

# Bringing 3D COTS DRAM Memory Cubes to Space

Anthony Agnesina  
Department of ECE  
Georgia Institute of Technology  
Atlanta, GA 30332  
404-894-0373  
agnesina@gatech.edu

John Carson  
3D Stacking Systems  
Irvine Sensors Corporation  
Costa Mesa, CA 30332  
714-444-8725  
jcarson@irvine-sensors.com

James Yamaguchi  
3D Stacking Systems  
Irvine Sensors Corporation  
Costa Mesa, CA 30332  
714-444-8785  
jyamaguchi@irvine-sensors.com

Jean Yang-Scharlotta  
NASA Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA 91109  
818-354-0412  
jean.yang-scharlotta@jpl.nasa.gov

Christian Krutzik  
3D Stacking Systems  
Irvine Sensors Corporation  
Costa Mesa, CA 30332  
714-444-8700  
ckrutzik@irvine-sensors.com

Sung Kyu Lim  
Department of ECE  
Georgia Institute of Technology  
Atlanta, GA 30332  
404-894-0373  
limsk@ece.gatech.edu

**Abstract**—This paper details the architectural choices and implementation challenges faced in the building and validation of a space-qualified 3D DRAM memory system, in an effort to offer high memory capacity, increased bandwidth, fault tolerance and improved size-weight-and-power characteristics needed for harsh space mission environments. Our novel horizontal 3D stacking technology called “Loaf-Of-Bread” (LOB) is used to integrate multiple Commercial-Off-The-Shelf (COTS) DRAM memory dies into a cube structure (3D-M<sup>3</sup>). A custom Radiation-Hardened-By-Design (RHBD) controller sitting underneath the cube supplements the 3D-M<sup>3</sup> in addressing COTS radiation weaknesses by including advanced SEU and SEFI mitigation features such as error detection and correction, scrubbing, device data rebuilding and die management. We developed a custom DDR physical layer (PHY) for 14 independent dies to connect the 3D-M<sup>3</sup> to its controller.

Validation and functional evaluation of the ASIC controller will be conducted prior to tape-out on a custom FPGA-based emulator platform integrating the 3D-stack. The selected test methodology ensures high-quality 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 case of technology roadmap changes or unsatisfactory radiation results.

## TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. CUBE DESIGN CONSIDERATIONS .....	1
3. CONTROLLER ARCHITECTURE .....	3
4. DESIGN FOR FLEXIBILITY .....	6
5. FPGA TEST BOARD.....	7
6. CONCLUSIONS.....	9
ACKNOWLEDGMENTS .....	10
REFERENCES .....	10
BIOGRAPHY .....	10

## 1. INTRODUCTION

The on-board capabilities of space computing have not kept up with the needs of current and future important NASA mission scenarios, which today cannot be accomplished in a robust and cost-effective fashion. In particular, the Deep Space exploration program has unfulfilled and unique

requirements such as high-performance data processing or autonomous operations, along with extreme needs for energy management, efficiency, fault tolerance and resilience. Although technology development efforts are underway into developing a High performance Space Computing/Next Generation Space Processor (HPSC/NGSP), a radiation hardened multi-core computing processor expected to fuel missions until 2030, they do not address the limitations of current on-board memory systems. 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. The 3D memory cube can potentially be utilized in many aspects of the spacecraft beyond just the main compute/processor. Other spacecraft systems such as instruments, detectors, sensors, and communication, may utilize our high density, high bandwidth memory cube to enable high performance operations.

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 approach is to leverage COTS memory devices to maximize memory density, bandwidth and speed, by integrating them into a 3D memory cube 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 complete memory module can be directly connected to a host processor and act as a hot-pluggable DDR3 module of large capacity (8GB). The controller has been tested on an FPGA from the Xilinx Virtex-6 family and a complete test board structure is in development to subject the cube to radiation sources. This paper is intended to be a natural continuation and complement the work from the same authors presented in the introductory article [1].

## 2. CUBE DESIGN CONSIDERATIONS

After a long wait, three-dimensional integrated circuits (3D ICs) are now becoming a fairly mainstream technology, with in particular applications into consumer 3D memory cubes such as the High Bandwidth Memory (HBM) [2] and Hybrid Memory Cube (HMC) [3], where multiple memory dies

are vertically integrated and connected through thousands of through-silicon vias (TSVs), to offer “more than Moore” improvements in terms of size, capacity, speed, and power consumption. This type of memory parts could truly enable the aforementioned space missions but unfortunately none of these are currently ready for space. In particular, it is unclear how TSVs behave when deployed in cryogenic temperatures environments [4].

### Leveraging TSV-based 3D IC?

Very little work has been found on 3D IC reliability analysis for space applications. Gouker et al. [5] studied radiation effects on 3D IC for the first time and characterized Single Event Transients (SETs) for a small-scale logic circuit implemented into a 3-tier 3D IC using a 180nm SOI process. Each tier contains a simple circuitry that measures SET pulse width and is connected with  $1.25\mu\text{m}$  TSVs. Their heavy-ion testing based on krypton shows that the pulse width distribution across the tiers exhibits non-trivial tier-to-tier variations. They observed that various materials used in the 3D IC lead to different levels of energy deposited during the radiation testing. The authors also conducted a similar study using proton and neutron beams on an SRAM module implemented in the same 3D 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. In addition, the designs used are nowhere close to commercial memory cubes as such is limited and mainly focused on processing-in-memory architecture research. 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 would be 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 of designing both a radiation hardened memory chip and a radiation hardened controller chip is fairly complex and costly. Also, since the overall design would be relying on a specific memory, it would limit the ability of the module to interchange memory.

### The COTS Solution

These issues rationalize supporting the path to utilize a 3D- $M^3$  of COTS memory with a custom RHBD controller. COTS devices have had a lot of scrutiny and interest in the space community. Extensive testing on SDRAM COTS has been undertaken over the recent years to understand their performance in a space environment. NASA has successfully used SDRAMs COTS in spacecraft for critical applications such as for the Compute Element within the Mars Science Laboratory Curiosity rover.

As presented in [1], with proper component screening and radiation testing, the use of COTS devices within our memory structure provides a low cost and effective solution to capture

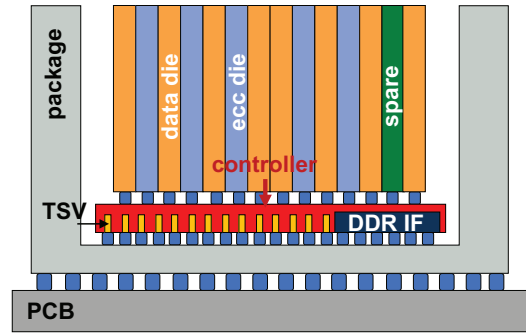


Figure 1. Our 3D-Memory package

the benefits offered by state-of-the-art commercial devices. Leveraging COTS memory devices stacked in a “Loaf-of-Bread” (LOB) fashion theoretically achieves comparable performances to the HMC, plus in particular increased density, possible intermixing of different memory types such as 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.

### Our Stacking Method

In order to stack the COTS memory devices factors such as die orientation, number of dies within the cube, die density, and IO pad out to the controller chip were analyzed. Our final stack configuration contains fourteen DDR dies stacked in a LOB fashion as shown in the system-level description in Figure 1. Eight DDR x16 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 (3D- $M^3$ ) flip bumped to the logic controller tier. Interconnection of the active circuitry on the top of the controller to the bottom IO pads is made through TSVs. These TSVs will be filled with copper to eliminate the 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 [6]. 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 3D- $M^3$  to additional reflow processes during the assembly as well as requiring TSVs, a single cavity package has been chosen over a dual-cavity package because it has the advantage of allowing to test the completed 3D memory module prior to assembly into the package. Moreover the dual cavity requires an extremely high IO count to pass through the substrate which poses technical risks.

In the LOB cube configuration detailed in Figure 2, it is more readily achievable for the controller to interact with each individual die within the cube, compared to a vertical stacking method like in traditional TSV-based stacks. A Re-Distribution Layer (RDL) brings the necessary interconnections from each die IOs to the bottom, where an Under Bump Metal (UBM) is used to form the IOs of the cube. The LOB structure can be a tall cube depending upon the die X-Y footprint. The list of benefits for the LOB structure listed in [1] such as point-to-point logic-to-memory interconnects, low thermal impedance, IO connectivity able to support individual die access for recover and rebuild process, are important and very desirable for mitigating radiation effects,




Figure 2. Our proposed “Loaf-of-Bread” structure

which makes the LOB structure a clear choice for the 3D-M<sup>3</sup> design.

### 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 approximately a 45% increase in single bit upset cross-section and a 17% logic upset decrease [7]. 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.

The reviews from [8], [9], [10] showed that in general DDR3 devices, typically from Micron, Samsung or Hynix, are strong candidates for implementation into space memory applications. No occurrences of destructive failures or single-event latchup were noted for any of the devices tested. A general observation is that DDR devices are more sensitive to single-event functional interrupts (SEFIs) than single-event upsets (SEUs). All devices tested showed relatively good resistance to total ionizing dose (TID) effect, with devices able to be used from ~100 krad (Si) up to ~400 krad. In Micron devices, idle current increases dramatically as the TID approaches the 100 krad range, whereas the Samsung and Hynix devices tend to show a more gradual increase in idle current with increased TID exposure. A noticeable trend in terms of tolerance to SEUs seemed to play out: during SEE Proton testing, SEUs and Column/Row SEFIs showed an increase towards low linear energy transfer. Typically, Hynix devices showed the best resistance to SEUs than Samsung and Micron devices. In some instances, the Samsung and Hynix devices showed orders of magnitude better cross sections than the Micron devices.

## 3. CONTROLLER ARCHITECTURE

The top-level block diagram of our controller is shown in Figure 3. The architectural choices are based on the main requirements that the controller has to integrate smoothly with a host DDR interface so that the memory could be actively used by the processor — namely serve its requests with very low latency, as well as control and perform the necessary maintenance (error-mitigation) operations on the

14 independent dies memory stack.

The necessary features required for the controller and addressed in the RTL code are in particular:

- read/write capability to the stack,
- wide 14 DDRs physical interface (PHY) to the cube,
- DDR interface to Host,
- EDAC,
- access to each individual die within the cube for individual power management, ON-OFF or rebuilding,
- housekeeping,
- 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.

### Host Interface Considerations

*Latency Concerns*—The overall latency of the core processing is critical for the RTL design. Indeed, 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 is brought back from the cube to the host. Therefore, the latency of the path Host DDR PHY to MUX to Stack PHY to CAS Latency of the cube to EDAC and back to the Host has to be reasonable (not much more than twice the CAS latency) to be able to be handled by the host integrated memory 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.

*Flexibility Concerns*—A programmable bus width interface on the host side allows data widths densities (baseline of x128) to be achieved depending on the host memory controller specifications. Note that due to the buffering of the internal DDR3 die, the 3D-M<sup>3</sup> 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.

*Integration Concerns and Idle Detector*—As it is expected for the controller to be “transparent” to the host, namely 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 using an “Idle” gpio input signal is provided to allow the host system to notify the cube that the DDR3 bus will be idle – this will be required if a rebuild is required and the host desires to have it completed immediately. It does however present additional requirements to the host processor memory controller but it is designed as an optional feature. 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 the operation of our controller from Normal mode (i.e. a slave that serves the requests of the host in a passthrough like fashion) to Maintenance mode (the controller takes over control to perform powerup, initialization, calibration, scrubbing, rebuild, etc.). Due to the fixed timing of a DDR3 interface, the *Idle detector* must insure that no host operations occur

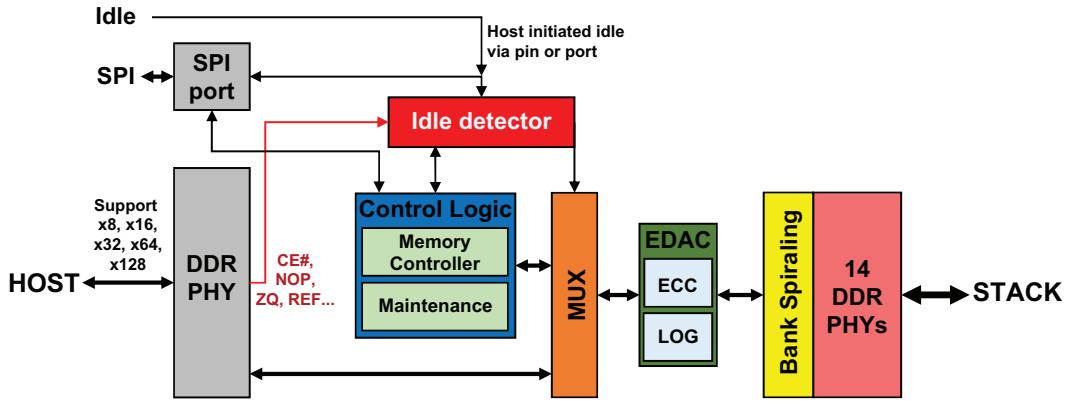


Figure 3. Controller architecture

within the duration of the operation, i.e. the future. This is a difficult task as host operations occur asynchronously from the standpoint of the command execution. For this reason, the memory cube requires some timing adjustments on the part of the host memory controller to allocate a few hundreds of nanoseconds extra on operations such as refresh (by adjusting tRFC to a higher value), ZQ, CE enable, or NOP commands to allow the memory cube to perform its internal management functions. For example, if the CE active timing is extended by 200ns then our controller would be guaranteed 200ns to clean up its execution on detection of CE going active — this would allow the module to perform internal functions whenever CE was inactive.

#### Cube Interface Considerations — PHY Layer

As the main interface between the logic tier and the memory cube, the PHY provides a high-speed electrical interface to the DRAM dies, integrating signal integrity features such as equalization, PLL, voltage and temperature compensated delays. The PHY also takes care of necessary features such as power-up/on-the-fly calibration, lane training, write Leveling, etc.

Because we want each die to be independent (for example for individual ON/OFF, calibration or rebuild), one PHY layer is implemented for each die. In order to account for the fly-by routing of traditional dual in-line memory modules (DIMMs), traditional DDR3 PHYs implement very advanced timing procedures, which add several clock cycles of latency 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 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 [11], in order to leverage an existing but simple FPGA PHY infrastructure that can be easily customizable to our needs.

#### Error Mitigation Features

A goal of the 3D 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 1. In the following subsections, we describe the details of some of the mitigation and detection techniques.

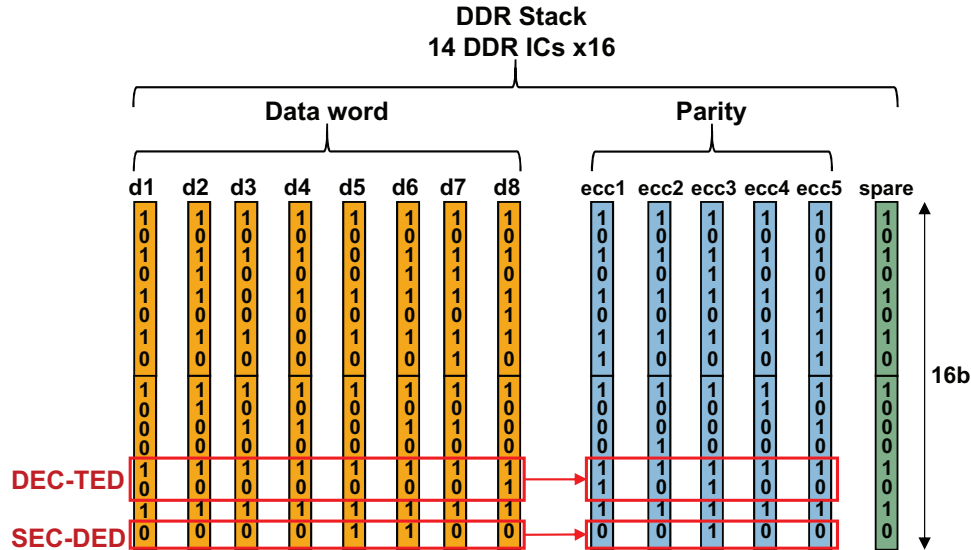
*EDAC Implementation*—We implement a byte wide Single Error Correction - Double Error Detection (SEC-DED) across 8 data devices (or similarly a DEC-TED across 2 rows, i.e. 16 bits data/10 bits ECC). This requires five additional DDR devices as shown in Figure 4. This method incurs a 20% power penalty compared to a more powerful shortened Reed Solomon RS(22,16) (i.e. can correct up to 3 symbols with 2 additional devices over RS code parity) 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 Figure 4, 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 that it can potentially correct a full die including nibble errors on one die if all the other dies are without 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 [12]. They present a high die cost but provide very high performance (encoding and decoding at clock speed is possible) which is important to reduce the overall latency of the controller.

SEUs can be easily mitigated by our SEC-DED scheme. On the other hand, detection of SEFIs is challenging. Indeed SEFI failures can affect the entire device preventing proper readout of data: temporary burst errors, to faulty rows or columns, or even corrupting an entire die which is equivalent to 16 bit errors in our case. If the SEC-DED provides good protection against SEFIs that cause device failures in the sense that it can provide full recovery of a failed device, this approach 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

**Table 1. Failure mode mitigation**

Failure Mode	Mitigation	Detection
SEU in memory array	EDAC	EDAC algorithm
SEU buildup	Scrubbing	Scrubbing over entire memory array
Leakage	Reduce refresh rate	By ground based testing of estimated leakage, plus programmable refresh rate combined with temperature sensor
SEFI — data failure	Power cycle and rebuild	EDAC errors exceeding programmable threshold level in same die
SEFI — high current	Power cycle and rebuild	Current draw anomaly at die
Device failure	Activate cold spare	Device fails BIST after power cycle



**Figure 4.** DDR stack device layout with SEC-DED code (8b data/5b ECC) or DEC-TED (16b data/ 10b ECC)

with the original data to determine failing bits and increment appropriate bit error count register. By tracking the level and frequency of high error symbols, we can determine if it is likely that a device has been affected by a 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. For this reason, it is critical that the rate of SEFIs be much lower than the rebuild time.

**Scrubbing**—As typical error rates are on the order of  $10^{-10}$  per bit-day [10], 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 non-intrusive, low-rate, background scrubbing (it is typical to perform scrubbing 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 that 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 as will remain tolerant to another SEU. 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 re-read. If data is still in error on re-read, 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.

**Current Monitoring**—As part of the power conditioning circuitry, we implement a current monitor that 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 such that current draw can be compared, i.e. 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.

**Built-In-Self-Test**—The BIST module is required to perform device initialization and zeroization with varying patterns to simplify the scrub process. To enhance the capabilities, the module also contains logic to read out and compare the data to provide full self-test capabilities. This module can also 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 powerup type tests to validate the

operation of the module. For extended test capability, we also implemented a powerful March X algorithm [13] that can test address decoding faults, stuck-at faults, transition faults and some coupling faults that can be triggered if necessary and if the cube has not to be readily available as this procedure has a high time complexity. The BIST is also designed to support pattern offsets for each die such that parallel testing can be performed where each die gets an offset.

*Rebuild*—The rebuild process consists of a single state machine controller that can be muxed into any of the DDR device channels. The rebuild process requires a read/write cycle for the entire array and it is estimated to require less than 5 seconds at maximum priority. The rebuild logic also contains various debug registers that can be used for ground-level test and characterization as well as in-flight diagnostics. To prevent SEU buildup the logic tier performs a continuous scrub of the 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. The rebuild time will be much lower than the SEFI rate. The logic tier supports power cycling and/or resetting of individual die that become affected by a SEFI failure. With a SEC-DED the array loses correction capabilities during rebuild. In either case, it is critical that the rebuild occur as quickly as possible.

*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 performs an address mixing such that data is stored in alternate banks among the dies. An additional “spiraling” at system level, as shown in Figure 1, which is done by alternating data dies with ECC dies in the cube can be executed to further expand this concept.

*Software Conditioning*—Software conditioning proposed in [9] 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 appear to dramatically improve with software conditioning. However, [9] showed that most device SEFIs could be removed by applying the *C1* nondestructive procedure. In order to clear certain SEFI failures, a full device power cycle may be required however [14]. The *C1* procedure consists of three operations normally performed during initialization only:

- rewriting the load mode registers of the DRAM die,
- resetting the internal DLL of the DRAM die, and
- re-performing the ZQ calibration.

Our controller can perform periodical dynamic software conditioning procedure when required, and independently between dies. In order to implement the procedure, the block in charge of initialization inside each PHY was modified so that to re-trigger the initialization steps and skip the unwanted calibration after initialization ended.

*Scheduling-based Mitigation Features*—An interesting observation by [9] 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 in order to close unused rows as soon as possible.

### *Verification*

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 operations. 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 (8Gb x16 DDR3 Verilog model from Micron) through PCB traces simulated by wire delay models. Simulations (see Figure 5) 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. In order to specifically test the interaction with the host, we developed a host emulator including an integrated memory controller and PRBS-based address/data pattern generators.

## 4. DESIGN FOR FLEXIBILITY

Our design methodology focused on allowing interchangeability of the memory device to migrate to different densities or different 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 3D-M<sup>3</sup> 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 re-routing of signal on the RDL can provide compatibility for a certain range of die sizes.

### *Controller Flexibility*

The controller has been designed in an open, modular architecture so that it can support multiple operating modes. The host data width (originally x128) can be reduced to target more classical widths (x32 or x64), at the expense of underutilizing the full capability of the cube. The controller can also be parameterized to target different commercial DRAM dies by simply changing various timing parameters (CAS latency, tREFI, etc.) and load mode register values. As an alternative to DRAM, DDR3 MRAM chips showed extremely high TID [15]. Thus, a STT-MRAM die can also be switched to, with slight modifications in the controller as changing timings, triggering anti-scribble during calibration, scrambling (precharge all) before power-down and disabling refreshes.

### *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 IO design and finally to the cube interface board design. The ability to readily modify the RDL design from one die type to another can be relatively difficult due to the different die sizes, IO pad spacings & pitch, and IO pad sizes. This is especially true when only a two metal RDL is used. Due to the need to fanout the RDL circuitry, adjacent die to the one modified die are sacrificed in order to achieve the necessary fanout. This is advantageous to attempt to utilize different memory dies to match a set IO pattern on the final RH controller. This will allow the use of single RH controller design that can accommodate different memory for the 3D memory cube.

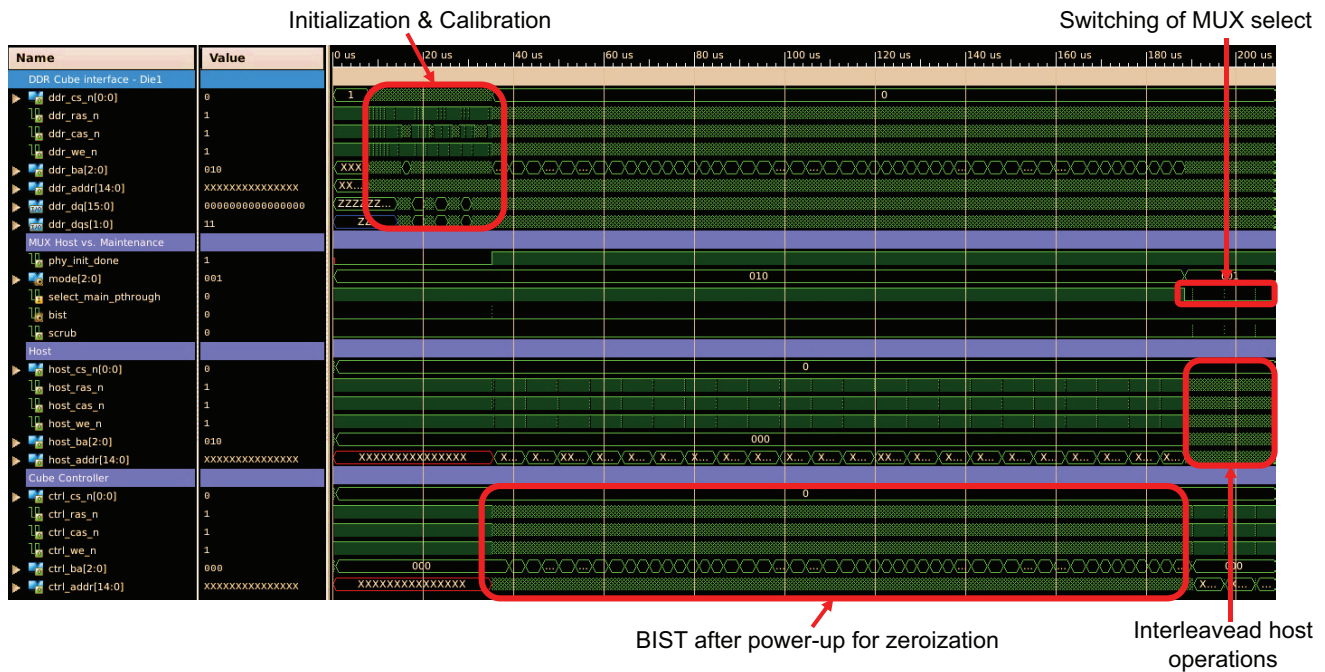


Figure 5. Simulation of the controller

Cost wise, this makes sense as a majority of the cost for the final 3D memory module is in the RH ASIC controller. In order to make this concept feasible, it is necessary to move to a four-metal layer RDL in order to effectively route out the necessary interconnects from the die to match the IO pattern on the RH controller. This would allow for two circuit layers, a power layer, and a ground layer. This also would 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 due to the fact that the RDL process is sequential, the chances of defect densities increase with increasing layer counts. This can be minimized by being able to relax 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.

## 5. FPGA TEST BOARD

To provide a breadboard validation of the design prior to fabrication of the final custom ASIC controller tier, we develop an FPGA-based breadboard. The FPGA emulating the function of the controller will connect to an actual 3D-M<sup>3</sup>.

### 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 IO requirements of nearly 1,000 pins. The high IO requirement is driven by the need of the controller to communicate with 14 memory dies as well as with the host memory controller. A Xilinx XC6VLX550T-1FF1760C FPGA with 1200 IOBs has been chosen as a target device. The high IO utilization constrains placement and results in potential routing delays that will reduce operating speed. Typical FPGA implementation re-

quirements for pin placement, such as allocating only center banks for address and control pins, cannot be respected due to the IO limitations. This is however an “FPGA only” type issue as the final ASIC design would overcome these limitations due to custom IO placement and area-array TSVs. Therefore, the FPGA design initial focus is on supporting minimum operating speeds of the DRAM dies for the first pass (300 MHz as baseline goal with DLL on).

Manual placement of the IOBs primitives (I/O-DDRs, IODE-LAYS, 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 Figure 3 on the Virtex-6, with the exception of the DDR host PHY which is currently in development stage. We achieved timing closure at 300 MHz. Fabric utilization of the FPGA resources is shown in Table 2. Figure 6 shows the mapping of the controller on the designated FPGA. The utilization is expected 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 Finite State Machine implementation.


### FPGA Test Platform

The block diagram of the test board is shown in Figure 7. In addition to integrating the FPGA-emulated controller and 3D-M<sup>3</sup>, the test board includes the following features:

- power supply current monitor for applicable rails,
- an external DDR3 interface for connection to host processor,
- connector-less logic probe pads for at least 1 DDR of the 3D-memory cube to enhance debug-ability,
- FPGA side channel IO for real-time performance monitoring and logging (e.g. a USB3.0 interface is used to perform memory dumps and allow for uploading

**Table 2. 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 %
IDELAYE1s	280	19 %



**Figure 6.** Floorplan of the FPGA-mapped controller design: 14 PHYs, Control Logic, and EDAC

particular memory patterns to simulate error conditions and faults),

- micro jumpers for at least one DDR device to allow “hard” error testing, and
- voltage switching for DDR layers.

As the FPGA fabric itself does not support device power switching, external FETs will be utilized under the control of the FPGA. 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 adequate operation margins. This task is actually more critical for the breadboard as the interconnect between memory cube and FPGA emulator is physically much further apart. This requires careful

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 ( $\leq 5$  mm from die pad).

#### Custom Substrate for Cube Mounting

The preliminary concept for integrating the stack onto the FPGA based logic tier board is shown in Figure 8. Low-profile spring-arrays (interposers) in the 4 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 and fan-out for the tight BGA pitch of the memory cube. Secure contact over the connector areas is made by pressure distribution compression rings.

The most critical area of the substrate is the cube BGA interface. The BGA pad size is estimated at  $150\mu\text{m}$  with  $150\mu\text{m}$  clearance. The pads will be non-solder mask defined to provide a well-controlled surface area for each ball. The stack uses a staggered BGA pattern which dictates that the IOs route out from one side (see Figure 9) for all the layers since the vertical direction is blocked by the BGA pattern itself. To maintain signal integrity (match impedance and electrical length), all DDR signal lines from one device are routed on the same layer in a stripline configuration. Due to the staggered pattern, the BGA escape pattern requires the use of blind vias and 7 signal layers to allow for each layer to route out unobstructed by adjacent layers.

#### Host Emulator

Given that no space processor is currently available to test and demonstrate the memory module, a host emulator has to be used for the final test setup in order 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 concept 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 has the benefit of providing 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.

#### Radiation Test Environment

In an effort to ultimately subject the 3D memory cube to radiation testing (both TID and SEE), we propose a test setup for radiation testing. For performance concerns, the interconnection between the 3D memory cube/interface substrate to the FPGA emulator requires the close proximity of the FPGA to the cube. However, this configuration does not lend itself to be subjected 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 Figure 10. The setup interconnects the two boards through a series of cables. It is speculated that there will be approxi-



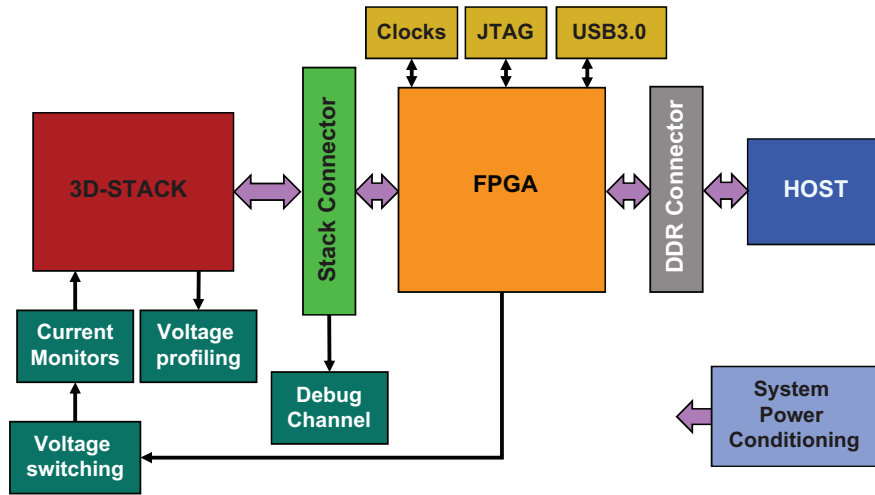


Figure 7. Test board architecture




Figure 8. Cube substrate to FPGA board connection concept




Figure 9. Cube substrate composite of signal layers

mately 1500 IOs that will need to be 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' to 20' range to adequately isolate any test electronics from the radiation source. For the SEE testing, shorter cables can be used as the radiation source is more pinpointed.

## 6. CONCLUSIONS

In this research 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 can avoid the use of proprietary TSV dies, achieve high memory capacity and good radiation tolerance. The controller has been tested in hardware and a complete test board structure is in development to subject the cube to radi-

ation sources and assess the performances of the controller. Our modular architecture allows flexibility in the choice of the memory die. Our cube development design methodology is intended to be a building block that provides a path to additional opportunities such as integration of non-volatile memory or computing resources.

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




Figure 10. FPGA radiation testing environment

## 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] Anthony Agnesina, Amanvir Sidana, James Yamaguchi, Christian Krutzik, John Carson, Jean Yang-Scharlotta, and Sung Kyu Lim. A Novel 3D DRAM Memory Cube Architecture for Space Applications. In *Proceedings of the 55th Annual Design Automation Conference, DAC '18*, pages 24:1–24:6, New York, NY, USA, 2018. ACM.
- [2] D. U. Lee, K. W. Kim, K. W. Kim, H. Kim, J. Y. Kim, Y. J. Park, J. H. Kim, D. S. Kim, H. B. Park, J. W. Shin, J. H. Cho, K. H. Kwon, M. J. Kim, J. Lee, K. W. Park, B. Chung, and S. Hong. 25.2 A 1.2V 8Gb 8-channel 128GB/s high-bandwidth memory (HBM) stacked DRAM with effective microbump I/O test methods using 29nm process and TSV. In *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pages 432–433, Feb 2014.
- [3] J. T. Pawlowski. Hybrid Memory Cube (HMC). In *2011 IEEE Hot Chips 23 Symposium (HCS)*, pages 1–24, Aug 2011.
- [4] Sung Kyu Lim. Bringing 3D ICs to Aerospace. *JICCE*, 12(2):117–122, June 2017.
- [5] P. M. Gouker, B. Tyrrell, M. Renzi, C. Chen, P. Wyatt, J. R. Ahlbin, S. Weeden-Wright, N. M. Atkinson, N. J. Gaspard, B. L. Bhuya, L. W. Massengill, E. Zhang, R. Schrimpf, R. A. Weller, M. P. King, and M. J. Gadlage. SET Characterization in Logic Circuits Fabricated in a 3DIC Technology. *IEEE Transactions on Nuclear Science*, 58(6):2555–2562, Dec 2011.
- [6] Yinghong Lin and David Joy. A new examination of secondary electron yield data. 37:895 – 900, 11 2005.
- [7] M. Park et al. Soft error study on DDR4 SDRAMs using a 480 MeV proton beam. In *2017 IEEE International Reliability Physics Symposium (IRPS)*, pages SE–3.1–SE–3.6, April 2017.
- [8] D. M. Hiemstra. Guide to the 2016 IEEE Radiation Effects Data Workshop Record. In *2017 IEEE Radiation Effects Data Workshop (REDW)*, pages 1–7, July 2017.
- [9] Martin Herrmann, Kai Grürmann, Fritz Gliem, Hagen Schmidt, Gilbert Leibelng, Heikki Kettunen, and Veronique Ferlet-Cavrois. New SEE Test Results for 4 Gbit DDR 3 SDRAM. 2012.
- [10] F. Gliem M. Herrmann, K. Grürmann and H. Schmidt. Radiation Tests of Candidate Memory Devices DDR3 SDRAM and NAND Flash, 2013.
- [11] Xilinx. UG086 Memory Interface Generator, User Guide v3.6, 2010.
- [12] V. Gherman, S. Evain, N. Seymour, and Y. Bonhomme. Generalized parity-check matrices for SEC-DED codes with fixed parity. In *2011 IEEE 17th International On-Line Testing Symposium*, pages 198–201, July 2011.
- [13] T. Koshy and C. S. Arun. Diagnostic data detection of faults in RAM using different march algorithms with BIST scheme. In *2016 International Conference on Emerging Technological Trends (ICETT)*, pages 1–6, Oct 2016.
- [14] K. Grürmann F. Gliem, D. Walter and M. Herrmann. Memory Technology Trends and Qualification Aspects, 2012.
- [15] J. Heidecker. MRAM Technology Status, 2013.

## BIOGRAPHY



**Anthony Agnesina** received the *Diplôme d'Ingénieur* from 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 pursuing the Ph.D. degree with the School of Electrical and Computer Engineering. His current research interests include 3D memory architectures, computer-aided design of VLSI circuits, and applied machine learning to electronics.



**James Yamaguchi** is VP of Irvine Sensors Corporation's 3D Electronics and Mass Storage group. He received his B.A. in Chemistry in 1977 and his B.S. in Chemical Engineering from California State University, Long Beach. He joined Irvine Sensors Corporation in 1993. Mr. Yamaguchi has over 39 years of hands-on processing experience in the areas of 3D packaging, thin film deposition,

electroplating, high density MCM fabrication and PWB fabrication. He has extensive R&D experience in fabrication technologies, process design and integration, facility operations and technology transfer. He is co-inventor on 14 patents and has co-authored 6 journal articles.



**Christian Krutzik** is a Senior Electrical Engineer at Irvine Sensors Corporation. He received his Masters Degree in Electrical and Computer Engineering from University of California, Irvine. Mr. Krutzik has over 18 years of experience involving electrical design and evaluation of 3D stacked modules and system miniaturization. Projects have included SSD design and development, compact

ITB 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. Mr. Krutzik is currently involved with the design and development of a secure SSD product line and has experience with developing and debugging SATA, PCIe, and other high speed interfaces. Mr. Krutzik 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. Further experience at Irvine Sensors 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. Mr. Krutzik also has experience with multiple electrical simulation tools such as Hyperlynx, HSPICE, and other FEA tools. Software experience includes C, Python, Matlab, Javascript, PHP, HTML, assembly, and Labview on both Linux and Windows platforms.



**John Carson** is the President & CEO and a founder of Irvine Sensors Corporation. Upon graduation from MIT in 1961, Mr. Carson joined Baird Atomics, Inc. in Cambridge and Waltham, Massachusetts, where he became Project Engineer and Assistant Program Manager on a space-based infrared surveillance program. He left Baird Atomics in 1967 to form a consulting company for sensor systems design and development, with clients that included North American Aviation, Honeywell, Lockheed, Ford Aeronautics, and Grumman. Over the years, this company evolved into Irvine Sensors Corporation. Mr. Carson 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. Mr. Carson 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 and has served on the Scientific Advisory Board for the Egg Factory, an incubation firm headquartered in Roanoke, Virginia.



**Jean Yang-Scharlotta (M07)** received her B.S. degree in chemical engineering from The University of Texas in Austin, in 1990 followed by M.S. and Ph.D. degrees in chemical engineering from Stanford University in 1996. She conducted engineering and research at the Los Alamos National Laboratory from 1987-1991 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 department manager at the Advanced Micro Devices focusing on 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 in the components assurance office at JPL. Dr. Yang-Scharlotta 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** received the B.S., M.S., and Ph.D. degrees from UCLA in 1994, 1997, and 2000, respectively. He joined the School of Electrical and Computer Engineering, Georgia Institute of Technology in 2001, where he is currently Dan Fielder Endowed Chair Professor. His current research interests include modeling, architecture, and electronic design automation (EDA) for 3D ICs.

His research on 3D IC reliability is featured as Research Highlight in the Communication of the ACM in 2014. Dr. Lim is 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 (2012) and the IEEE International Interconnect Technology Conference (2014). He has been an Associate Editor of the IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems since 2013.