

A COTS-Based Novel 3-D DRAM Memory Cube Architecture for Space Applications

Anthony Agnesina¹, *Graduate Student Member, IEEE*, James Yamaguchi, Christian Krutzik, John Carson, Jean Yang-Scharlotta², *Member, IEEE*, and Sung Kyu Lim, *Senior Member, IEEE*

Abstract—The first mainstream products in three-dimensional integrated circuit (3-D IC) design are memory devices where multiple memory tiers are horizontally integrated to offer manifold improvements when compared with their 2-D counterparts. Unfortunately, none of these existing 3-D memory cubes are ready for harsh space environments. This article introduces a new memory cube architecture for space, based on the vertical integration of multiple commercial-off-the-shelf, 3-D stacked, dynamic random-access memory (DRAM) memory devices with a custom radiation-hardened-by-design controller. Our solution offers high memory capacity, increased bandwidth, fault tolerance, and improved size-weight-and-power characteristics needed for space missions. Validation and functional evaluation of the application-specific integrated circuit (ASIC) controller will be conducted prior to tape-out on a custom FPGA-based emulator platform integrating the 3-D stack. The selected test methodology ensures high-quality register transfer level (RTL) as well as allows to subject the cube structure to radiation testing. The proposed design concept allows for flexibility in the choice of the DRAM die in the case of technology road-map changes or unsatisfactory radiation results.

Index Terms—Aerospace electronics, dynamic random-access memory (DRAM), memory management, radiation effects, three-dimensional integrated circuit (3-D IC).

I. INTRODUCTION

THE on-board computing capabilities of spacecraft are a major limiting factor for accomplishing many classes of future missions. In particular, the deep space exploration program requires effective execution of data-intensive operations, such as terrain relative navigation, hazard detection and avoidance, and autonomous planning and scheduling. These lengthy and wearing tasks require high-bandwidth and low-latency memory systems to maximize processor usage and provide rapid access to observational data captured by high-data-rate instruments (e.g., hyperspectral infrared imager and interferometric synthetic aperture radar), as well as

necessitate electronic devices that can exhibit extreme operational robustness, efficiency, fault tolerance, and energy management aspects. Although technology development efforts are underway into developing a high-performance space-computing/next-generation space processor (HPSC/NGSP), a radiation hardened multicore computing processor expected to fuel missions until 2030, they do not address the limitations of current on-board memory systems.

Recently, manufacturers such as Micron, Hynix, Samsung, AMD, Sunnyvale, CA, USA, and Intel have exploited the 3-D stacking technology to create the next generation of 3-D stacked memories, employing through-silicon vias (TSVs) for interconnection between dies. This leads to increased bandwidth, lower latency, and power consumption, thereby mitigating the “memory wall.” Unfortunately, these devices are not ready for the space applications. There is also presently a limited memory capacity available for these type of devices (<8 GB) and 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 double data rate (DDR) operation nor will they offer the selectability required for radiation hardness.

Progressive testing of commercial-off-the-shelf (COTS) devices for space applications has opened up new windows of opportunity for increased functionality of space electronics. Not only does this increase the component selection choices but it also drives down the procurement time and system cost.

Our proposed approach is to leverage COTS memory devices to maximize memory density, bandwidth, and speed, by integrating them using our new stacking technology into a 3-D memory cube (3-D-M³) supplemented with a controller chip. Features integrated in the controller chip to address COTS deficiencies in terms of radiation tolerance include error correction and detection (EDAC), scrubbing, device data rebuilding, and die-level reboot and swap. Our cube would prove to be a complementary memory system for HPSC and other space-computing systems targeting high performance. Our complete memory system can address the computational performance, energy management, and fault tolerance needs of space missions. It can potentially be utilized in many aspects of the spacecraft beyond just the main compute/processor, such as in instruments, detectors, sensors, and communication systems.

This article is viewed as a comprehensive summary and a substantial extension of the work presented in [1] and [2].

II. CUBE DESIGN CONSIDERATIONS

After a long wait, three-dimensional integrated circuits (3-D ICs) are becoming a mainstream technology, with applications

Manuscript received January 9, 2020; revised April 21, 2020 and May 13, 2020; accepted May 22, 2020. Date of publication July 20, 2020; date of current version August 26, 2020. This work was supported in part by the National Aeronautics and Space Administration (NASA) Small Business Innovation Research (SBIR) Grant under Contract NNX17CP02C and in part by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. (Corresponding author: Anthony Agnesina.)

Anthony Agnesina and Sung Kyu Lim are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: agnesina@gatech.edu).

James Yamaguchi, Christian Krutzik, and John Carson are with Irvine Sensors Corporation, Costa Mesa, CA 92626 USA.

Jean Yang-Scharlotta is with the NASA Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA.

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVLSI.2020.3002211

1063-8210 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

into consumer memory such as the high bandwidth memory (HBM) [3] and hybrid memory cube (HMC) [4], where multiple memory dies are vertically integrated and connected through thousands of TSVs, to offer “more than Moore” size, capacity, speed, and power consumption improvements. These memory parts could truly enable the aforementioned space missions but unfortunately none are ready for space. In particular, it is unclear how TSVs behave in cryogenic temperatures (50–100 K) and high radiations (>5 Sv/day) environments, such as Mars and icy moon surfaces [5].

A. Leveraging TSV-Based 3-D IC?

1) *Radiation Impact on TSV Reliability*: Little work exists on 3-D IC reliability analysis for space applications. Gouker *et al.* [6] studied radiation effects on 3-D IC for the first time and characterized single event transients (SETs) for a small-scale logic circuit implemented into a 3-tier 3-D IC using a 180-nm silicon on insulator (SOI) process. Each tier contains a simple circuitry that measures SET pulsewidth and is connected with 1.25- μm TSVs. Their heavy-ion testing based on krypton shows that the pulsewidth distribution across the tiers exhibits nontrivial tier-to-tier variations. They observed that various materials used in the 3-D IC lead to different levels of energy deposited during the radiation testing. For instance, the effect of nuclear recoils is strongest in tungsten, which is typically used for contacts and vias. This issue will exacerbate if tungsten TSVs are used instead of copper, which is more popular due to its advantage in coefficient of thermal expansion (CTE)-related mechanical stress problems. The authors also conducted a similar study using proton and neutron beams on an static random access memory (SRAM) module implemented in the same 3-D IC technology. Unlike in the logic circuit, their SRAM module suffers less from tier-to-tier effects when the un-thinned bottom tier is rad-hardened with a special SOI wafer. In both studies, however, their focus is not on TSVs, and no rad-hard design method is presented. This is critical as TSVs, a large object in 3-D ICs, can become an easy target for particle attack and cause single event effects (SEEs) and total ionization dose (TID) issues to themselves and nearby objects. In addition, the designs used are nowhere close to commercial memory cubes. Moreover, existing work on actual commercial memory cubes as such is limited and mainly focused on processing-in-memory architecture. If some efforts have been done to reduce power and bandwidth overhead, none of these solutions address radiation hardening of memory cubes nor their operations under extreme temperature.

One solution to alleviate the lack of radiation hardening of these stacks is to incorporate an HBM or HMC-type cube with a custom radiation-hardened-by-design (RHBD) controller. That would give the advantage of utilizing state-of-the-art high bandwidth and low-power memory stacks, as well not needing to fabricate a full stack. Unfortunately, it is unclear as to the availability of the memory stacks alone without the logic tier, and the proprietary nature of these tiers is a showstopper for agencies like NASA. Moreover, if building a module with both the memory and controller being RHBD theoretically provides the highest robustness and performance, the downfall is that designing both is fairly complex and costly. Also, since the design would be relying on a specific memory, it would limit the ability of the module to interchange memory.

2) *Thermo-Mechanical Effect Mitigation for TSVs*: Using polymers for the TSV liner such as benzo-cyclo-butene (BCB)—instead of more rigid silicon dioxide—is a possible

solution for the mechanical stress issue. However, BCB liners are harder and more expensive to make. Another popular solution to alleviate TSV mechanical stress is enlarging the TSV keep-out-zone. This reduces the impact of thermo-mechanical issues on nearby devices, interconnects, and neighboring TSVs, but its area overhead can be significant. The design solution is to evenly space out TSVs in the full-chip layout and avoid creating spots with high TSV concentration. This, however, needs to be done carefully because confining TSV location to a certain location may compromise the overall performance of the design.

3) *Low Temperature Impact on TSVs*: One of the major issues of 3-D ICs on earth is thermal. Power density increases in 3-D ICs due to the closer proximity of devices in the structure, which poses a major challenge in heat dissipation. Mars or icy moon surface, on the other hand, pose an entirely new and opposite challenge: heat is not a major issue, but TSVs will shrink faster than silicon substrate due to a negative thermal load and CTE mismatches.

However, the cryogenic temperature provides an excellent operation environment for memory cubes and other electronics when compared with room or hot temperature. Dynamic random-access memory (DRAM) cells retain data much longer and thus require much less frequent refresh. The mobility of charge carriers improve in MOSFETs, so they switch faster and achieve higher gain. Interconnect resistance reduces with lower temperature, so communication becomes faster. Device reliability improves due to a sharper ON/OFF ratio, and interconnects experience orders of magnitude lower electromigration. Moreover, studies show that radiation damage reduces under cryogenic temperature [7].

B. COTS Solution

The mentioned issues rationalize a 3-D-M³ of COTS memory with a custom RHBD controller. COTS devices have a lot of scrutiny and interest in the space community. Extensive testing on SDRAM COTS has been undertaken over the recent years and NASA successfully used them in spacecraft for critical applications such as for the Compute Element within the Mars Science Laboratory Curiosity rover. With proper component screening and radiation testing, the use of COTS devices within our cube provides a low-cost and effective solution to capture the benefits offered by state-of-the-art commercial devices.

III. OUR 3-D MEMORY CUBE SOLUTION

Selecting the most radiation-tolerant, high-performance COTS memory available, stacked using the proposed technology and integrated with the controller chip, will render the 3-D architecture resilient for space environment while retaining state-of-the-art performance in terms of bandwidth and density. More than theoretically achieving comparable performances to the HMC, our solution also offers increased density, possible intermixing of different memory types such as magnetoresistive random-access memory (MRAM), NAND Flash, etc. into a single cube, increased fault-tolerance aspects such as individual die access, as well as avoidance of TSVs use for interconnection between dies.

A. New Stacking Method

The HMC and HBM utilize 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

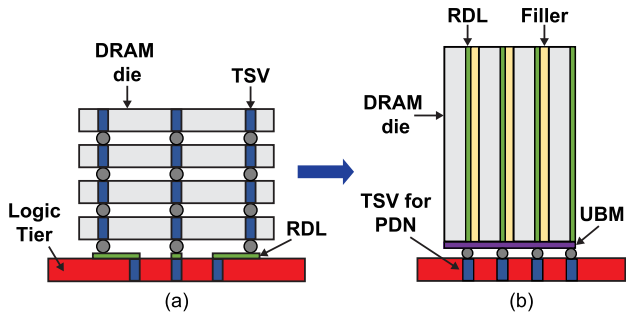


Fig. 1. Our proposed (a) “LOB” structure versus (b) traditional “pancake” style.




Fig. 2. Our 3-D-memory package.

the density. In our proposed “Loaf-of-Bread” (LOB) cube configuration detailed in Fig. 1, it is more readily achievable for the logic tier to interact with each individual die within the cube, when compared with vertical stacking. The LOB stack optimizes many design aspects, including the following:

- 1) absolute minimum electrical impedance due to the short and point-to-point logic-to-memory interconnects;
- 2) low thermal impedance;
- 3) I/O connectivity able to support individual die access for recover and rebuild process;
- 4) power input pad for power switching of individual die;
- 5) wide word access for increased bandwidth;
- 6) avoidance of TSVs for interconnection between dies.

The face-to-face connection between the cube pads and logic pads reduces the interconnection length, and minimizes capacitive and inductive coupling. Therefore, switching currents and reflections are reduced, and so is power consumption due to a lower capacitive loading (estimated <math><3\text{ pF}</math>) and the possibility to disable ODT (savings estimated to 1.3 W for the cube [8]) without suffering significant signal integrity issues.

B. Our 3-D-Memory Package

Our stack contains the logic controller placed under the 3-D-M³ with 14 DDR dies stacked in an LOB fashion, as shown in Fig. 2. Eight DDR \times 16 devices make up the data word, five extra devices store the ECC, and a spare die is included in case of die failure replacement. A DDR interface connects to the external host DDR interface. A single-cavity package surrounds the cube structure flip bumped to the logic tier. Interconnection of the active circuitry on the top of the controller to the bottom I/O pads is made through TSVs. These TSVs will be filled with copper to eliminate issues associated with tungsten filled vias, a high-Z material producing radiation effects that can be detrimental to the operation of the device as proton generation can affect the operation of the logic sections of the device [9]. A lid will be welded to the top of the package to completely enclose the module hermetically and for radiation shielding purposes. Despite subjecting the 3-D-M³ to

additional reflow processes during the assembly and requiring TSVs, a single-cavity package was chosen over dual-cavity because it allows testing the 3-D memory module prior to assembly into the package. Moreover, the dual cavity requires an extremely high I/O count to pass through the substrate which poses technical risks.

C. DDR3 Die Selection

To meet the present and future needs of space missions for high speed and high density memory, it was determined that DDR would be the best choice. DDR4 device versus a DDR3 device from the same manufacturer showed a 45% increase in single bit upset cross section and a 17% logic upset decrease [10]. This would tend to conclude that the DDR4 devices are more susceptible to bit errors than their DDR3 counterparts but have a higher resilience to logic upsets, which could prove beneficial given the very high level of ECC of the module. However, we selected the DDR3 technology as it is still the most mature of the DDR series and has had a fair amount of review for radiation tolerance whereas DDR4 has not had a lot of attention in this regards as of this date.

Reviews from [11] and [12] showed DDR3 devices from Micron, Samsung, or Hynix are strong candidates for implementation for space. No occurrences of destructive failures or single-event latchup were noted. A general observation is that DDR devices are more sensitive to single-event functional interrupts (SEFIs) than single-event upsets (SEUs). All devices showed good resistance to TID effect, from 100 krad (Si) up to 400 krad. In Micron devices, idle current increases dramatically as the TID approaches 100 krad, whereas Samsung and Hynix devices show a more gradual increase with increased TID exposure. A noticeable trend in tolerance to SEUs seemed to play out: during SEE proton testing, SEUs and column/row SEFIs showed an increase toward low linear energy transfer. Typically, Hynix devices showed best resistance to SEUs. In some instances, Samsung and Hynix devices showed orders of magnitude better cross sections.

D. Interfacing

The interfacing between 3-D-M³ and controller is a DDR interfacing. The 3-D-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 (\sim 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 CTE mismatches. The bottom interfacing of the controller to the package is a typical BGA (ball size of 250 μm) with smaller I/O density (mainly a \times 128 DDR interface and power/ground pads). Interconnection of the active circuitry from the top of the controller to the bottom I/O pads is made through TSVs.

A redistribution layer (RDL) applied to the active layers allows 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, we must 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 3-D-M³. Package attachment technique to the final system board is ceramic column grid array.

E. Comparison With HMC and HBM2

We estimate peak bandwidth, density, and worst-case power consumption of our memory cube using the characteristics of

TABLE I
COMPARISON OF OUR SOLUTION WITH THE HMC AND HBM2

Metrics	3D-M ³	HMC	HBM2
Package Size (mm)	~ 25x25x10	31x31x4.2	8x12x0.745
Density	8 GB	2 GB	8 GB
Peak Bandwidth	30 GB/s	151 GB/s	256 GB/s
Power	~ 8 W	~ 20 W	~ 30 W
Density×BW/Power	30	15	68

an 8-Gb Micron DDR3-1866L ×16 die of maximum speed grade. Table I compares our solution with the HMC and the latest HBM2, using the figures obtained in their datasheets. We estimate the power consumption of the 3-D-M³ to be about 5 W at full speed as one die consumes ~380 mW [8]. Adding the power consumption of the logic tier, the total power consumption is expected to remain under 8 W. Compared with other solutions, ours has notable advantages which are as follows.

- 1) A competitive Density × BandWidth/Power ratio.
- 2) As each die can be accessed individually, intermixing memory types (DDR, Flash, MRAM, etc.) into a single cube is possible using heterogeneous stacking techniques.
- 3) Achieves high memory densities without the need for TSVs.
- 4) If needed, much more taller stacks can be built to achieve higher bandwidth and density, without endangering the stability of the cube structure.

Our ambition is not to compete with consumer memory cubes, leveraging TSV-based, face-to-face or monolithic 3-D technologies which achieve higher power-performance-area thanks to fine-grained integration. Compared with existing memory systems used in space, our cube offers higher performance using only commercially available devices.

IV. CUBE FABRICATION AND VALIDATION

A. Manufacturing

The chips purchased as individual die are known good die burned in and pretested prior to stacking. We process each die by applying an RDL to bring the necessary interconnections from the die I/Os to an edge of the die. Traces are in gold with a thin layer of a titanium tungsten alloy as the barrier/adhesion metal and interdielectric coating material is a polyimide. These materials have low outgassing properties and high glass transition temperatures, therefore able to cure properly in very thin thicknesses which is critical in maintaining the runout of the layer reroute in relation to the bus metal pattern. A multilayer RDL made of two metal layers has been completed and the layer diced to size (see Fig. 3). The RDL process is a sequential process that is circuitry is built-up layer-by-layer. Using multiple layers of RDL allows route-ability of the interconnections as well as improves the distribution and impedance of power and ground connections.

After completing the RDL on the wafers, they are grounded to the desired thickness using a diamond grinder. The dies are then diced to their final lamination size and laminated to form the cube structure using an adhesive, as shown in Fig. 4. We integrate silicon bypass capacitor layers adjacent to the DDR dies for high-frequency performance of the power delivery network (PDN). Silicon fillers added on the ends of the cube protect the extreme active die and give more space




Fig. 3. Completed two-metal RDL.




Fig. 4. Layers laminated into cube structure.




Fig. 5. Expanded view of the 3-D cube.

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 by sputtering. We use a 3-metal UBM consisting of an adhesion metal followed by a solderable metal which is then capped with a thin layer of gold to prevent the solderable metal from oxidizing. The adhesion metal also acts as a barrier between the subsequent solderable metal and the aluminum pad on the device to prevent metal migration which can form undesirable intermetallics. 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 Fig. 5, and the detailed final pattern of the cube face is shown in Fig. 6.

B. Cube Design Validation

1) *Thermal Simulations:* To ensure 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 (13 active dies) show less than 25 °C rise in temperature. This is due to the LOB configuration which provides a direct thermal path through the silicon (a high thermal conductivity




Fig. 6. 3-D cube face with bus metal.




Fig. 7. Thermal comparison of (a) “pancake” versus (b) LOB.




Fig. 8. Comparison of signal path in (a) “pancake” and (b) LOB.

material) to remove heat. Thermal simulations (see Fig. 7) show that heat can be dissipated readily through the bottom and top of the package to maintain a safe operating temperature, whereas the heat in the “pancake”-style configuration cannot readily escape through the silicon. This is due to the fact that each die RDL creates a high thermal barrier between die, as shown in Fig. 8. The junction-to-case thermal resistance is approximately $3\text{ }^{\circ}\text{C}/\text{W}$ and junction-to-board about $5\text{ }^{\circ}\text{C}/\text{W}$. This shows a significant thermal path exists from junction-to-board. This provides the thermal management solution for utilizing the PCB path for additional 3-D-memory module cooling. As typical space missions require an electronic component operating range of $0\text{ }^{\circ}\text{C}$ – $60\text{ }^{\circ}\text{C}$, the simulated thermal performance provides an acceptable operating temperature level with $10\text{ }^{\circ}\text{C}$ margin for the thermal management solution.

2) *Electrical Simulation*: We perform an electrical simulation of the 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. The IBIS model is shown in Fig. 9. Performing a simulation using a DDR output with $35\text{-}\Omega$ driver and no termination resistor at the input (i.e., ODT disabled), a 16-bit PRBS input yields a clean and open eye diagram, as expected due to our architecture with short interconnect.

V. CONTROLLER DESIGN

The major role of the controller is 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. The top-level diagram of the controller is shown in Fig. 10. Our architecture is destined to be fully flexible and reusable, integrating an individual DDR memory controller to enable various bus protocols to communicate with the cube. For example, a few space agencies will communicate with the memory cube using the Serial RapidIO protocol.

Our architectural choices answer the need of a smooth integration of the controller with a host DDR interface, so that the memory can be actively used by the processor and perform necessary maintenance/error-mitigation operations on the 14 independent dies’ stack. The necessary features required for the controller and addressed in the register transfer level (RTL) code are in particular:

- 1) read/write capability to the stack;
- 2) wide 14 DDRs physical interface (PHY) to the cube;
- 3) DDR interface to host;
- 4) EDAC and housekeeping;
- 5) access to each individual die within the cube for individual power management, ON-OFF, or rebuilding;
- 6) SEFI error handling.

The specifications of the DDR interface between the host and controller have important repercussions on the architecture and latency of the memory controller.

A. Host Interface Considerations

1) *Latency Concerns*: The latency of the core processing is critical, as typical DDR3 devices have CAS latencies ranging from 7 to 16 clock cycles depending on the data rate. When the host memory controller issues a read to our memory module, the command is first issued to the cube, and the data are brought back from the cube to the host. Therefore, the latency of the path Host DDR PHY to MUX to Stack PHY to Cube CAS Latency to EDAC, and back to the Host has to be reasonable to be able to be handled by the host controller, and also for pure performance concerns. Minimizing the latency of each traversed block is then critical as the overall latency increases quickly with each block met on the path. Fortunately, the memory bandwidth is not affected by the additional processing.

2) *Flexibility Concerns*: A programmable bus width interface on the host side allows multiple data widths densities (baseline of $\times 128$) depending on the host specifications. It is noted that due to the buffering of the internal DDR3 die, the 3-D- M^3 presents a reduced load to the host processor, which in turn allows higher densities to be achieved that might not otherwise be possible due to loading and PCB layout issues.

3) *Integration Concerns and Idle Detector*: For the controller to be “transparent” to the host, function without interacting with it and without provision for extensive sideband control, an idle detector is included to determine when internal operations can be interleaved with regular host operations and interject scrubbing and other management functions. A discrete approach with an “Idle” gpio input signal is provided to allow the host system to notify the cube that the DDR3 bus will be idle—this is required if a rebuild is needed and the host desires to have it completed immediately. Likewise, an SPI port allows for additional housekeeping readout (query device operation, full error statistics, temperature, etc.), configuration of the cube as well as initiate additional tests.

The idle detector controls the multiplexer MUX to switch operation of our controller from normal mode




Fig. 9. IBIS model interface DDR-logic.

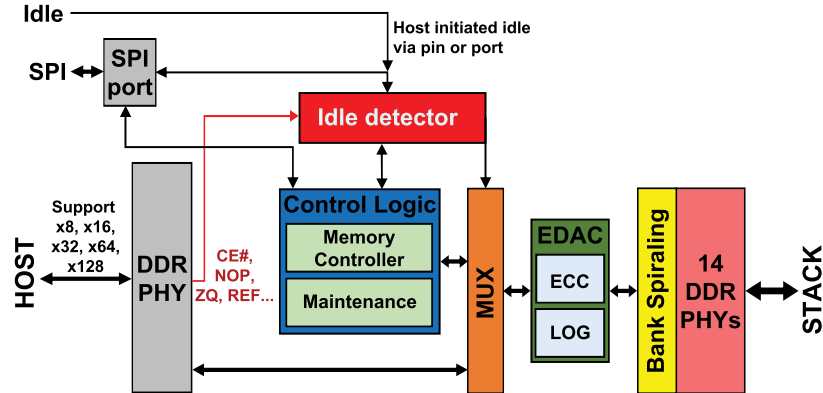


Fig. 10. Controller architecture.

(i.e., a slave-serving requests of the host in a pass-through-like fashion) to maintenance mode (the controller takes over control to perform power-up, initialization, calibration, scrubbing, rebuild, etc.). Due to the fixed timing of a DDR3 interface, the idle detector must insure that no host operations occur within the duration of the operation that is the future. This is a difficult task as host operations occur asynchronously from the standpoint of the command execution. Thus, the memory cube requires timing adjustments on the part of the host memory controller to allocate a few hundred extra nanoseconds on operations such as refresh (by adjusting t_{RFC} to a higher value), ZQ, CE enable, or NOP commands to allow the memory cube to perform its internal management functions.

B. Cube Interface Considerations—PHY Layer

As the main interface between the logic tier and the 3-D- M^3 , 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 crosstalk interference as well as provide tight delay matching.

Because each die must be independent, one PHY is implemented per die. To account for the fly-by routing of traditional dual in-line memory modules (DIMMs), DDR PHYs implement advanced timing procedures, adding several clock cycles to the CAS latency. As latency is a concern for our application, using them in their current form is not acceptable. Moreover, in the final application, the memory stack will be integrated on top of the application-specific integrated circuit (ASIC) controller, which eliminates any tight timing requirements due to the close proximity of connections.

Because our code will first be tested and validated on an FPGA, where complex training procedures are used to optimize the sampling delays and compensate for timing

errors introduced by the FPGA and physical interconnections, we relaxed the baseline performance requirements to allow working at a reduced frequency and therefore use a simplified lower latency PHY. As a back-end for coding the custom PHY, some parts of the Xilinx PHY MIG for Virtex-5 are utilized [13], to leverage an existing but simple FPGA PHY infrastructure that can be easily customizable to our needs.

C. Memory Controller

Our memory controller offers the following features.

- 1) A memory access scheduler with a two-level hierarchical organization as displayed in Fig. 11 selects a memory request to send to DRAM memory. A crossbar redirects the requests to the appropriate per-bank queue, which hold requests and their status until they are picked by the scheduler. The first level consists of separate per-bank schedulers keeping track of the state of each bank and selecting the highest priority request from its bank request buffer. The second level consists of an across-bank channel scheduler selecting the highest priority request among all the pending requests presented by the bank schedulers. When a request is scheduled by the memory access scheduler, its state is updated in the bank request buffer, and it is removed from the buffer when the request is served by the bank. In each bank scheduler, a finite state machine (FSM) translates memory requests to DRAM commands while respecting DRAM timings. The channel scheduler implements a simple yet efficient first-ready first-come first-serve (FR-FCFS) scheduling [14] to schedule the requests issued to the controller with good Quality of Service as well as maximum throughput.
- 2) A refresh controller performs a variable rate CAS before RAS refresh whose rate can be changed to compensate for the alteration of DRAM cells' retention time due to ambient temperature and aging of the dies, or even

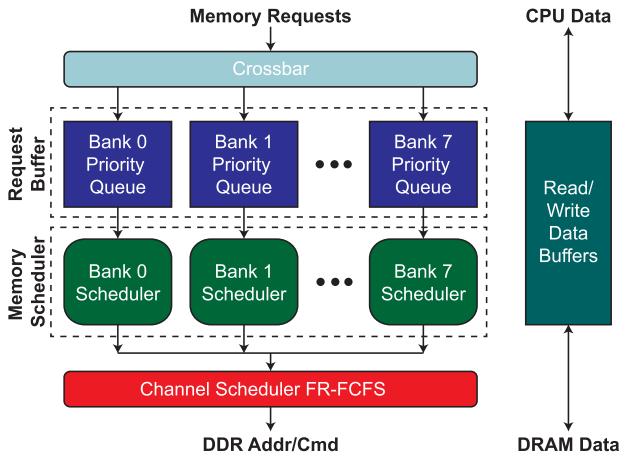


Fig. 11. Memory access scheduler for efficient FR-FCFS scheduling.

 TABLE II
FAILURE MODE MITIGATION

Failure Mode	Mitigation	Detection
SEU in array	EDAC	EDAC algorithm
SEU buildup	Scrubbing	Scrub over entire array
Leakage	Reduce refresh rate	Ground based testing of leakage, programmable refresh rate with T° sensor
SEFI: data failure	Power cycle & rebuild	EDAC errors exceed programmable threshold
SEFI: high current	Power cycle & rebuild	Current draw anomaly
Device failure	Activate cold spare	Device fails BIST after power cycle

counter the “rowhammer” effect caused by aggressive row activations.

- 3) The controller follows open-page policy achieving full efficiency in case of repeated reads/writes to the same page. It switches to close-page policy to reduce power consumption in case of random access applications.
- 4) The controller can put the idle memory in power-down mode until a request occurs.

D. Error Mitigation Features

A goal of the 3-D memory module is to provide robust protection against all forms of radiation induced failures of COTS DDR devices. Our RHBD controller provides means of mitigation for each of the error modes, as shown in Table II. In the following, we provide details on some of the mitigation and detection techniques.

1) *EDAC Implementation*: We implement a byte-wide single error correction-double error detection (SEC-DED) across eight data devices (or similarly a DEC-TED across two rows, i.e., 16-bits data/10-bits ECC). This requires five additional DDR devices as shown in Fig. 12. This method incurs a 20% power penalty compared to a more powerful shortened Reed Solomon RS(22,16) but provides a much lower latency as the SEC-DED calculation is a simple XOR tree. In the case of SEC-DED, 16 encoders/16 decoders are needed to fully encode/decode the 128-bit data word. The data and ECC are interleaved within the 128-bit data word as shown in Fig. 12, so that retrieving each 8-bit mask allows to detect the faulty locations in each die. By re-encoding the decoded word, we provide flags to tell if a correctable was detected or if an error was unable to be corrected. This encoding scheme allows the controller to correct one error across each row so

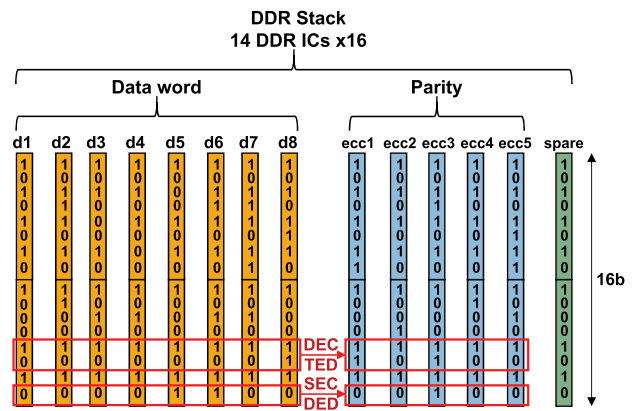


Fig. 12. DDR stack device layout with SEC-DED code (8-b data/5-b ECC) or DEC-TED (16-b data/10-b ECC).

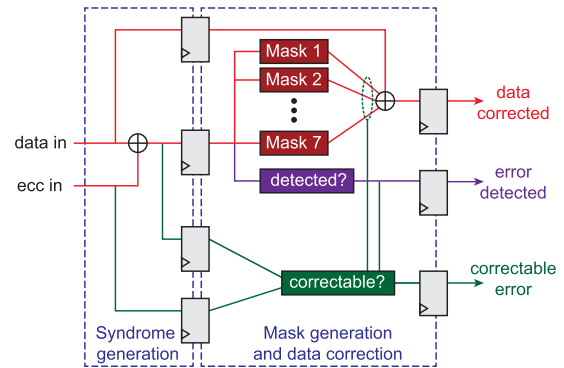


Fig. 13. Pipelining of HSIAO decoder.

that it can potentially correct a full die including nibble errors on one die if all the other dies have no errors.

HSIAO codes are chosen to define systematic linear block codes for SEC-DED. Their fixed code word parity enables the construction of low-density parity-check matrices and fast hardware implementations [15]. They present a high die cost but provide very high performance (encoding and decoding at clock speed) which is important to reduce the overall latency of the controller. Our HSIAO decoder architecture shown in Fig. 13 is pipelined internally with two registering stages to break the computation logic and achieve improved frequency.

SEUs are easily mitigated by our SEC-DED scheme. On the other hand, detection of SEFIs is challenging. Indeed, SEFIs can affect the entire device preventing proper readout of data: temporary burst errors, to faulty rows or columns, or even corrupting an entire die. If SEC-DED provides good protection against SEFIs that cause device failures (i.e., provides full recovery of a failed device), it still is vulnerable to SEUs when a device fails and during device rebuild. To detect SEFIs, the controller contains status and control registers. The EDAC stores a copy of the original data so that the corrected data can be XOR’d with the original data to determine failing bits and increment appropriate bit error count registers. By tracking the level and frequency of high error symbols, we can determine if it is likely that a device has been affected by an SEFI. For the case of SEFI failure the assumption will be made that only one device fails at a time, multiple device failures will cause stack malfunction.

2) *Scrubbing*: As typical error rates are on the order of 10^{-10} per bit-day, or roughly 1 per device-day, SEU buildup is not hard to prevent by regular scrubbing. As the

probability of multiple bit errors within the same ECC code word exponentially increases over the SEU rate, it is possible to perform a nonintrusive, low-rate, background scrubbing (e.g., after each refresh) that is still orders of magnitude less than the expected error rate. To simplify the scrubbing process, rather than maintaining a clean/dirty bit for each ECC code word, the entire DDR device is zeroized on power up with a built-in-self-test (BIST). This will include adding the proper ECC data such the scrub algorithm can proceed and operate properly for uninitialized and unused memory.

Another failure mode due to radiation exposure seen in DDR are stuck bits. Alone a stuck bit does not pose a problem as the ECC will correct it. Radiation testing has shown that stuck bits are not always persistent and can be removed with a device reset. To perform validation during scrub, if a bit error is encountered, the scrubber will store bit error location, repair, write back, and reread. If the data are still in error on reread, the scrubber will repeat the process. The cycle will repeat until a maximum repeat count is hit. The system can trigger a device rebuild if it determines excessive stuck bits are present.

3) *Current Monitoring*: As part of the power conditioning circuitry, a current monitor tracks the relative current draw of each device within the stack. As all the stacked devices are identical, only a relative current profile needs to be monitored to compare current draw that is the absolute value is not critical which relaxes design constraints of the current monitor. The current monitor uses a sufficient RC time constant to average current draw for each measurement interval.

4) *Built-In-Self-Test*: The BIST module is required to perform device initialization and zeroization with varying patterns to simplify the scrub process. The module contains logic to read out and compare the data to provide full self-test capabilities. It can be triggered to operate a full device memory scan as needed. Our BIST implements pattern options for fast in-system testing, such as address, checkerboard, or 1's/0's. Each of these patterns has its strength and weaknesses but they are typically sufficient for power-up type tests to validate the operation of the module. The efficiency versus test time tradeoff of the algorithm chosen is very important as a BIST of the entire memory array can take very long for Gb capacities. As an example, a single-pass address pattern would require less than 4 s at 1600 Mb/s. For extended test capability, we also implemented a powerful March X algorithm [16] that can test address decoding faults, stuck-at faults, transition faults, and some coupling faults. It has a relatively small time complexity of $6n$ where n is the number of cells in the memory array. It follows the following scheme pattern ($w :=$ write & $r :=$ read), where arrows specify the addressing order:

$$\Downarrow (w0); \Uparrow (r0, w1); \Downarrow (r1, w0); \Downarrow (r0).$$

It can be triggered if necessary and if the cube has not to be readily available as this procedure has a high time complexity. Our BIST also supports pattern offsets for each die such that parallel testing can be performed where each die gets an offset.

5) *Rebuild*: The rebuild process consists of a single state machine that can be mixed into any of the DDR device channels. The flow diagram is shown in Fig. 14. The rebuild process requires a read/write cycle for the entire array and it is estimated to finish in less than 5 s at maximum priority. The rebuild logic contains debug registers for ground-level test and characterization and in-flight diagnostics. To prevent SEU buildup the logic tier performs a continuous scrub of the

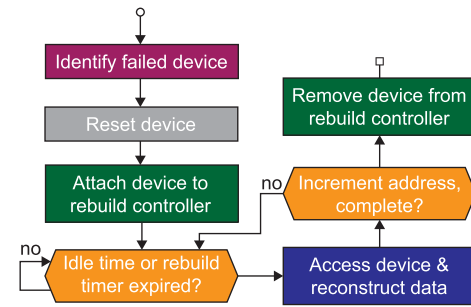


Fig. 14. Flow diagram of DDR device rebuild process.




Fig. 15. Typical DDR bank layout and spiral addressing within exemplar eight-layer stack.

entire memory space at a programmable rate. Through proper EDAC implementations, active rebuilding of a DDR device can be performed without incurring any down time.

6) *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 the data are stored in a “spiral” fashion (in alternate banks among the dies) when viewed through the stack, as shown in Fig. 15. An additional “spiraling” at system level, as shown in Fig. 2, done by alternating data dies with ECC dies in the cube, can be executed to further expand this concept.

7) *Software Conditioning*: Software conditioning proposed in [12] can help address SEFIs without data loss. The degree in which software conditioning can help appears to be manufacturer dependent as Samsung devices did not dramatically improve with software conditioning. However, Herrmann *et al.* [12] showed that most device SEFIs could be removed by applying the $C1$ nondestructive procedure. To clear certain SEFI failures, a full device power cycle may be required however [17]. The $C1$ procedure consists of three operations normally performed during initialization only:

- 1) rewriting the load mode registers of the DRAM die;
- 2) resetting the internal DLL of the DRAM die;
- 3) reperforming the ZQ calibration.

Our controller can perform periodical dynamic software conditioning when required, and independently between dies.

8) *Scheduling-Based Mitigation Features*: An interesting observation in [12] recommends to adopt a closed page policy because idle banks are more sensitive to SEEs. For this reason, whenever timing is not critical, our controller issues a Precharge All after every maintenance operation to close unused rows as soon as possible.

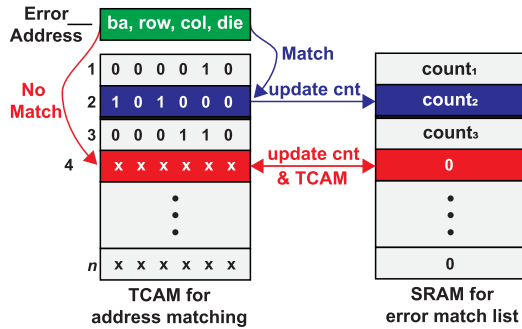


Fig. 16. TCAM-based diagnostic log.

9) *TCAM-Based Diagnostic Log*: A diagnostic log module keeps track of the addresses and characteristics of the errors found by the EDAC. To automatize rebuild and sparing, the Log also keeps track of error counts. His architecture based on a ternary content-addressable memory (TCAM) is shown in Fig. 16. When a new error occurs, the address of the error is compared to the ones already encountered and recorded in the TCAM. In case of a match, the error count is incremented in the SRAM. If there is no match, the address is added to the TCAM. When any of the error counts exceeds a certain threshold, specific maintenance operations are triggered, e.g., sparing if too many errors are found on one die.

10) *Sparing*: 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 are steered from the deficient die to the unused die.

VI. DESIGN FOR FLEXIBILITY

Our design methodology focuses on allowing interchangeability of the memory device to migrate to different densities or technology (e.g., DRAM to MRAM) as well as allow interchangeability of the host processor. Throughout the development, efforts have been made to allow upgradability of the memory devices used in the 3-D-M³ by maintaining a common footprint such that the controller chip can be reused without too much changes in its architecture. Physical limitations on die size exist, but proper rerouting of signal on the RDL can provide compatibility for a certain range of die sizes.

A. Controller Flexibility

The controller has been designed in an open, modular architecture so that it can support multiple operating modes. The host data width ($\times 128$) can be reduced to target more classical widths, at the expense of under-utilizing the full capability of the cube. The controller can also be parameterized to target different commercial DRAM dies by changing various timing parameters (CAS latency, tREFI, etc.) and load mode register values. As alternative to DRAM, DDR3 MRAM chips showed extremely high TID [18]. Thus, an STT-MRAM die can also being switched to, with slight modifications in the controller as changing timings, triggering anti-scribble during calibration, scrambling before power-down, and disabling refreshes.

B. Die Flexibility and RDL Redesign

With the possibility of needing to redesign the RDL layer to accommodate a different DDR3 die, reviews indicated that changes to the RDL design will cascade to the cube bus I/O

TABLE III
DESIGN METRICS OF THE PRELIMINARY 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		

design and to the cube interface board design. Modifying the RDL from one die type to another is difficult due to the different die sizes, I/O pad spacings and pitch, and I/O pad sizes. This is especially true when only a two metal RDL is used. Due to the need to fan-out the RDL circuitry, adjacent die to the one modified die are sacrificed to achieve the necessary fan-out. This is advantageous to attempt to utilize different memory dies to match a set I/O pattern. This will allow the use of single RH controller design that can accommodate different memory. This is sensible cost-wise as a majority of the cost for the memory module is in the RH ASIC controller. To make this concept feasible, we need to move to a four-metal RDL to effectively route out the necessary interconnects from the die to match the I/O pattern on the RH controller. This would allow for two circuit layers, a power layer, and a ground layer, and allow relaxed circuit spacing (i.e., wider traces and spaces) which could relax the necessary circuitry metal thicknesses which in turn may reduce the necessary dielectric thicknesses between layers. A possible drawback with increasing the number of metal layers is that as the RDL process is sequential, the chances of defects increase with increasing layer counts. This can be minimized relaxing some of the design parameters which would help reduce the possibility of defects occurring during the processing. Newer dielectric materials (e.g., photo-definable dielectrics) may also produce more repeatable results, which may help reduce defect density during processing.

VII. CONTROLLER IMPLEMENTATION AND VERIFICATION

A. ASIC Design

To estimate gate counts and performance, we implement a detailed GDSII layout for a preliminary controller RTL, excluding PHYs. 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 3-D-M³. A summary of the preliminary design is shown in Table III. 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 static timing analysis to verify timing power requirements. The controller operates up to 800 MHz. But, 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

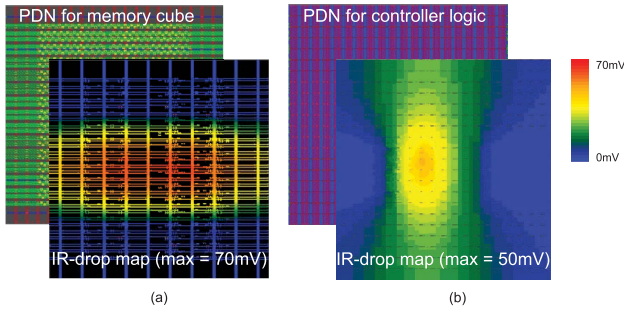


Fig. 17. Full-chip PDN and IR-drop map for (a) memory cube and (b) controller logic.

588.8 mW will likely increase with the addition of postfailure mechanisms and PHY.

We build independent PDNs for the memory cube and controller logic. For the DDR3 PDN, the power distribution is spread evenly across the 18 power pin locations per DDR3 die. Therefore, a total of $18 \times 13 = 234$ current sources draw current from the ideal supply at 1.5 V. Each pin then draws 20 mW. The DDR3 PDN has to run from the VDDQ/VSSQ pins in the bottom of the package to the top metal layer to connect to the respective DDR3 power/ground pins. For power drop analysis, planned rows of power/ground connections are used for a dedicated DDR3 PDN. A total of four metal layers are used in the logic tier, dictating the DDR3 pads to Metal 4. In our DDR3 PDN design, the central region has no M1 VDDQ/VSSQ lines. Therefore, the DDR3 dies around the center will experience higher IR drop than the ones at the top/bottom edges. Overall wire density for the DDR3 PDN is small for a 10 mm \times 10 mm footprint and there is sufficient room to fit in the logic PDN and signal routing. The regular arrangement of M3 wires is used to have uniform and wide connections for all power/ground pins of DDR3. Initial PDN simulations with Cadence Voltus showed that less than 5% IR-drop (supply voltage of 1.5 V) could be achieved for the power path from BGA pad, to TSVs, through four metal layers route, and finally to the DDR device pad [see Fig. 17(a)].

The logic PDN planning is much more straightforward. The power/ground C4 bumps supply VDD/VSS to M1 TSV landing pads unlike 2-D IC PDN where they supply VDD/VSS to top metal which is then distributed. The advantage of having connections from the bottom is that top metal layers routing resources are not required for the PDN. Instead, a wide line on M1 and M2 along with the cell power/ground rails can connect all logic cells to PDN. The logic PDN shares metal resources with the already present DDR3 PDN. The metal stripes are carefully planned to have efficient use of metal resources while allowing sufficient room for signal routing. Logic PDN IR drop analysis follows the standard state-of-the-art 2-D IC flow with power assigned to each standard cell based on power simulation and the power-grid/pin locations. The cells behave as current sources which cause an IR drop at the VDD rails on the cell. The cell power values are obtained from power analysis with Synopsys PrimeTime. Fig. 17(b) shows the IR drop map for the logic PDN. Logic supply voltage is at 1.5 V. Since most logic is concentrated slightly left off-center, maximum IR drop occurs in that region. However, this drop is well below 5% of VDD that is 75 mV. The careful design planning of VDD/VSS supplies from M1 ensures that the power connections are close to the currents source (standard cells), leading to low IR drop. The maximum IR drop in this




Fig. 18. IROC SoCFIT simulation.

TABLE IV
DESIGN METRICS OF THE FPGA LOGIC CONTROLLER

Resource	Utilization	Utilization %
Slices	5,686	6%
Slice LUTs	10,674	3 %
Slice Registers	13,063	1 %
BRAM36E1s	11	1 %
Bounded IOBs	809	67 %
BUFIODQS	28	19 %
BUFGs	4	12 %
IDELAYCTRLs	18	50 %
IODELAYE1s	280	19 %

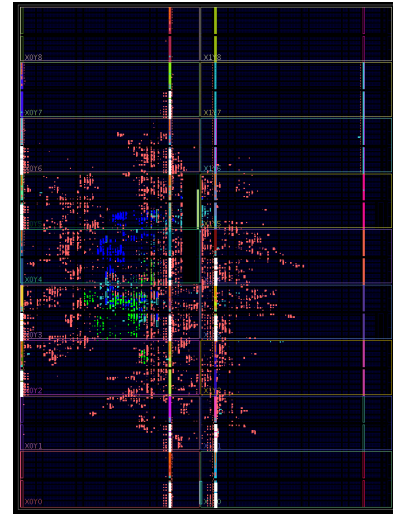


Fig. 19. Floorplan of the FPGA-mapped controller design: 14 PHYs, control logic, and EDAC.

case is 50 mV. It is noted that the IR drop map is much denser than the DDR3 PDN case. This is because logic PDN covers almost all M1 cell rails and hence has more components on the map.

Final gate level netlist is obtained after layout from Cadence Innovus. Timing information for each cell in standard delay format (SDF) is obtained from Synopsys PrimeTime. These two design data along with standard library description in Verilog format are fed to IROC SoCFIT to generate soft error rate (SER) evaluation as shown in Fig. 18. The results of the SER process simulation, using the GF 130-nm libraries, indicate the total SER is 1.54 FIT/device where FIT are measured in units of errors per billion hours of use. These results are positive indicators for the design concept.

B. FPGA Implementation of the RTL Controller

Though the memory controller architecture easily fits within a commonly available FPGA, the initial design did exhibit some performance issues owing to high I/O requirements of nearly 1000 pins, due to the need of the controller to communicate with 14 memory dies as well as with the host memory controller. A Xilinx XC6VLX550T-1FF1760C

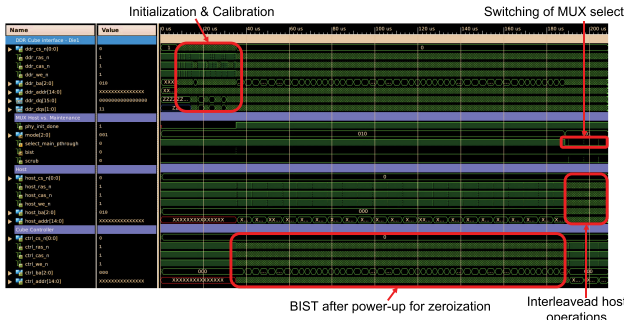


Fig. 20. Simulation of the controller.

FPGA with 1200 IOBs was chosen as target device. The high I/O utilization constrains placement and results in potential routing delays that reduce operating speed. Typical FPGA requirements for pin placement (e.g., allocate only center banks for address and control pins) cannot be respected due to the I/O limitations. This is however an “FPGA only” issue. The final ASIC design overcomes these limitations due to custom I/O placement and area-array TSVs. Therefore, the FPGA design focus is on supporting minimum operating speeds of the DRAM dies (300 MHz as baseline goal with DLL ON).

Manual placement of the IOBs primitives (I/O-DDRs, IODELAYS, and read data capture flip-flops) has been completed for the 14 PHYs to ensure proper timing closure. The location of these flip-flops and the routes between the IDDR and fabric flip-flops must be carefully matched. The Xilinx ISE tool was used to map the controller architecture shown in Fig. 10 on the Virtex-6, with the exception of the DDR host PHY currently in development. We achieved timing closure at 300 MHz. Fabric utilization of the FPGA resources is shown in Table IV and Fig. 19 shows the corresponding mapping. The utilization is expected to change substantially for the ASIC implementation with the addition of logical radiation-hardening techniques to detect and protect against soft errors, such as distributed triple modular redundancy, protection of redundant logic, and fault-tolerant FSM implementation.

C. Simulation

We carried out extensive Verilog simulations using ISE Simulator from Xilinx on the RTL code to test and verify the functionality of the controller in handling different memory operations and its ability to perform the required maintenance. The test benches include mechanisms to properly test SEE and SEFI type failures. In our test benches, the logic controller connects to the 14 PHYs, themselves connected to their corresponding DRAM die (8 Gb × 16 DDR3 Verilog model from Micron) through PCB traces simulated by wire delay models. Simulations (see Fig. 20) show correct operation of the controller, including DRAM initialization and calibration, read/write operation with refreshes, bank spiraling, variable rate refresh, self-Refresh low power mode, ZQ Calibration, BIST, software conditioning as well as effective working of the idle detector in switching from normal to maintenance modes. To specifically test the interaction with the host, we developed a host emulator with an integrated memory controller and PRBS-based address/data pattern generators.

VIII. BOARD

To provide breadboard validation of the design prior to fabrication of the final custom ASIC controller tier, we develop

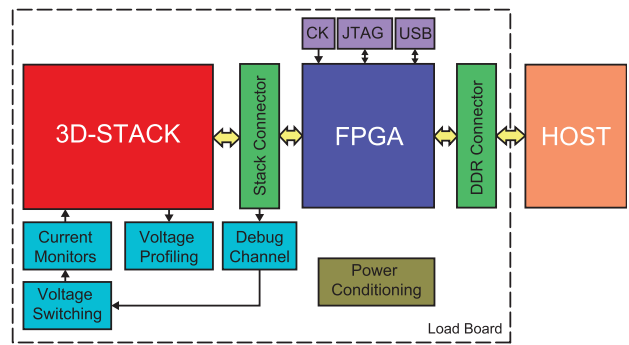


Fig. 21. Test board architecture.




Fig. 22. Cube substrate to FPGA board connection concept.

an FPGA-based breadboard. The FPGA emulating the function of the controller will connect to an actual 3-D-M³.

A. FPGA Test Platform

The block diagram of the test board is shown in Fig. 21. In addition to integrating the FPGA-emulated controller and 3-D-M³, the test board includes the following features:

- 1) power supply current monitor for applicable rails;
- 2) external DDR3 interface for connection to host processor;
- 3) connector-less logic probe pads to enhance debugability;
- 4) FPGA side channel I/O for real-time performance monitoring and logging (e.g., a universal serial bus 3.0 (USB3.0) interface is used for memory dumps and allow for uploading particular memory patterns to simulate error conditions and faults);
- 5) micro jumpers to allow “hard” error testing;
- 6) voltage switching for DDR layers.

As the FPGA fabric itself does not support device power switching, external FETs will be utilized. The FPGA will support analog-to-digital converters to allow power monitoring.

The test board requires schematic design, layout, and careful signal integrity analysis to insure proper operation margins. This task is critical as the interconnect between cube and FPGA emulator is physically much further apart. This requires impedance matching, length matching, and crosstalk evaluation. When the final design is migrated to an ASIC, the cube interconnect electrical concerns are virtually eliminated due to the extremely short interconnect (≤ 5 mm from die pad).

B. Custom Substrate for Cube Mounting

The concept for integrating the stack onto the FPGA board is shown in Fig. 22. Low-profile spring arrays (interposers)




Fig. 23. Cube substrate composite of signal layers.




Fig. 24. 3-D simulation of VDDQ plane and DDR cube layer reroute.

in the four quadrants allow a ceramic high-density build-up substrate to connect to the load board on the peripheries of the stack substrate. This provides thermal expansion stress relief, fan-out for the tight BGA pitch (necessary surface area for mounting hardware and alignments, the compressive forces required by the interposer are absorbed away from the stack BGA interface) of the memory cube as well as sufficient electrical performance. This solution has also the benefit of easy swapping of parts. Secure contact over the connector areas is made by pressure distribution compression rings.

1) *BGA Interface*: The most critical area of the substrate is the cube BGA interface. The pad size is estimated at $150\ \mu\text{m}$ with $150\text{-}\mu\text{m}$ clearance. The pads will be nonsolder mask defined to provide a well-controlled surface area for each ball. The stack uses a staggered BGA pattern which dictates that the I/Os route out from one side (see Fig. 23) for all the layers since the vertical direction is blocked by the BGA pattern itself. To maintain signal integrity, all DDR signal lines from one device are routed on the same layer in a stripline configuration. Due to the staggered pattern, the escape pattern requires the use of blind vias and seven signal layers to allow for each layer to route out unobstructed by adjacent layers.

2) *Power Distribution Analysis*: The substrate requires careful design of the PDN as the co-fired metallization used in the ceramic buildup structure provides a high resistivity. The typical sheet resistance of the tungsten planes is $10\text{--}15\ \text{m}\Omega/\text{square}$. To test all features of the cube it is necessary to individually connect the power rail of each die. The challenging part is connecting the voltage planes to each DDR device as each layer is roughly on a 1-mm pitch. Thus, the plane layout is going through multiple iterations of layout-simulation-adjust to minimize the plane resistance. Each DDR plane connection has a narrow area just beneath the cube—this gets further reduced by the necessary via anti-pads. To estimate the total resistance to the layer and properly account for the location of the BGA power balls, we perform a full 3-D simulation




Fig. 25. Eye diagram of data signal DQ FPGA to DDR.

including the power plane, vias, and RDL on the die. The results shown in Fig. 24, describing one of the fourteen VDDQ layers, indicate about $140\text{-m}\Omega$ resistance from the connector to the die. Moreover, the narrow area under the BGA has the highest current density. The ceramic substrate has relatively high thermal conductivity which should help to reduce the self-heating of the plane layers.

3) *Electrical Analysis*: An extensive electrical analysis of the substrate design was performed to validate the substrate performance. Since the substrate being designed is a co-fired ceramic structure with thick film tungsten metal traces and via fills, it was critical to evaluate the electrical performance because tungsten has a much higher resistivity than copper. Copper, typically used for PCBs, is approximately $5\times$ better in conductivity than tungsten. Therefore it was necessary to incorporate various design criteria (e.g., doubling up on the power planes) to compensate for the difference in conductivity. Simulations were performed at $1333\ \text{MT/s}$ ($666\ \text{MHz}$). The main simulation setup consists of two drivers and two receivers to allow simulation of crosstalk. The substrate is modeled as a coupled stripline to match the ground-signal-signal-ground configuration of the ceramic substrate. Lengths were pulled from the board layout using the longest trace lengths. Simulation for the FPGA to DDR link is shown in Fig. 25. The data link in this direction uses ODT at the stack and a DCI FPGA driver set for roughly $40\text{--}50\ \Omega$. Results indicate a 440-ps aperture and an adequate voltage margin of $180\ \text{mV}$.

Also reviewed for the substrate design was the via performance. Due to the tight pitch of the stack, small vias ($\sim 100\ \mu\text{m}$) are used. As a risk mitigation, it was desired to design the substrate to support a cube design with no internal capacitors in case fabrication or performance became issues. To support off-stack capacitors, simulations showed that the best location would be directly underneath the stack that is on the bottom of the substrate directly on the vias. This is however limited by the small vias and their corresponding inductance. Fig. 26(a) shows the vias' path setup. To validate the initial calculations, a 3-D analysis was performed using the Fast Henry field solver. The design dimensions of the substrate were used to represent the via structure in detail. A simulation is run for the full via field (multiple sets of triplets) as shown in Fig. 26(b). Results indicate $\sim 500\text{-pH}$ inductance. Combined with on-die DDR capacitance, this provides an appropriate solution. Note that the baseline configuration is for in-stack capacitors which will reduce the overall inductance by orders of magnitude due to the close proximity within the stack.

C. Host Emulator

As no space processor is currently available to test and demonstrate the memory module, we use a host emulator for




Fig. 26. (a) Power via modeling setup and (b) 3-D simulation of via field per die.

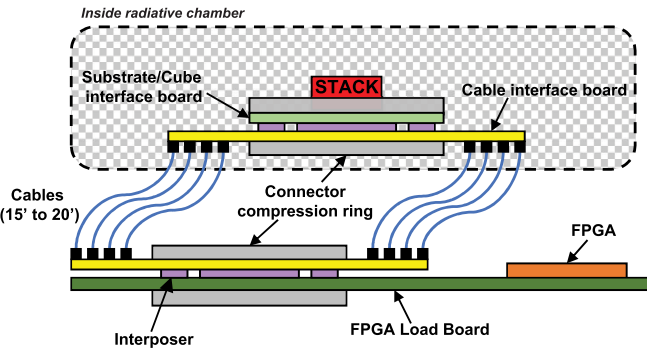


Fig. 27. FPGA radiation testing environment.

the final test setup to exercise the memory controller and perform pattern testing, error reporting and handle user interface tasks. A few approaches for host emulation were considered. One is to use an off-the-shelf motherboard with DIMM slots with a custom adapter to allow the FPGA emulator to connect to the DIMM slots via a cable. This provides a very flexible interface and a proof-of-concept of operability with a standard memory controller (as the motherboard will accommodate an actual CPU). A limit, however, is the length of additional cabling required as well as minimum speed setting which may interfere with timing control if required settings are out of range (most motherboard designs are tailored for the high-end not the low end). A dual channel is required to support the full 128-bit width and the open source MemTest86 program provides necessary software for performing memory testing on a motherboard platform. The other solution would be to use an FPGA-based host emulator. This option is the most flexible in terms of adapting to required pinout and cabling requirements but it has more limited test features.

D. Radiation Test Environment

In an effort to ultimately subject the 3-D-M³ to radiation testing, we propose a custom test setup. For performance concerns, the interconnection between the 3-D-M³/interface substrate to the FPGA emulator requires close proximity of the FPGA to the cube. However, this configuration does not lend itself to radiation testing as the components on the FPGA board would not be shielded from the radiation source and it would be too expensive to procure radiation hardened components for the FPGA board. A low-speed approach for radiation testing using two boards is proposed and shown in Fig. 27. The setup interconnects the two boards through a series of cables. There will be about 1500 I/Os interconnected

between the two boards. For TID testing, the DUT has to be in front of the radiation source so cable lengths in the 15’–20’ range will adequately isolate any test electronics from the radiation source. For SEE testing, shorter cables can be used as the radiation source is more pinpointed. The extraction of system level results and the analysis of the detailed error behaviors of the proposed system for space applications is targeted for future work.

IX. CONCLUSION

This research provides an advanced technology enabling new space missions not possible with the current bulky, low-bandwidth, low-density, and power-hungry memory devices. We developed and demonstrated a radiation tolerant stacked memory array based on state-of-the-art chip stacking and radiation mitigation technologies. Our module can be directly connected to a host processor and act as a highly reliable DDR3 module thanks to its integrated RHBD controller. Using our LOB technology to stack COTS memory dies, we avoid the use of proprietary TSV dies, achieve high memory capacity and good radiation tolerance. The controller is tested in hardware and a complete test board is in development to subject the cube to radiation sources and assess the performances of the controller. Our modular architecture allows flexibility in the choice of the memory die and our cube development design methodology is intended to be a building block that provides a path to additional opportunities such as integration of nonvolatile memory or computing resources.

We believe the technology developed in this work can find utility in pure commercial applications such as high-performance computing. Our 3-D-M³ technology is applicable for direct attachment to graphical processing units substrates to improve performance (i.e., die-to-die interconnection, and DDR interfaces not be required to go off package). These stacks use bare silicon die that are rarely dynamically burned-in, resulting in infant mortality that can take out an entire memory stack. The proposed 3-D architecture and controller design will provide fault tolerance to overcome this problem to enable much taller, and hence, more capable stacks.

REFERENCES

- [1] A. Agnesina *et al.*, “A novel 3D DRAM memory cube architecture for space applications,” in *Proc. 55th ACM/ESDA/IEEE Design Autom. Conf. (DAC)*, Jun. 2018, pp. 1–6.
- [2] A. Agnesina, J. Yamaguchi, C. Krutzik, J. Carson, J. Yang-Scharlotta, and S. K. Lim, “Bringing 3D COTS DRAM memory cubes to space,” in *Proc. IEEE Aerosp. Conf.*, Mar. 2019, pp. 1–11.
- [3] D. U. Lee *et al.*, “25.2 a 1.2 V 8Gb 8-channel 128GB/s high-bandwidth memory (HBM) stacked DRAM with effective microbump I/O test methods using 29 nm process and TSV,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2014, pp. 432–433.
- [4] J. T. Pawlowski, “Hybrid memory cube (HMC),” in *Proc. IEEE Hot Chips 23 Symp. (HCS)*, Aug. 2011, pp. 1–24.
- [5] S.-K. Lim, “Bringing 3D ICs to aerospace: Needs for design tools and methodologies,” *J. Inf. Commun. Converg. Eng.*, vol. 1, no. 2, pp. 117–122, Jun. 2017.
- [6] P. M. Gouker *et al.*, “SET characterization in logic circuits fabricated in a 3DIC technology,” *IEEE Trans. Nucl. Sci.*, vol. 58, no. 6, pp. 2555–2562, Dec. 2011.
- [7] R. K. Kirschman, “Low-temperature electronics,” *IEEE Circuits Devices Mag.*, vol. 6, no. 2, pp. 12–24, Mar. 1990.
- [8] *TN-41-01: Calculating Memory System Power For DDR3*, Micron, Boise, ID, USA, 2007.
- [9] Y. Lin and J. C. Joy, “A new examination of secondary electron yield data: Surface and interface analysis,” *Int. J. Devoted Develop. Appl. Techn. Anal. Surf., Interfaces Thin Films*, vol. 37, no. 11, pp. 895–900, 2005.

- [10] M. Park *et al.*, "Soft error study on DDR4 SDRAMs using a 480 MeV proton beam," in *Proc. IEEE Int. Rel. Phys. Symp. (IRPS)*, Apr. 2017, p. SE-3.
- [11] D. M. Hiemstra, "Guide to the 2007 IEEE radiation effects data workshop record," in *Proc. IEEE Radiat. Effects Data Workshop (REDW)*, Jul. 2017, pp. 1–4.
- [12] M. Herrmann *et al.*, "New SEE test results for 4 Gbit DDR3 SDRAM," in *Proc. RADECS Data Workshop*, 2012, pp. 1–5.
- [13] *UG086 Memory Interface Generator, User Guide V3.6*, Xilinx, San Jose, CA, USA, 2010.
- [14] S. Rixner, W. J. Dally, U. J. Kapasi, P. Mattson, and J. D. Owens, "Memory access scheduling," *SIGARCH Comput. Archit. News*, vol. 28, no. 2, pp. 128–138, May 2000.
- [15] V. Gherman, S. Evain, N. Seymour, and Y. Bonhomme, "Generalized parity-check matrices for SEC-DED codes with fixed parity," in *Proc. IEEE 17th Int. On-Line Test. Symp.*, Jul. 2011, pp. 198–201.
- [16] T. Koshy and C. S. Arun, "Diagnostic data detection of faults in RAM using different march algorithms with BIST scheme," in *Proc. Int. Conf. Emerg. Technol. Trends (ICETT)*, Oct. 2016, pp. 1–6.
- [17] F. Gliem *et al.*, "Memory technology trends and qualification aspects," IDA, Braunschweig, Germany, Tech. Rep., 2012. [Online]. Available: https://indico.esa.int/event/67/contributions/3053/attachments/2453/2826/1535_-_Memory_technology_trends.pdf
- [18] J. Heidecker, "MRAM technology status," NASA Jet Propuls. Lab., California Inst. Technol., Pasadena, CA, USA, Tech. Rep. 13-3 2/13, 2013.



Anthony Agnesina (Graduate Student Member, IEEE) received the Diplôme d'Ingénieur from the CentraleSupélec, Gif-sur-Yvette, France, in 2016, and the M.S. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2017, where he is currently working toward the Ph.D. degree at the School of Electrical and Computer Engineering.

His current research interests include 3-D memory architectures, computer-aided design of VLSI circuits, and applied machine learning to electronics.



James Yamaguchi received the B.A. degree in chemistry in 1977, and the B.S. degree in chemical engineering from California State University at Long Beach, Long Beach, CA, USA, in 1979.

He joined Irvine Sensors Corporation, Costa Mesa, CA, USA, in 1993, where he is currently the Vice President (VP) of the 3-D Electronics and Mass Storage Group. He has over 39 years of hands-on processing experience in the areas of 3-D packaging, thin film deposition, electroplating, high-density MCM fabrication, and printed wiring board (PWB)

fabrication. He has extensive research and development experience in fabrication technologies, process design and integration, facility operations, and technology transfer. He is a co-inventor on 14 patents and has coauthored six journal articles.



Christian Krutzik received the master's degree in electrical and computer engineering from the University of California at Irvine, Irvine, CA, USA.

He is currently a Senior Electrical Engineer at Irvine Sensors Corporation, Costa Mesa, CA, USA. He has over 18 years of experience involving electrical design and evaluation of 3-D stacked modules and system miniaturization. Projects have included solid-state devices (SSD) design and development, compact 1TB universal serial bus (USB) SSD drives, NAND flash characterization for secure erase products, FPGA development for LiDAR applications, miniaturized and wearable biotelemetry sensors, and acoustic processing/telemetry over RF using hearing-aid technology. He was also involved with a project to evaluate NAND flash devices for data remanence which included low-level analysis of flash devices and intricate knowledge of standard SSD processor capabilities. He is also involved with the design and development of a secure SSD product line and has experience with developing

and debugging serial advanced technology attachment (SATA), PCIe, and other high-speed interfaces. Further experience at Irvine Sensors Corporation includes stacked computer systems running Linux, various RF interfaces, high-density stacked memories, storage devices such as SSD's, FPGA's, embedded firmware, software, and DSP-based systems. He also has experience with multiple electrical simulation tools such as HyperLynx, HSPICE, and other finite element analysis (FEA) tools. Software experience includes C, Python, MATLAB, Javascript, PHP, HTML, assembly, and Labview on both Linux and Windows platforms.

and debugging serial advanced technology attachment (SATA), PCIe, and other high-speed interfaces. Further experience at Irvine Sensors Corporation includes stacked computer systems running Linux, various RF interfaces, high-density stacked memories, storage devices such as SSD's, FPGA's, embedded firmware, software, and DSP-based systems. He also has experience with multiple electrical simulation tools such as HyperLynx, HSPICE, and other finite element analysis (FEA) tools. Software experience includes C, Python, MATLAB, Javascript, PHP, HTML, assembly, and Labview on both Linux and Windows platforms.



John Carson is currently the President, the Chief Executive Officer (CEO), and a Founder of Irvine Sensors Corporation, Costa Mesa, CA, USA. Upon graduation from MIT in 1961, he joined Baird Atomics, Inc., Cambridge, MA, USA, and Baird Atomics, Inc., Waltham, MA, USA, where he became a Project Engineer and an Assistant Program Manager on a space-based infrared surveillance program. He left Baird Atomics, Inc., in 1967, to form a consulting company for sensor systems design and development, with clients that included North American Aviation, Los Angeles, CA, USA, Honeywell, Charlotte, NC, USA, Lockheed, Calabasas, CA, USA, Ford Aeronautics, Newport Beach, CA, USA, and Grumman. Over the years, this company evolved into Irvine Sensors Corporation. He has been responsible for the design and development of space-based military and NASA sensor systems, commercial spectrometers, commercial laser printers, 3-D electronics, and various neuromorphic sensor and electronics systems for the military. He is the holder of 23 patents. In addition to his Irvine Sensors duties, he has chaired the Industry Advisory Board for the Caltech Center for Neuromorphic Systems Engineering, Pasadena, CA, USA, and has served on the Scientific Advisory Board for the Egg Factory, an incubation firm headquartered in Roanoke, VA, USA.



Jean Yang-Scharlotta (Member, IEEE) received the B.S. degree in chemical engineering from The University of Texas at Austin, Austin, TX, USA, in 1990, and the M.S. and Ph.D. degrees in chemical engineering from Stanford University, Stanford, CA, USA, in 1996.

She conducted engineering research at the Los Alamos National Laboratory, Los Alamos, NM, USA, from 1987 to 1991 on radioisotope chemistry prior to graduate research in self-assembled monolayers at Stanford University. From 1996 to 2008, she was a Device Technology Engineer then the Department Manager at the Advanced Micro Devices focusing on the introduction and scaling of new flash memory device technologies. She ushered in the first three generations of the MirrorBit™ devices and technologies. She has been a Senior Microelectronics Specialist at the Components Assurance Office, NASA Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, CA, USA. She holds more than 70 U.S. patents and is the author of many publications on memory devices and technologies. Her research interests include memory technologies, electronics and materials in radiation and extreme environments, physics of failure, interfacial phenomenon, and technology development.



Sung Kyu Lim (Senior Member, IEEE) received the B.S., M.S., and Ph.D. degrees from the University of California at Los Angeles (UCLA), Los Angeles, CA, USA, in 1994, 1997, and 2000, respectively.

He joined the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA, in 2001, where he is currently the Dan Fielder Endowed Chair Professor. His current research interests include modeling, architecture, and electronic design automation (EDA) for 3-D ICs. His research on 3-D IC reliability is featured as a

Research Highlight in the *Communications of the ACM* in 2014.

Dr. Lim was a recipient of the National Science Foundation Faculty Early Career Development (CAREER) Award in 2006. He received the Best Paper Awards from the IEEE Asian Test Symposium in 2012 and the IEEE International Interconnect Technology Conference in 2014. He has been an Associate Editor of the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS since 2013.