high performance COTS memory available, stacked using the proposed technology and integrated with the controller chip, will render the 3D architecture resilient for space environment while retaining state of the art performance in terms of bandwidth and density.

2.1 Overall Structure

Figure 1 shows an illustration of the proposed architecture for the complete 3D-Memory Package. It is composed of the following components:

- A single cavity ceramic package surrounding the 3D-M³ flip bonded to the logic tier.
- A 3D-M³ made of 14 memory dies arranged vertically in a “Loaf-of-Bread” (LOB) cube configuration, a structure whose benefits are presented in the following section.
- The logic controller placed under the cube which includes a Serializer/Deserializer (SerDes) interface. The complete architecture of the controller is presented in Section 4.

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor, or affiliate of the United States government. As such, the United States government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for government purposes only.

DAC '18, June 24–29, 2018, San Francisco, CA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5700-5/18/06...\$15.00

<https://doi.org/10.1145/3195970.3195978>

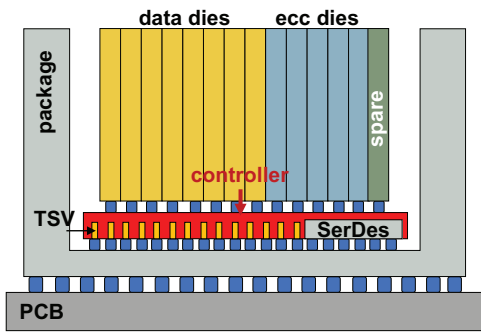


Figure 1: Our 3D-Memory Package

- A printed circuit board (PCB) interconnect to mount multiple such memory solutions or our memory system along with a host processor to allow Near Memory applications.

2.2 Interfacing

The interfacing between the 3D-M³ and the controller is a DDR interfacing. The 3D-M³ is Micro flip-chip Ball Grid Array (ball size of 100 μm) attached to the controller due to the high chip I/O density (~ 1288 pins) and small pitch required for keeping the 14 die DDR interface to a reasonable size. The flip-chip process is also used to maintain as much height to the solder pillars to help with stresses induced at the interface due to Coefficient of Thermal Expansion mismatches. The bottom interfacing of the controller to the package is a typical BGA (ball size of 250 μm) with smaller I/O density (~ 361 pins). Interconnection of the active circuitry from the top of the controller to the bottom I/O pads is made through TSVs.

A Re-Distribution Layer (RDL) applied to the active layers allows the interconnection of the die I/Os to one edge of the layer. The RDL brings out traces to the cube face at the same pitch of the I/O pads on the die (75 μm). Therefore, it is necessary to create a staggered I/O pattern on the face of the cube to expand the pitch at 200 μm to allow for 100 μm diameter tin-lead BGA balls to attach the 3D-M³. Package attachment technique to the final system board is Ceramic Column Grid Array. To hermetically seal the package and for radiation shielding purposes, a one-piece lid constructed from high-Z material will be welded to the top of the package.

2.3 Comparison

Peak bandwidth, density and power consumption of our memory solution are estimated using an 8 Gb Micron DDR3-1866L x16 die of maximum speed grade. Table 1 compares our solution with the Hybrid Memory Cube (HMC) [4]. We estimate the power consumption of the 3D-M³ to be about 5W at full speed as one die consumes ~ 380 mW [2]. Adding the power consumption of the logic tier with SerDes, the total power consumption is expected to remain under 8W.

Compared with other solutions, ours has notable advantages:

- A higher Density \times BandWidth/Power ratio,
- As each die can be accessed individually, intermixing memory types (DDR, Flash, MRAM...) into a single cube is possible using heterogeneous stacking techniques,

Table 1: Comparison of our solution with the HMC

Metrics	3D-M ³	HMC
Package Size (mm)	$\sim 25 \times 25 \times 10$	$31 \times 31 \times 4.2$
Density	8 GB	2 GB
Peak Bandwidth	30 GB/s	151 GB/s
Power (worst case)	~ 8 W	~ 20 W
Density \times BW/Power	30	15
Ball Count	361	896

- Achieves high memory densities without the need for TSVs.

3 CUBE DESIGN AND FABRICATION

3.1 Die Selection

DDR3 SDRAM technology is chosen over DDR4 despite being slower and more power hungry mainly because TID and SEE/Heavy Ions-Protons tests on DDR3 are well documented, while radiation testing on DDR4 devices is lacking. Moreover, the dies have to be readily obtained in a wafer format as we need to perform the RDL process at wafer level.

3.2 Stacking Method

The HMC utilizes TSVs and a “pancake” style stacking of dies, which does not require traversal along the cube face and lends itself to a lower cube height profile, but limits the number of dies that can be stacked together, limiting the density. For direct access to each individual die within the cube, achieving wide word access and improved error robustness, the 3D-M³ is built in a LOB style versus a pancake style, where all connections come directly out the edge. The LOB stack optimizes many design aspects, including:

- absolute minimum electrical impedance due to the short and point-to-point logic-to-memory interconnects,
- low thermal impedance,
- I/O connectivity able to support individual die access for recover and rebuild process,
- power input pad design for power switching of individual die,
- wide word access for increased bandwidth,
- avoids the use of TSVs for interconnection between dies.

In particular, the face-to-face connection between the cube pads and logic pads has the effect to reduce the interconnection length, which minimizes capacitive and inductive coupling. Therefore, switching currents and reflections are reduced, and so is power consumption due to a lower capacitive loading (estimated <3 pF) and the possibility to disable ODT (savings estimated to 1.3W for the cube [2]) without suffering significant signal integrity issues.

3.3 Cube Manufacturing

To process the memory dies for stacking, we apply on each die a RDL in order to bring the necessary interconnections from the die I/Os to an edge of the die. A multi-layer RDL made of two metal layers has been completed and the layer diced to size (see Figure 2). Using multiple layers of RDL allows route-ability of the interconnections as well as improves the distribution and impedance of power and ground connections.




Figure 2: Completed 2-Metal RDL Layer




Figure 4: 3D Cube Face with Bus Metal




Figure 3: Expanded View of 3D cube

After completing the RDL on the wafers, they are grounded to the desired thickness using a diamond grinder. Subsequently, the dies are diced to their final lamination size and then laminated to form the cube structure using an adhesive. We also integrate silicon bypass capacitor layers adjacent to the DDR dies to provide high-frequency performance for the power delivery network (PDN). Silicon fillers are added on the ends of the cube to protect the extreme active die and give more space to apply the bus metal. The face of the cube where the lead extensions go to the edge are processed with an isolation dielectric to allow for formation of interconnect pads. As the attachment of the memory cube to the logic tier is flip-chip, we apply an Under Bump Metal (UBM) to the bond pads. The bus metal is applied to the active face of the cube to form the bond pads for cube assembly. The final view of the cube along with its dimensions are shown in Figure 3, and the detailed final pattern of the cube face is shown in Figure 4.

3.4 Simulation Results

To ensure that the module operation is thermally acceptable, we perform thermal analysis for a stacked configuration of dies using SolidWorks Simulation Pro software. Results for fully loaded operation of the memory stack show less than 25°C rise in temperature. This is due to the LOB cube configuration which has the advantage of providing a direct thermal path through the silicon

to remove heat. As typical space missions require an electronic component operating range of 0 to 60°C, the simulated thermal performance provides an acceptable operating temperature level with 10°C margin for the thermal management solution.

We perform a preliminary electrical simulation of the anticipated DDR interface to validate performance with ODT disabled. Using FastHenry and FastCap field solvers, we extract the parasitic values of the DDR route to construct a simulation deck simulated using Hyperlynx LineSim. Performing a simulation using a DDR output with 35 Ω driver and no termination resistor at the input (i.e. ODT disabled), a 16-bit PRBS input yields a clean and open eye-diagram.

4 LOGIC CONTROLLER IMPLEMENTATION

The major roles of the controller are to communicate with the 14 DDR dies stacked above it and support basic DDR functions such as read, write and refresh. To address COTS radiation weaknesses we include error correction, scrubbing, device data rebuilding and die management features. Additionally, a high-speed serial interface provides enhanced attributes, such as reduced pin count and simplified connectivity. A preliminary RTL code of the controller has been completed and fully verified on sign-off ASIC and a FPGA platform is currently under design phase. The architecture of the logic memory controller is presented in Figure 5 and contains the modules described in the following sections.

4.1 DDR PHY Interface

As the main interface between the logic tier and the 3D-M³, the PHY provides a high-speed electrical interface to the DRAM. Due to the LOB cube configuration, we avoid the T-branch or Fly-By topologies found in typical DDR modules, where respectively loading is high or routing delays are introduced by traces. Here, driving point-to-point nets individually from the PHY, we can control the loading and cross-talk interference as well as provide tight delay matching.

4.2 Bank Spiraling

As the stacked devices are in close proximity, it is possible for radiation to strike the stack such that all layers get affected within the same bank. To minimize potential failures in similar memory cell areas of the dies, the controller remaps addresses so that data is stored in a “spiral” fashion when viewed through the stack, as shown in Figure 6.

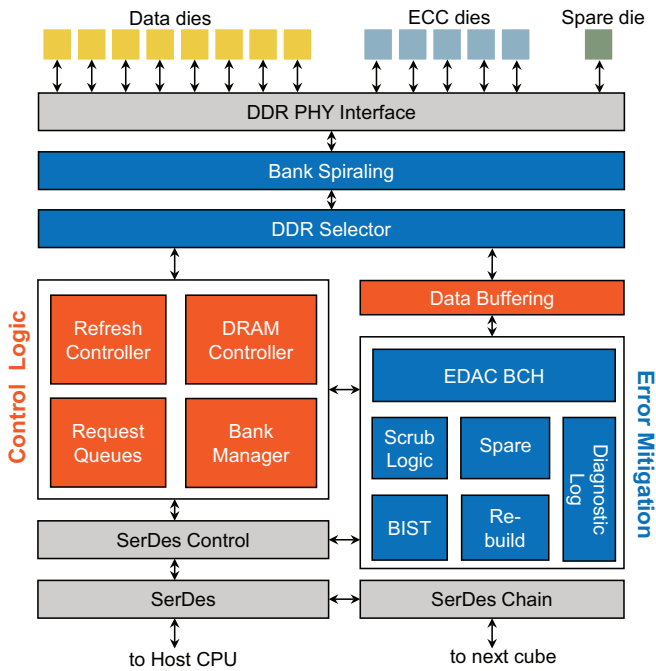


Figure 5: Architecture Diagram of the Logic Controller

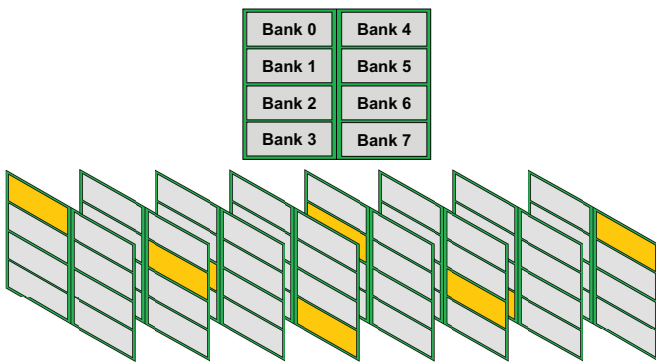


Figure 6: Typical DDR Bank Layout and Spiral Addressing within exemplar 8-layer Stack

4.3 DDR Selector

The DDR selector is a multiplexor that allows individual die to be either reset or power-cycled at run-time to allow for SEFI mitigation and stuck-bit buildup.

4.4 Data Buffering

The datapath includes a Data Buffer stage which holds the data read or written to the DRAM, until the request is completely serviced.

4.5 Control Logic

Our Control Logic offers the following low latency features:

- The general DRAM Controller is made of a Finite State Machine which translates memory requests to DRAM commands while

respecting the DRAM timings. We implement a simple yet efficient First-Ready First-Come First-Serve scheduling [5] to schedule the requests issued to the controller with good Quality of Service as well as maximum throughput. The controller can follow an Open-page Policy which achieves full efficiency in case of repeated reads and writes to the same page.

- The Request Queues are per-bank queues which hold the requests and their status until they are picked by the scheduler.
- A refresh controller performs a variable rate Cas-Before-Ras refresh whose rate can be changed in order to compensate the alteration of DRAM cells' retention time due to ambient temperature and the aging of the dies, or even to counter the "rowhammer" effect caused by aggressive row activations.
- A Bank Manager monitors the status of the banks so that unnecessary closing/opening are avoided.

The following low power features are offered in our Control Logic unit:

- The controller can switch to a Close-Page Policy to reduce power consumption in case of random access applications.
- The controller can put the idle memory in Power-Down mode until a request occurs.

4.6 Error Mitigation

As presented in many DDR3 device radiation studies [1], DDR3 devices are susceptible to the following radiation effects:

- (1) Single Event Upsets (SEU),
- (2) SEU buildup over time,
- (3) Stuck bits, and
- (4) Single Event Functional Interrupts (SEFI).

Fortunately, many tested DDR3 devices do not exhibit Single-Event Latchup (SEL) effects up to $61\text{MeV}\cdot\text{cm}^2/\text{mg}$ [1]. In order to achieve reliable operation of DDR3 in space environment, the controller provides a means of mitigation for each of the error modes. The following features are implemented in our Error Mitigation unit:

- The EDAC module provides necessary error detection and correction to mitigate SEUs as well as failed device logic. Our advanced ECC scheme uses a block-based Bose-Chaudhuri-Hocquenghem (BCH) encoding which provides a Double Error Correction-Triple Error Detection capability [3] and requires a total of 80 parity bits for our 128-bit (8 dies x16) data word. The parity bits are stored across 5 additional DDR dies.
- The Diagnostic Log keeps track of the logs for the Scrub Logic, BIST and Rebuild Logic modules, such as the addresses and characteristics of the errors found by the EDAC.
- The Scrub Logic decides whether to scrub a row, bank, or the entire die. To prevent SEU buildup or stuck bits, the logic tier performs a continuous scrub of the entire memory space at a programmable rate.
- The Rebuild module consists of a single State Machine controller that can be muxed into any of the DDR device channels, and performs a read/write cycle for the entire array.
- The Built-In Self Test (BIST) logic initializes and "zeroizes" each DDR device with a varying pattern, performs basic health testing and handles run-time evaluation of potential fault locations.

Table 2: Design Metrics of the ASIC logic controller

Parameter	Results	Parameter	Results
Footprint (mm)	10×10	Target period	2.426 ns
# Cells	79,698	Longest Path Delay	1.213 ns
Cell area	966,584 μm^2	Internal power	246.2 mW
# Nets	79,879	Switching power	342.6 mW
Wirelength	7,982 mm	Leakage power	2.004 μW
# I/Os to DDR3 Cube	1,288	Total power	588.8 mW
# TSVs to package	361		

- A Sparing module is included to automatically decide if, according to the error reports, we need to swap a deficient die with an entire spare cold die. Data is steered from the deficient die to the unused die.

4.7 Chaining & Serial I/O

The Serial I/O (SIO) interface allows multiple stacks to be chained together for scalability and higher memory capacity. This also greatly reduces the PCB layout requirements as well as simplifies the connectivity and reduce the SIO lanes required by the host processor.

To expand the operational flexibility, each stack can be configured to be operated independently or automatically in an n-modular redundancy configuration. As all DDR reads are performed in parallel, the only additional latency is from the voting logic and SIO data transfer. Another interesting application of the device chaining is to allow for automated stack failure rebuild by allowing for a stack to be configured as an XOR parity unit.

We plan to use Serial RapidIO as interconnection protocol between our cubes, in order to assure a SpaceVPX Fault-Tolerance and allow highly scalable and peer-to-peer possibilities.

4.8 Verification

Extensive Verilog simulations (see Figure 7) were carried out on the post implemented ASIC/FPGA designed RTL to test and verify the functionality of the controller in handling different memory operations and also the correction of errors in the data (the test benches include mechanisms to properly test Single-Event Effects and SEFI type failures). The EDAC’s capability to retrieve correct data upon a full DDR failure was simulated and verified for the 14-stack cube configuration. In addition, cold-spares swapping where the controller replaces a bad die with corrected data loaded into a spare die, was tested. This entire set of analysis was done with using 8 Gb DDR3 Verilog models from Micron.

4.9 ASIC Design

To begin estimating gate counts and performance concerns, we implement a detailed GDSII layout for a preliminary controller RTL, excluding SerDes and PHY. We select a foundry-grade 130 nm technology as it is the closest technological representative to the RHBD process among available resources. The controller RTL is synthesized and placed-and-routed to obtain the final sign-off layout. The design consists of controller logic cells and two PDNs to supply the power to logic cells and to the 3D-M³. A summary of the preliminary design is shown in Table 2.

The routing schemes of the design consist of three layers. Signal TSVs providing I/Os to the package routed to logic cells, control and data signals of the controller routed to DDR3 cube interface on the top, and logic cells routed among themselves.

We perform a Static Timing Analysis to verify the timing and power requirements. The controller can operate up to 800 MHz. However, I/Os of the controller must be driven through long wires from/to scattered signal TSVs and signal pins of DDR3 cube interface. Due to scattered external I/O pins and TSVs, the timing results are sensitive to their locations. Optimal configurations of signal and power pins are to be investigated to improve performance. The power consumption measured at 588.8 mW will likely increase with the addition of post-failure mechanisms and SerDes plus PHY.

For power drop analysis, planned rows of power/ground connections are used for a dedicated DDR3 PDN. A total of 4 metal layers are used in the logic tier, dictating the DDR3 pads to Metal 4. Initial PDN simulations showed that less than 5% IR-drop could be achieved for the power path from BGA pad, to TSVs, through 4 metal layers route, and finally to the DDR device pad (see Figure 8). Also, using IROC SoCFIT, we evaluate the Soft Error Rate of the design to 1.54 FIT/device. These results are positive indicators for the design concept.

4.10 FPGA Prototyping

Validation of the RTL code through the use of an FPGA provides a complete prototype design ready for migration to a final ASIC that can be integrated with the 3D-M³. However, it is anticipated that there will be some front-end design to migrate FPGA-based IP to ASIC IP such as the SerDes and PHY modules. What’s more, required for the FPGA implementation but not for the ASIC, are complex PHY training procedures to optimize the sampling delays and compensate for timing errors introduced by the FPGA and physical interconnections. In the final application, cube interconnect electrical concerns are virtually eliminated due to the extremely short interconnect.

We use the Xilinx DDR PHY IP Core [6] as DDR interface to communicate with the 3D-M³. As the IP has a limitation of a maximum load of 9 memory components, our solution to address 14 memory dies incorporates multiple PHYs. Thorough manual placement of the PHYs and controller has been completed due to tight timing issues that occur in such a timing constrained design. Synthesis and Implementation of the designed controller has been completed on a Xilinx Virtex Ultrascale XCVU440 FPGA with a maximum achievable frequency of 200 MHz and resource utilization of about 40k look-up tables, 50k flip-flops and 200 Block RAM36k, where PHYs are responsible for more than 70% of the utilization.

4.11 Design of the Test Board

The test board consists of the 3D-M³ mounted on an interposer substrate connected to the FPGA. For connection of the substrate to the cube, we use high density Z-ray spring pins as shown in Figure 9. The spring pin provides sufficient electrical performance and has the benefit of easy swapping of parts. Using such an interposer allows substrate connection to the load board on the peripheries of the substrate. A support frame integrated into the substrate (i.e. bonded at PCB level) helps reducing stress at the stack interface

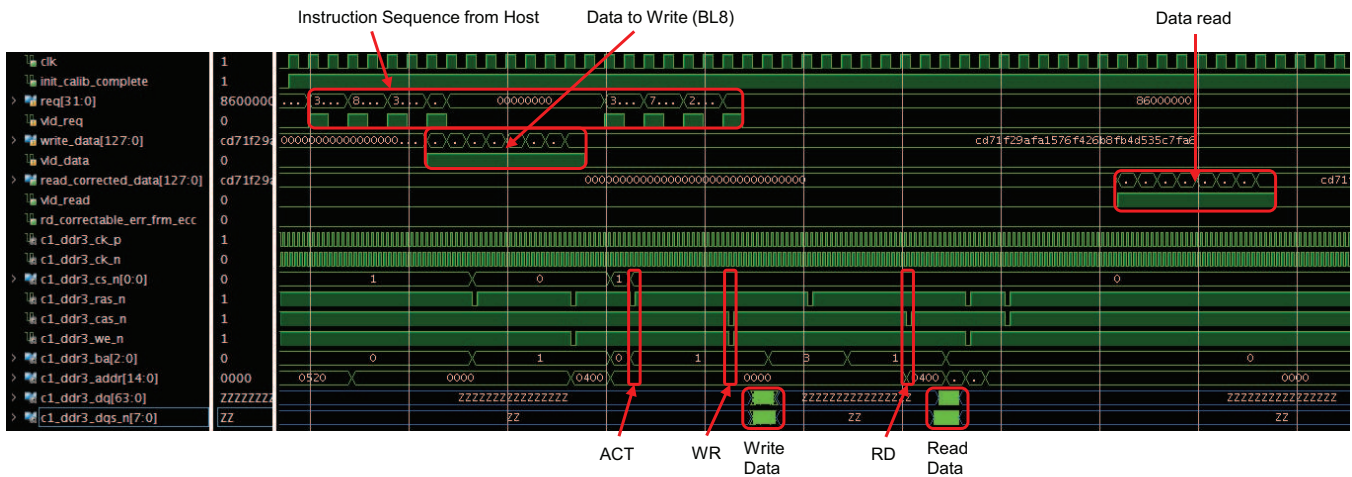


Figure 7: Simulation of the controller

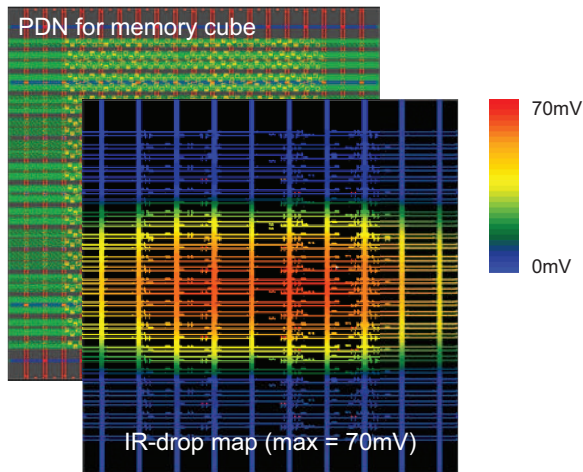


Figure 8: Full-chip PDN and IR-drop map for memory cube

and provides the necessary surface area for mounting hardware and alignments. This approach also allows the compressive forces required by the interposer to be absorbed away from the stack BGA interface. Using an interposer in the 4 quadrants provides a balanced structure and can handle the required I/O count.

5 CONCLUSION

The most important aspect of this research is to provide an advanced technology enabling new space missions not possible with the current bulky, low bandwidth, low density and power hungry memory devices. The integration of a COTS, 3D DRAM stacked device with a RHBD logic controller supplemented by a high-speed serial interface, provides a versatile, fault-tolerant and scalable architecture. Our vertical stacking method allows individual die access and can provide high-density intermixing of different memory types. The proposed design concept and design path provides the flexibility to interchange DDR3 die for the 3D-M³ portion of the module. This




Figure 9: Stack mounting using spring interposer

makes it possible to revisit the die selection in parallel with the controller design in case future radiation testing or memory device manufacturer's roadmap alterations require a component change.

ACKNOWLEDGMENTS

This research is funded by the NASA SBIR Grant under the contract number NNX17CP02C. A portion of the research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] M. Herrmann, K. Grumann, and F. Gliem. Heavy Ion SEE Test of 4 Gbit DDR3 SDRAM Devices Test report, 2014.
- [2] Micron. TN-41-01: Calculating Memory System Power For DDR3, 2007.
- [3] R. Naseer and J. Draper. DEC ECC design to improve memory reliability in Sub-100nm technologies. In *2008 15th IEEE International Conference on Electronics, Circuits and Systems*, pages 586–589, Aug 2008.
- [4] J. T. Pawlowski. Hybrid Memory Cube (HMC). In *2011 IEEE Hot Chips 23 Symposium (HCS)*, pages 1–24, Aug 2011.
- [5] S. Rixner et al. Memory Access Scheduling. *SIGARCH Comput. Archit. News*, 28(2):128–138, May 2000.
- [6] Xilinx. UltraScale Architecture-Based FPGAs Memory IP v1.4.