# Design and Analysis of 3D-MAPS (3D Massively Parallel Processor with Stacked Memory)

Dae Hyun Kim, Krit Athikulwongse, Michael B. Healy, Mohammad M. Hossain, Moongon Jung, Ilya Khorosh, Gokul Kumar, Young-Joon Lee, Dean L. Lewis, Tzu-Wei Lin, Chang Liu, Shreepad Panth, Mohit Pathak, Minzhen Ren, Guanhao Shen, Taigon Song, Dong Hyuk Woo, Xin Zhao, Joungho Kim, Ho Choi, Gabriel H. Loh, Hsien-Hsin S. Lee, and Sung Kyu Lim

**Abstract—**This paper describes the architecture, design, analysis, and simulation and measurement results of the 3D-MAPS (3D massively parallel processor with stacked memory) chip built with a 1.5 V, 130 nm process technology and a two-tier 3D stacking technology using 1.2 $\mu$m-diameter, 6 $\mu$m-height through-silicon vias (TSVs) and 3.4 $\mu$m-diameter face-to-face bond pads. 3D-MAPS consists of a core tier containing 64 cores and a memory tier containing 64 memory blocks. Each core communicates with its dedicated 4KB SRAM block using face-to-face bond pads, which provide negligible data transfer delay between the core and the memory tiers. The maximum operating frequency is 277 MHz and the maximum memory bandwidth is 70.9 GB/s at 277 MHz. The peak measured memory bandwidth usage is 63.8 GB/s and the peak measured power is approximately 4 W based on eight parallel benchmarks.

**Index Terms—**3D Multiprocessor-memory stacked systems, 3D integrated circuits, Computer-aided design, RTL implementation and simulation

## 1 Introduction

Three-dimensional integrated circuits (3D ICs) are expected to provide numerous benefits. If a traditional two-dimensional (2D) IC is redesigned in multiple tiers vertically stacked in a 3D IC, the footprint area of the 3D IC reduces significantly. Because of this smaller form factor achieved by 3D stacking and 3D interconnections enabled by through-silicon vias (TSVs) and/or face-to-face (F2F) bond pads, 3D ICs have shorter wirelength than 2D ICs. Since the shorter interconnections improve the performance of the chip and reduce dynamic power consumption, 3D ICs have better performance and/or lower power consumption than 2D ICs. In addition, if the area of a 2D IC is very large, fabrication cost of its 3D implementation could be lower than that of the 2D IC [1], [2]. 3D ICs also enable heterogeneous integration by which circuit components built with different technologies can be integrated into a single 3D chip [3], [4].

As redesigning a single chip in multiple tiers provides many benefits listed above, stacking multiple chips, which are otherwise mounted and interconnected on a printed circuit board (PCB), in a 3D IC and connecting them using TSVs or F2F bond pads also provides the same kinds of benefits [5]–[9]. In particular, stacking multiple chips removes the limitation on the communication bandwidth among the chips, thereby enabling extremely high inter-chip (inter-tier) communication bandwidth. In addition, inter-chip signal transfer delay through TSVs or F2F bond pads is much shorter than that through PCB interconnects. Therefore, stacking multiple chips can resolve the memory bottleneck issue existing in the computer systems.

Although 3D ICs are expected to provide very high inter-tier communication bandwidth, the development and implementation of the architectures and applications that can fully exploit the wide bandwidth is still under research. An application expected to fully use the wide bandwidth is so called *wide-I/O* memory [10]. The communication bandwidth between a processor chip and a memory chip on a PCB is usually limited by the maximum number of the pins of the chips. If the processor and memory chips are stacked and connected through TSVs or F2F bond pads, however, the memory bandwidth can be increased and the memory access delay can be decreased dramatically. Since more and more applications such as image/video processing and large database management are demanding wider memory bandwidth, wide-I/O memory can fulfill the memory bandwidth requirement [5], [11]–[13].

In this paper, we present the architecture, design and analysis methodologies, and measurement results of our 3D-MAPS (3D massively parallel processor with stacked

- *D.H. Kim, K. Athikulwongse, M.B. Healy, M.M. Hossain, M. Jung, I. Khorosh, G. Kumar, Y.-J. Lee, D.L. Lewis, T.-W. Lin, C. Liu, S. Panth, M. Pathak, M. Ren, T. Song, D.H. Woo, X. Zhao, H.-H.S. Lee, and S.K. Lim are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332. E-mail: {daehyun, krit, mbhealy, mhossain7, moongon, ilyakhorosh, gokul.kumar, yjlee, dean, twlin, cliu, shreepad.panth, mohitp, mzren87, taigon.song, dhwoo, xinzhao, leehs}@gatech.edu; limsk@ece.gatech.edu.*
- *G. Shen and G.H. Loh are with the School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332. E-mail: gshen@gatech.edu, loh@cc.gatech.edu.*
- *J. Kim is with the Department of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon 305-701, Korea. E-mail: joungho@ee.kaist.ac.kr.*
- *H. Choi is with Amkor Technology Korea, Inc., Seoul 133-706, Korea. E-mail: Ho.Choi@amkor.co.kr.*

*Manuscript received 02 Sep. 2012; revised 19 Aug. 2013; accepted 17 Sep. 2013. Date of publication 30 Sep. 2013; date of current version 12 Dec. 2014. Recommended for acceptance by K. Ghose.*
*For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.*
*Digital Object Identifier no. 10.1109/TC.2013.192*

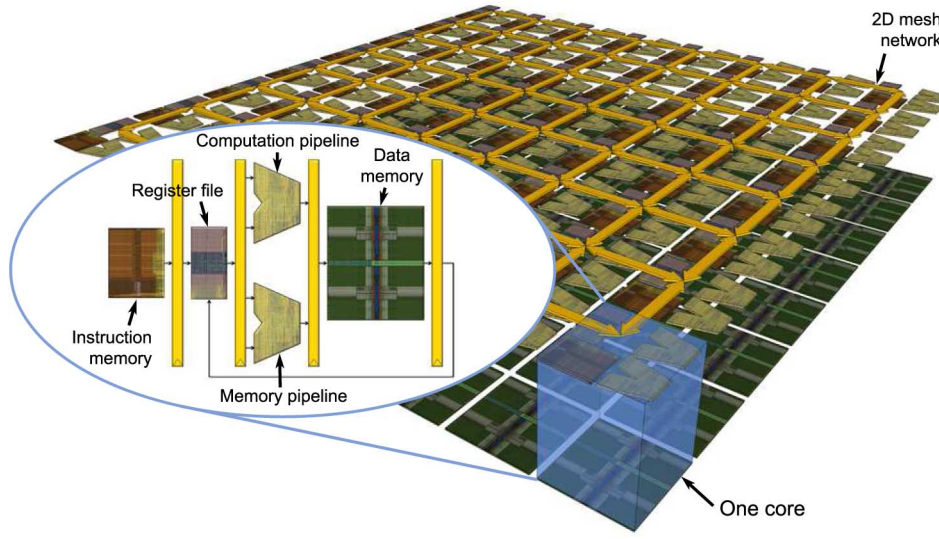


Fig. 1. 3D-MAPS architecture.

memory) chip, a fully-functioning multi-core processor with stacked memory, to demonstrate how core-memory 3D stacking can exploit very wide memory bandwidth in reality. 3D-MAPS is built with GlobalFoundries 130 nm process technology and Tezzaron two-tier 3D stacking technology that provides $1.2$ μm-diameter, $6$ μm-height TSVs for chip-to-package connections and $3.4$ μm-diameter face-to-face bond pads for inter-tier electrical connections. 3D-MAPS consists of two tiers, one (core tier) for processors and the other (memory tier) for memory. The core tier contains 64 cores in $8 \times 8$ array and the memory tier (256 KB SRAM) contains 64 4 KB memory tiles in the same $8 \times 8$ array, each of which is dedicated to its corresponding core at the same x- and y- location. The number of inter-tier connections for core-memory communication is 7,424, which is much greater than that of the wide-I/O single data rate standard (less than 800 connections) [10]. The memory operates at the same clock frequency (277 MHz) as the processors. We design the processor in such a way that we can run memory read/write operations *every clock cycle* to fully exploit the wide memory bandwidth. We also execute eight parallel benchmark applications to fully utilize the 64 cores and benefit from 3D stacking.

The contributions of this work are as follows: (1) This is the first work that presents a fully-functioning, general purpose, many-core 3D IC processor. This work presents measured memory bandwidth and power data based on fully verified parallel benchmark programs. (2) This work demonstrates how to exploit and extend the architectural features of a simple 2D processor to best exploit the highly parallel, single-cycle memory latency made possible in 3D stacked IC. (3) This work presents detailed descriptions on our design, analysis, and validation flows and techniques used to design 3D-MAPS. Since commercial or academic tools natively supporting 3D IC design and analysis do not exist, we demonstrate how to utilize and extend the features of the commercial tools designed for 2D ICs to handle tape-out quality of 3D IC designs.

This paper is organized as follows. We present the architecture design of 3D-MAPS in Section 2 and the 3D bonding and fabrication technologies used to manufacture 3D-MAPS in Section 3. In Section 4, we demonstrate the physical design flow and methodology used to design 3D-MAPS. Section 5 describes the benchmarks used to evaluate 3D-MAPS. We present the analysis methodologies and tools used to analyze 3D-MAPS in Section 6 and the package and board design in Section 7. In Section 8, we present measurement results. We perform a cost analysis for 3D-MAPS and discuss the applicability of our design methodology in Section 9, and summarize this paper in Section 10.

## 2 Architecture

In this section, we present the instruction set architecture, the single- and multi-core architectures, and the organization of the memory tier of 3D-MAPS. Fig. 1 shows an overview of the 3D-MAPS architecture.

### 2.1 Single-Core Architecture

The single-core architecture and the instruction set architecture (ISA) of 3D-MAPS are similar to those of the MIPS architecture, but we modify it to satisfy the system specifications and enable full utilization of the wide memory bandwidth. The most strict design specification was the area of a single core ($560$ μm by $560$ μm).[1] Therefore, we explored various design options such as pipeline depth, register file size, issue width, and arithmetic functions to support, and simplified the single core architecture and the ISA by excluding the components requiring large silicon area such as complex decoders, dynamic instruction schedulers, reorder buffers, branch predictors, floating-point units, and integer dividers.

Fig. 2 shows a simplified datapath of a single core of 3D-MAPS. Word size is 32-bit and pipeline depth is five. A single core has a 64-bit (two 32-bit instructions) 1.5 KB instruction memory, a 32-bit 4 KB data memory (in the memory tier), and a 32-bit dual-pump register file having 32 registers, three input (write) ports, and four output (read) ports. It also has a general-purpose ALU, a 16-bit multiplier, and four

1. The footprint area given to us from 2009 DARPA/Tezzaron 3D IC multi-project wafer run is $5$ mm $\times$ $5$ mm. The layout area excluding I/O area is almost $4.5$ mm $\times$ $4.5$ mm, so the maximum width of a single core is about 560um.

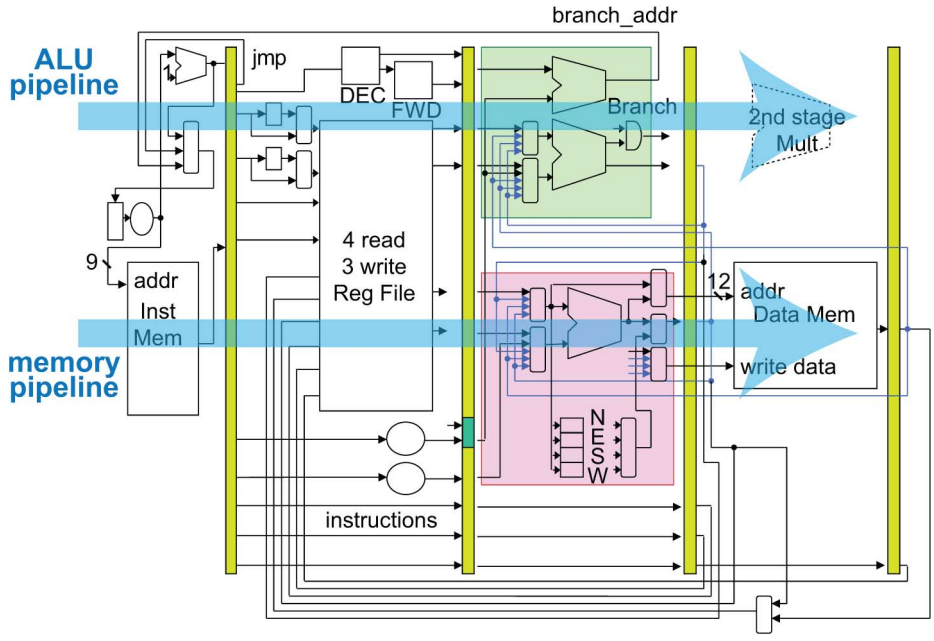


Fig. 2. Details of the single core architecture of 3D-MAPS.

(north/south/east/west) core-to-core communication ports, each of which has 33 input and 33 output bits (32-bit data and 1-bit control). Inter-core communication occurs in the third pipeline stage.

Issue width is two, so each 64-bit instruction bundle consists of two 32-bit instructions. We reserve one execution path for each operation type (memory/non-memory) so that each single core can access the memory every clock cycle, which maximizes the utilization of the wide memory bandwidth. However, our ISA also runs certain commonly-used non-memory instructions in the memory pipeline when the memory instruction is absent. In addition, the ISA also supports auto-increment to further increase memory bandwidth utilization by improving the ratio between the memory instruction and the non-memory instruction counts.

Since supporting out-of-order executions requires more complex logic, which occupies too large area to contain in the single core layout, we do not support them. Instead, we implement out-of-order executions at the software level by optimizing the instructions of the benchmarks. However, this is essentially very similar to running out-of-order executions using hardware supporting out-of-order executions (see Section 9 for more discussions).

### 2.2 Multi-Core Architecture

In order to reduce routing complexity, simplify the multi-core network, and minimize power consumed in the inter-core interconnections, we employ a point-to-point 2D mesh network controlled by explicit communication and synchronization instructions in the 3D-MAPS multi-core architecture. In particular, we choose the 2D mesh network instead of other alternatives such as the 2D folded torus network for the following reasons. Two of the most important design specifications in the design of 3D-MAPS are containing 64 cores and supporting the 3D-MAPS ISA. The width (or the height) of the layout of an optimized single core supporting this ISA is $560\ \mu m$, so the width of each routing channel between two adjacent rows of cores is approximately $10\ \mu m$. However, this channel is too narrow to provide sufficient routing resources for power/ground, clock, control signals, and additional 66 wires for core-to-core communication between the leftmost and the rightmost cores. Therefore, we choose the 2D mesh network for our multi-core architecture. To support our 2D mesh network topology, each core has buffers for sending (or receiving) data to (or from) its north, south, east, and west neighbors. To synchronize cores, we use an H-tree-shaped global barrier instruction.

As the core tier has 64 cores in $8 \times 8$ array, the memory tier also has 64 memory tiles in the same $8 \times 8$ array. Each memory tile is dedicated to its corresponding core placed at the same x- and y- location and communicates with it through 116 F2F bond pads. A memory tile consists of four banks placed in $2 \times 2$ array and each bank is an 8-bit 1 KB SRAM block. By controlling each bank separately, 3D-MAPS executes byte operations efficiently. Fig. 5 shows a single memory tile and the 116 F2F bond pads (represented by small dots) placed in the middle of the tile.

## 3 TSV and Stacking Technology

Table 1 shows the details of the 3D technology used to build 3D-MAPS. The device technology is based on 1P6M 130 nm process provided by Global Foundries. The supply voltage is 1.5 V. Tezzaron 3D technology stacks two dies using face-to-face (F2F) bonding. The thickness of the bottom die is $765\ \mu m$, but that of the top die is $12\ \mu m$ ($6\ \mu m$ for the metal layers and $6\ \mu m$ for the silicon substrate) due to thinning after bonding as illustrated in Fig. 3.[2] TSV resistance ($= 0.6\Omega$) is almost negligible, but TSV capacitance ($= 15fF$) is not.

2. This specific vertical stack-up was the choice made by Tezzaron, our chip manufacturer, and Amkor, our package manufacturer to optimize yield and reliability. We note that several other stacking and packaging options are possible. For example, the entire chip structure can be flipped and attached to the package using C4 bumps. However, this may complicate thermal issues in the logic tier. In addition, this manufacturing option was not available.

TABLE 1
Technology Specifications for 3D-MAPS

| | |
|---|---|
| 2D | |
| Process technology | Global Foundries 130nm |
| # metal layers | 6 |
| Supply voltage | 1.5V |
| 3D | |
| Die-to-die bonding | Face-to-face |
| # dies | 2 |
| TSV shape | Octagon |
| TSV type | Via-first |
| TSV diameter | 1.2um |
| TSV height | 6um |
| Minimum TSV landing pad width | 2.2um |
| Minimum keep-out zone spacing | 0.5um |
| Minimum TSV-to-TSV pitch | 2.5um |
| TSV resistance (R) | 0.6Ω |
| TSV capacitance (C) | 15fF |
| Face-to-face connection pad shape | Octagon |
| Face-to-face connection pad diameter | 3.4um |
| Face-to-face connection pad pitch | 5um |
| Face-to-face connection pad R & C | Negligible |

Via-first TSVs are inserted into the thin die. TSV diameter is $1.2\ \mu m$, and the height is $6\ \mu m$. The minimum Metal 1 TSV landing pad width is $2.2\ \mu m$. The minimum keep-out zone spacing that refers to the distance between the vertical surface of a TSV and a device is set to $0.5\ \mu m$. The minimum TSV-to-TSV pitch is $5\ \mu m$. The backside of the bottom ($= thick$) die is not used at all (it is attached to a dummy silicon substrate for packaging as illustrated in Fig. 3. The backside of the top die is used for wire bonding. Most of the TSVs in the top die are used for I/O, while others are used for dummy TSVs to satisfy the minimum TSV density rule.

Tezzaron provides F2F bond pads for the die-to-die communication. The Metal 6 layer is dedicated to F2F connections, so the actual number of metal layers that can be used for routing is five. The width of a F2F bond pad is $3.4\ \mu m$, and the pitch between two adjacent F2F bond pads is $5\ \mu m$. The resistance and capacitance of a F2F bond pad are negligible. The location of all F2F bond pads is aligned to a grid structure with $5\ \mu m$ pitch, and approximately one million ($= 1{,}000$ by 1,000 grid) F2F bond pad locations are available between the two dies. This allows very high degree of freedom in choosing F2F bond pads for signal and P/G routing.

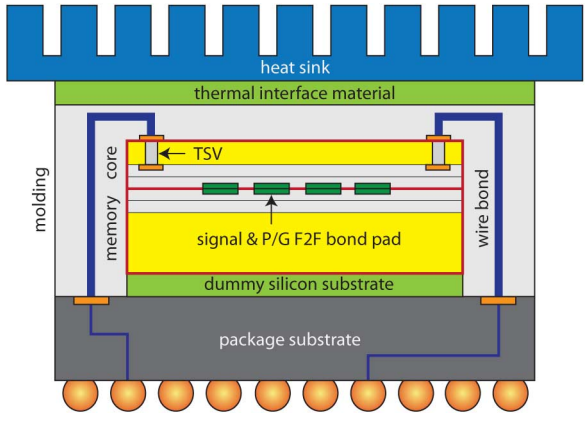


Fig. 3. 3D stacking, TSVs, face-to-face bond pads, and chip-to-package connections.

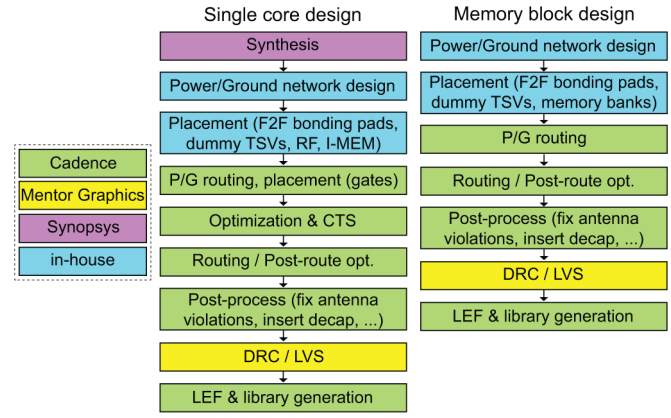


Fig. 4. 3D-MAPS physical design flow.

## 4 Physical Design of 3D-MAPS

### 4.1 Overview of 3D-MAPS Layout

The die footprint is $5\ mm \times 5\ mm$ and the area of a single core is $560\ m \times 560\ um$. The core-to-core spacing is $10\ \mu m$. Between a core and its memory tile, we have 668 F2F connections for power and ground (P/G). Therefore, the total number of F2F bond pads used for power delivery is 42,752 ($= 64 \times 668$). Since the resistivity of a single F2F bond pad is approximately $3m\Omega$, ignoring the contact resistance, the resistance through P/G F2F bond pads is almost negligible. We also have 116 F2F connections for signal between a core and its memory tile that include 32-bit data in/out, memory address, clock, and control signals. Therefore, the total number of F2F bond pads used for signal is 7,424 ($= 64 \times 116$). The longest path delay in the design is 3.61 ns at 1.5 V, which gives the maximum frequency of 277 MHz. 3D-MAPS contains about 50,000 TSVs and 50,000 F2F bond pads.

### 4.2 Single Core and Memory Tile Design

Fig. 4 shows our single core and memory tile design flow. Our 3D-MAPS design flow uses commercial tools from Cadence, Synopsys, and Mentor Graphics with our in-house tools to handle TSVs and 3D stacking. With the initial design constraints, the entire 3D netlist is synthesized by Design Compiler. The layout of each die is designed separately in SoC Encounter. The power distribution network is designed by our in-house tools. In addition, all the signal F2F pins are placed by our in-house tool, with proper alignments to the grid structure of Metal 6 bond pads. Clock tree synthesis is performed with proper boundary conditions at the clock F2F pins to the memory tier, followed by signal routing.

Fig. 5 shows the single core and single memory tile layouts. Each F2F bond pad in the single core layout is aligned with its corresponding F2F bond pad in the memory tile layout. A single core has 116 F2F signal and clock connections to its memory tile. The 116 core-to-memory connections consist of 32 data-out bits (from the core to its memory), 32 data-in bits (from the memory to the core), 40 address bits (10 bits per memory bank), eight control bits (two bits per memory bank), and four clock bits (one per memory bank). In the single core design, therefore, a primary 3D specific design task is constructing 3D routing topologies for the 116 signal and clock core-to-memory connections. The F2F bond pads located on

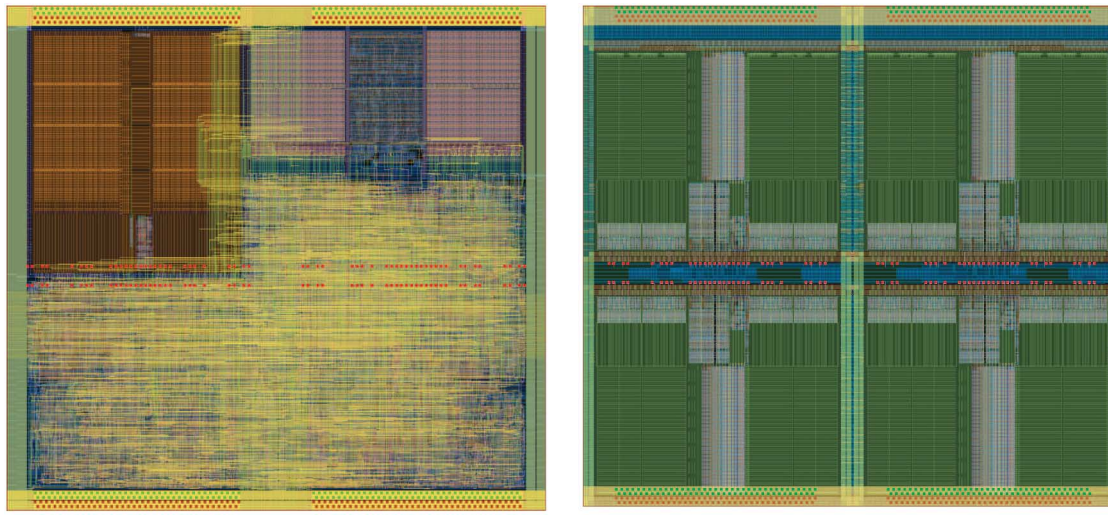


Fig. 5. Single core and memory tile layouts. Red squares in the middle are signal F2F bond pads, and the red/green squares on the top and the bottom are power/ground F2F bond pads.

Metal 6 are aligned to a $5\ \mu m$ pitch grid. Fig. 6 shows used and unused F2F bond pads.

To fully utilize existing commercial design tools, we pre-place F2F bond pads for all the core-to-memory connections as close to the pin locations of the memory blocks as possible. Placing F2F bond pads in this fashion enables us to resolve various design issues such as skew minimization in the 3D clock tree and timing optimization of 3D signal paths. For example, the longest wirelength between a F2F bond pad pin and its corresponding pin in the memory tier is approximately $10\ \mu m$, so the delay between the two pins is negligible. Therefore, we can perform clock skew minimization and timing optimization in the core tier only.

The next steps are similar to the traditional IC design flows: power/ground network design, placement, pre-clock tree synthesis (CTS) optimization, CTS, post-CTS optimization, routing, and post-routing optimization. For CTS, the F2F bond pad pins are defined as clock sinks. We also insert dummy TSVs in the single core design before standard cell placement to satisfy the TSV density design rule. Power delivery to the memory tier is done using P/G F2F bond pads placed on the top and bottom of each core as shown in Fig. 5.

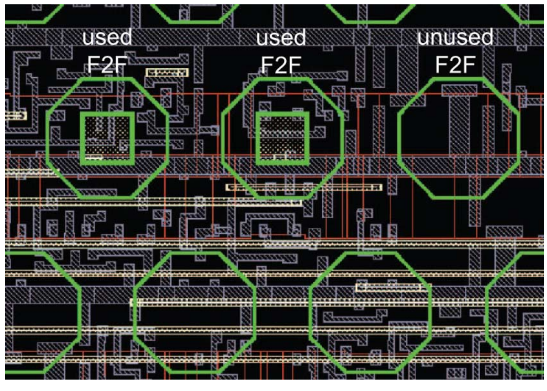


Fig. 6. Used and unused F2F bond pads. The pitch between two adjacent F2F bond pads is $5\ \mu m$.

### 4.3 Top-Level Design and Power Delivery Network

Fig. 7 shows the layout of the entire core and memory tiers. The design procedure for the core and memory tiers is straightforward. We first place I/O cells on the periphery of the core tier. In each I/O cell, we insert 204 *redundant* TSVs as shown in Fig. 8. This number is not chosen based on the current requirement, but mainly based on the area available. Next, all 64 cores and 64 memory tiles form $8 \times 8$ array in each tier. Since the F2F bond pad grid is pre-fixed, we shift each core and memory tile to align the F2F bond pads to the grid structure in Metal 6 layer.

The P/G ring in each core are connected with additional P/G wires that run in between the cores. These wires are connected to the P/G I/O cells. Clock routing is done at the full-chip level to connect the clock driver in the clock I/O cell to the clock entry point in each core with minimum skew.

## 5 Benchmarks

To demonstrate exploiting very wide memory bandwidth, we select eight memory-intensive benchmarks and parallelize them to assign evenly-distributed tasks to the 64 cores. Common methodologies in the selection, parallelization, and optimization of the benchmarks are as follows: 1) Core-to-core communication does not occur too often because it lowers the memory access frequency; 2) Boundary checking in the benchmarks is not complex; 3) Given data sets fit into the $8 \times 8$ core array. The eight benchmarks are as follows:

- AES Encryption: The input text is equally distributed to the 64 cores. Since the computation of the AES kernel occurs locally in each data block, this benchmark does not include core-to-core communication.
- Edge detection: The test image is split into an $8 \times 8$ array and assigned to the 64 cores. We use sobel operator for edge detection and avoid boundary processing problems by assigning to each core an image tile slightly larger than the original image tile.
- Histogram: We parallelize this benchmark by distributing the input data evenly to the 64 cores and accumulate

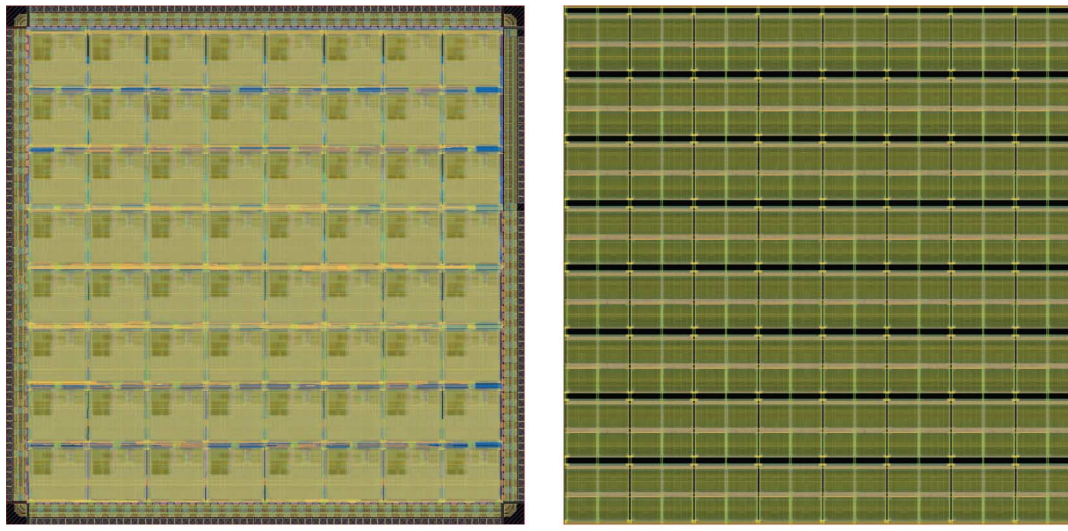


Fig. 7. Core die ($= 64$ cores) and memory die ($= 64$ SRAM tiles) layouts.

locally-computed histogram results across the cores to generate final results.

- K-means clustering: We use the algorithm presented in [14] to parallelize this benchmark. In particular, we calculate in parallel the distance between each data node and the center of the group to which the data node belongs. Then, we reassign each node to the group closest to it and go back to the distance computation step. We repeat the process until every node converges.
- Matrix multiplication: We use Cannon's algorithm [15] for distributed matrix multiplication. Each core works on a single element of the two source matrices at a time. After each element-wise multiplication, the matrix elements are systematically rotated among the neighboring cores so that every core receives a fresh pair of elements after each rotation. The final product is computed after a series of row-wise and column-wise rotations.
- Median filter: We apply $3 \times 3$ filter operations similar to [16] across a two-dimensional image. Each core processes a part of an input image split and assigned to it.
- Motion estimation: We use $16 \times 16$ macro blocks and find a motion vector minimizing the difference by shifting the macro blocks in a 2D image space. We parallelize this benchmark in a similar way to Median filter.
- String search: Input text is equally distributed to the 64 cores. Each core performs a search on the corresponding segment of the data. Upon completion of the local search, neighboring cores share their data to search the pattern in the overlapped region of the input text.

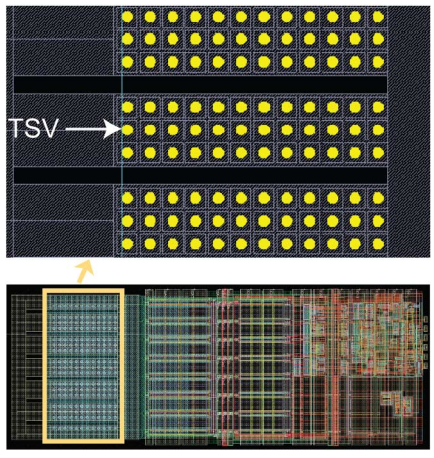


Fig. 8. Redundant TSVs inserted in an I/O pad.

## 6 Analysis of 3D-MAPS

In this section, we describe our strategy to extend existing commercial tools to analyze 3D-MAPS.[3]

### 6.1 Timing and Signal Integrity Analysis

Fig. 9 shows our flow for 3D static timing and signal integrity (SI) analysis. 3D SI-aware timing and SI analysis requires a netlist and a standard parasitic exchange format (SPEF) file for each tier. Therefore, we obtain the verilog netlists of the core and the memory tiers from Encounter and SPEF files of the tiers using QRC Extraction. We also create top-level netlist and SPEF files. The top-level netlist has two modules, one for the core tier and the other for the memory tier. We represent the face-to-face connections between the two tiers as wires connecting the two modules, and the TSVs between the core tier and the back side metal landing pads as wires connecting the core-tier module and the primary I/Os. The top-level SPEF file contains TSVs represented by the PI-model. We feed all the files to PrimeTime to perform SI-aware timing analysis. Similarly, we feed the same files to Cadence CeltIC to perform SI analysis.

Table 2 shows the timing-critical path of our single core, based on our 3D static timing analysis (STA) engine described above. The path is from a F/F of pipeline stage 2-3, through a MUX and an adder, to a F/F of pipeline stage 3-4. In the slack calculation, the STA engine uses clock arrival time at the begin

3. We do not perform on-chip thermal analysis mainly because our processor is low power and consumes up to 4 W as shown in Section 8. Our package-level solutions are enough to keep the processor low temperature using a simple air-cooled heatsink as shown in Fig. 14.

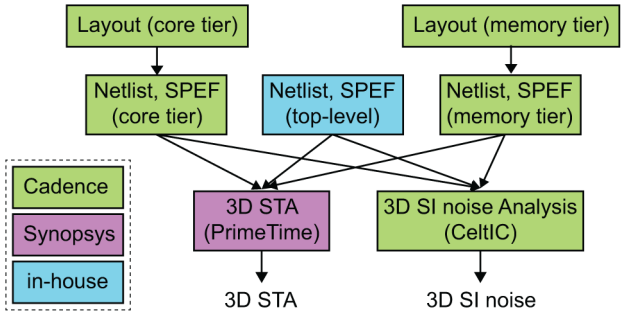


Fig. 9. Our 3D static timing and signal integrity analysis flow.

TABLE 2
The Critical Path in the Single Core. PS Denotes Pipeline Stage. Net Delay Is Negligible

| block | gates | delay (ns) |
|---|---|---|
| PS 2-3 | Scan D-F/F | 0.253 |
| MUX | AO22, AOI22, NAND2 | 0.668 |
| Adder | NOR2, NOR2, AOI21, OAI21 AOI21, BUF, XNOR2, XOR2, AO22 | 1.607 |
| PS 3-4 | Scan D-F/F | 0.487 |
| | slack | +0.019 |

point (1.270 ns) and the end point (0.802 ns) of the path. The path delay is 2.625 ns, excluding the setup time of the end point F/F. The maximum crosstalk noise value on the worst net is 674 mV, which is smaller than the noise limit of 750 mV.

In Tezzaron process technology, two tiers are bonded by face-to-face connections, so tier-to-tier capacitive coupling exists at the bonding interface. Therefore, capacitance extraction tools should consider the top metal layers in both tiers. However, we extract parasitic RC in each tier separately and ignore the tier-to-tier capacitive coupling. The main reason is that currently there is no commercial tool that can handle the extraction of this kind of parasitics. However, this is acceptable because we do not use Metal 5 in the memory tier, so the distance between the top surface of the Metal 5 layer in the core tier and that of the Metal 4 layer in the memory tier is approximately ten times greater than the distance between two adjacent metal layers.

### 6.2 Power and Power Supply Noise Analysis

Fig. 10 shows our 3D power analysis flow. We conduct 3D power consumption analysis as follows. We prepare netlists and standard delay format (SDF) files for both the core and the

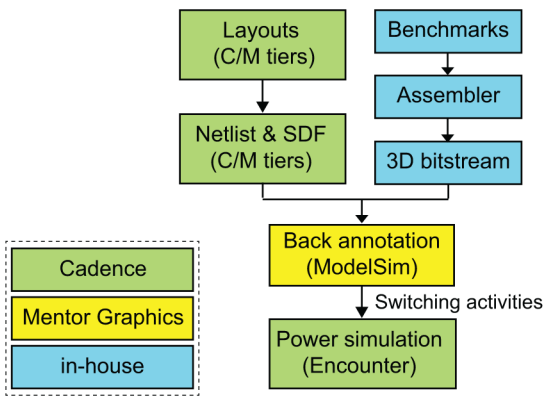


Fig. 10. Our 3D power analysis flow.

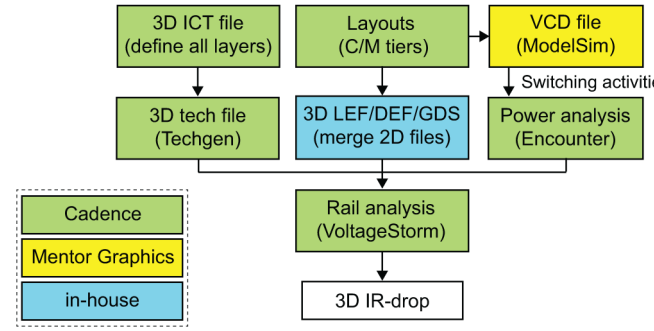


Fig. 11. Our 3D IR-drop analysis flow.

memory tiers. We also prepare bitstreams for each benchmark. With all these files, we run ModelSim to obtain switching activities of gates and nets. Finally, we perform power simulation using SoC Encounter with the netlist, SPEF, and switching activity files.

Fig. 11 shows our 3D IR-drop analysis flow. We perform 3D power noise analysis using Cadence VoltageStorm. The stand-alone VoltageStorm takes in a DEF file, technology files, and power dissipation files to generate both peak and average power noise values. For our 3D-MAPS design, we perform true 3D power noise analysis with VoltageStorm as follows. We create a 3D interconnect technology file (ICT) containing all the layers in the top and the bottom tiers and compile it using Cadence Techgen. We also construct a new LEF file containing instances specific to each tier. Then, we construct a new DEF file containing both the top and the bottom tiers using the new LEF file. Finally, we obtain 3D power noise values using VoltageStorm with the new technology, LEF, and DEF files as well as the power dissipation file obtained from our 3D power analysis.

Fig. 12 shows our 3D IR-drop analysis results, where we show the IR-drop map of a single core. The maximum IR drop occurs in the clock buffers, and the level is 60 mV. We also performed the full 64-core IR-drop analysis. The cores in the

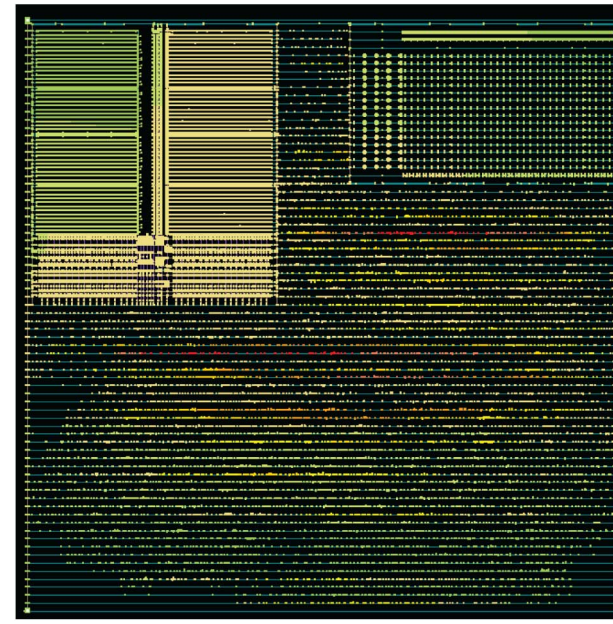

Fig. 12. IR-drop map of a single core. The maximum drop is 60 mV.

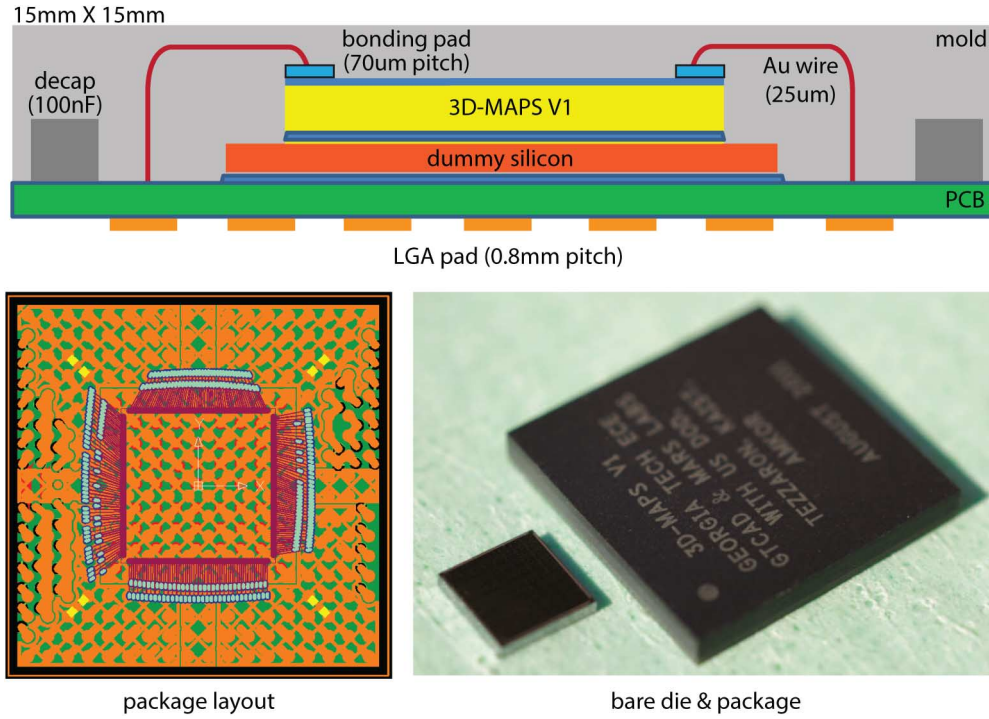


Fig. 13. Package design for 3D-MAPS.

middle of the layout experience the worst IR-drop of 78 mV, which is under the threshold of 150 mV.

### 6.3 DRC and LVS

Design rule checking (DRC) of 3D ICs consists of 2D and 3D DRC. 3D design rules include the minimum spacing between two adjacent TSVs, the minimum overlap between a TSV and its Metal 1 landing pad, the minimum TSV-to-device spacing, and so on. Since no inter-tier design rules exist, we run 2D and 3D DRC for each tier separately.

On the other hand, we perform 3D layout-versus-schematic (LVS) for both core and memory tiers simultaneously. Especially, the netlist extraction tool should be able to understand multi-tier layouts. For 3D LVS, we assign a unique number to each layer of the layouts. For example, we assign 32 and 42 to the poly layer and the active layer in the core tier, respectively, and 132 and 142 to the poly layer and the active layer in the memory tier, respectively. We also modify the extraction rule files to treat the same layers in the core tier and the memory tier separately. For example, the extraction rule includes both "a crossing of 32 and 42 forms a transistor" and "a crossing of 132 and 142 forms a transistor". Once a netlist is extracted from the given layouts, LVS is performed by a commercial LVS tool.

## 7 Package and Board Design

3D-MAPS is wire bonded to a four-layer, 0.8 mm-pitch land grid array (LGA) package as shown in Fig. 13. The LGA contains 324 land pads out of which 294 pads are used for power and ground to supply high current demand ($\sim 2.7$A) and 30 pads are used for clock and signal. The package is designed to accommodate high temperature ($\sim 90^{\circ}$C from our simulation) caused by high power density ($\sim 16$ W/cm$^2$). A dummy silicon substrate is inserted between the memory tier of 3D-MAPS and the package substrate to increase thermal conductivity from 3D-MAPS to the package. In addition, the center region of the LGA is implemented as a single large copper pad and dedicated to ground to decrease thermal resistance from the ground plane to the outside.

To verify the functionality of 3D-MAPS, we design a four-layer PCB test board having additional features. For example, to vary the power supply voltage from 0.9 V to 1.9 V, we add an additional power circuitry on the PCB. I/Os are connected to an FPGA test board (Xilinx Vertex-6) for verification. The average parasitic values of our signal package routes are $\mathrm{R} = 377.5\mathrm{m}\Omega$, $\mathrm{L} = 4.1\mathrm{nH}$, and $\mathrm{C} = 1\mathrm{pF}$. Fig. 13 shows the bare die of 3D-MAPS and its package.

To load data, execute benchmarks, read final data out, and verify the functionality of 3D-MAPS, we use scan chains and the test structure presented in [17]. Fig. 14 shows testing boards for 3D-MAPS. The scan cells are ordered into several scan chains that can be bypassed at the chain and core granularity. The scan chains are inserted and optimized using Design Compiler. The test architecture is composed of four

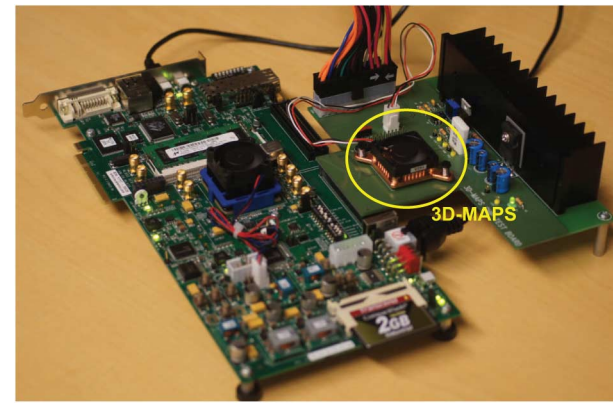


Fig. 14. Test board for 3D-MAPS.

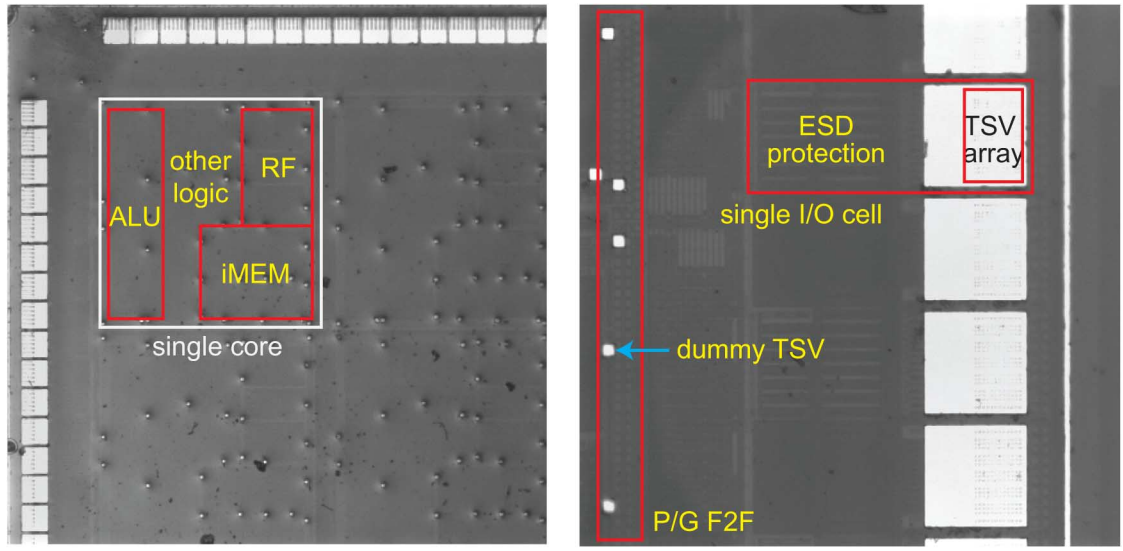


I/O cells & cores (IR image) I/O cells & ESD circuit (IR image)

Fig. 15. Die photographs.

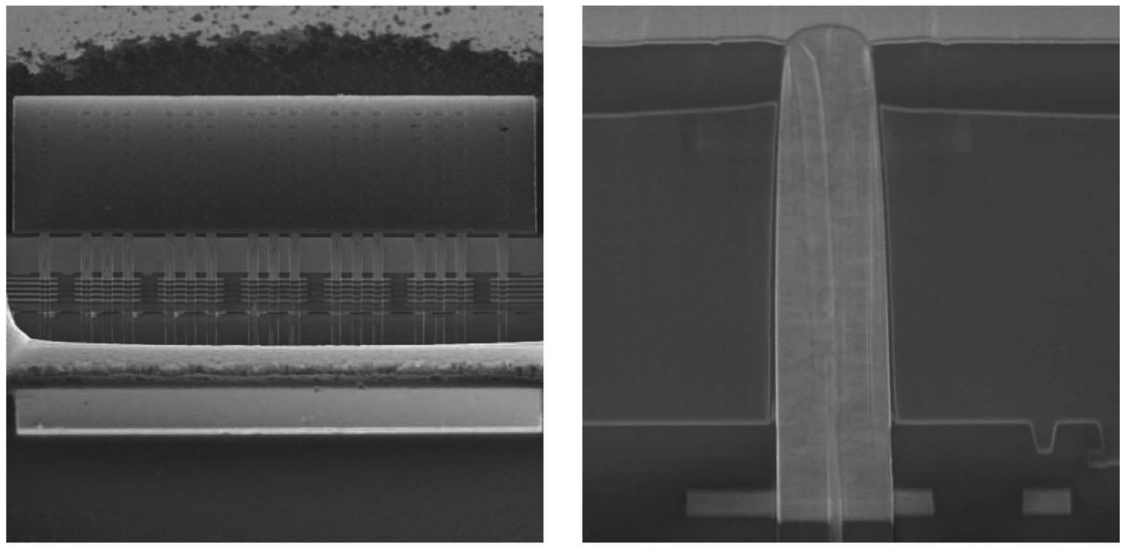

single IO cell, TSVs, BEOL of core die single TSV and its landing pads

Fig. 16. SEM image of Tezzaron TSVs and face-to-face bond pads.

independent test sectors that enable coarse-grained fault isolation; each sector can be tested and operated completely independently, up to and including the scan-in and scan-out package pins. A custom test controller is used to manage the configuration, test, and operation of the chip. 3D-MAPS contains 49,408 scan F/Fs (772 per core) and 16 chains (four per sector). The test system serves as the only access mechanism for this chip. An FPGA development kit serves as the external test driver delivering the driving bit-streams, managing the chip operation, and observing the results.

## 8 Die Shots and Measurement Results

Fig. 15 shows die photographs that contain logic blocks, TSVs, I/O cells, and F2F bond pads. Fig. 16 shows the details of TSVs and F2F bond pads between the two tiers. Fig. 17 shows an SEM image of TSVs in an I/O cell, front-side and back-side metal I/O pads, and P/G F2F bond pads.

Table 3 lists the peak memory bandwidth we measured in gigabytes per second (GB/s), instructions per cycle (IPC) per core, and measured power consumption of 3D-MAPS for each benchmark. The maximum memory bandwidth that we achieved is 63.8 GB/s (Median filter). The theoretical maximum memory bandwidth that 3D-MAPS can achieve is:

$$64 \times 4 \times (277 \times 10^6) = 70.912 \text{ GB/s} \quad (1)$$

where 64 is the total number of cores, 4 the word size (4 bytes), and 277 MHz the operating frequency. Therefore, the median filter benchmark uses up to 90% of the theoretical maximum memory bandwidth, whereas the string search benchmark uses 13%. The theoretical maximum memory bandwidth is higher than that of modern processors (e.g., Intel Core i7) whose maximum memory bandwidth is approximately 64 GB/s when DDR3-1333 memory is used. If we simply scale the clock frequency of 3D-MAPS up to 1333 MHz, the

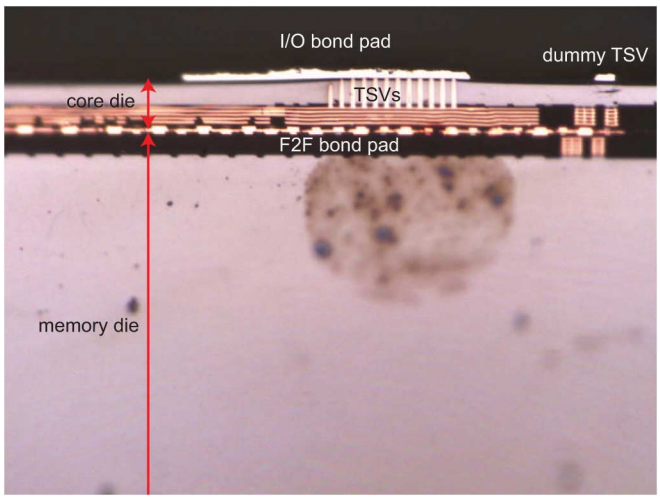


Fig. 17. SEM image of TSVs in an I/O cell, front-side and back-side metal I/O pads, and P/G F2F bond pads.

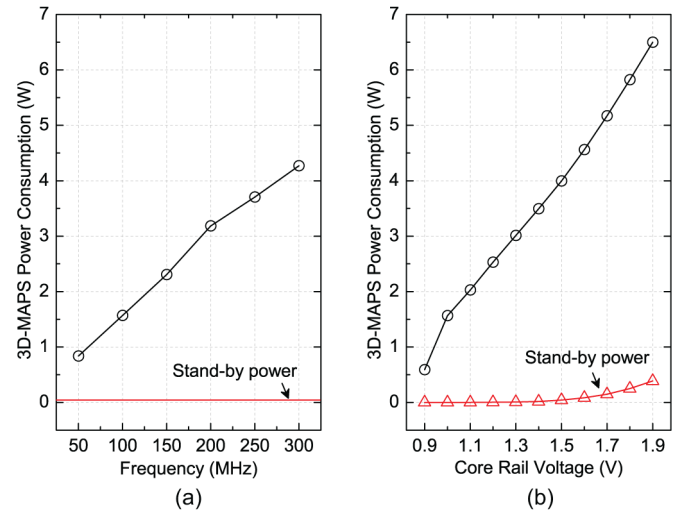


Fig. 18. (a) Frequency vs. power (at 1.5 V), (b) voltage vs. power (at 250 MHz) for the AES encryption benchmark.

maximum memory bandwidth becomes 341 GB/s, which is much higher than the modern memory bandwidth.

We measure power consumption of 3D-MAPS using a Watts-Up Pro power meter. The peak power consumption ranges from 3.5 W to 4.0 W as listed in Table 3. Although the median filter benchmark shows the highest memory bandwidth usage, the AES encryption benchmark has the highest peak power consumption because arithmetic and logic operations consume more power than memory operations. We also measure power consumption when the clock frequency and the core rail voltage vary. Fig. 18 shows power consumption of 3D-MAPS for the AES encryption benchmark when the clock frequency varies from 50 MHz to 277 MHz (at 1.5 V) and when the core rail voltage varies from 0.9 V to 1.9 V (at 250 MHz), respectively. The power consumption increases almost linearly with the frequency while the stand-by power is negligible. The core rail voltage vs. power graph also shows that the power consumption is almost linearly dependent on the supply voltage while the stand-by power consumption increases slightly. The only exception is at 0.9 V where the chip begins to suffer from near-threshold effects since the threshold voltage of our 130 nm PMOS is approximately 0.85 V.

## 9 Discussions

In this section, we present our analysis on the cost of 3D-MAPS designed in 2D and 3D. We also discuss the applicability of 3D-MAPS focusing on how we can extend 3D-MAPS when larger silicon area is given.

TABLE 3
Measured Memory Bandwidth, IPC Per Core, and Measured Power Consumption Results

| Benchmark | Bandwidth (GB/s) | IPC/core | Power (W) |
|---|---|---|---|
| AES encryption | 49.5 | 0.97 | 4.032 |
| Edge detection | 15.6 | 0.95 | 3.768 |
| Histogram | 30.3 | 0.90 | 3.588 |
| K-Means clustering | 40.6 | 0.94 | 4.014 |
| Matrix multiplication | 13.8 | 0.32 | 3.789 |
| Median filter | 63.8 | 1.62 | 4.007 |
| Motion estimation | 24.1 | 1.20 | 3.830 |
| String search | 8.9 | 0.65 | 3.876 |

### 9.1 Cost Analysis

In this section, we estimate and compare the cost of 2D-MAPS (2D version of 3D-MAPS) and 3D-MAPS chips using the cost analysis models presented in [18]–[20]. In our cost analysis, we assume that 2D-MAPS is laid out on a $10\ \text{mm} \times 5\ \text{mm}$ die and each memory block is placed right beside its corresponding core.[4] We also assume that die-to-wafer bonding is used and each die is tested before bonding, so only good dies are stacked. Therefore, the yield of a 3D-MAPS chip consisting of two dies is not $Y_c^2 \times Y_{\text{stacking}}$ but $Y_c \times Y_{\text{stacking}}$, where $Y_c$ and $Y_{\text{stacking}}$ are the yield of a single die and the stacking yield, respectively.

Based on the given design parameters and our assumptions in Table 4, the cost of a 2D-MAPS die $(= P_{\text{wafer}}/(N_c \cdot Y_c))$ is \$9.72 and that of a 3D-MAPS die is \$4.50. The total cost for building a 3D-MAPS chip is computed by summing the cost for two dies ($\$4.50 \times 2$), the TSV manufacturing cost for a logic die $(C_{\text{TSV}}/(N_c \cdot Y_c))$,[5] and the bonding cost ($C_{\text{bonding}}$). Assuming 100% stacking yield and almost zero stacking cost [19], a 3D-MAPS chip costs \$9.40 and a 2D-MAPS chip costs \$9.72, so a 3D-MAPS chip is cheaper than a 2D-MAPS chip by 3%. The reason that a 3D-MAPS chip is cheaper than a 2D-MAPS chip is because die yield is the most dominant factor for the cost of a 3D-MAPS chip.

If the defect density increases to 0.5, a 2D-MAPS chip and a 3D-MAPS chip cost \$11.35 and \$10.16, respectively, so the price gap goes up because the die yield is the dominant factor in the cost computation. On the other hand, if the TSV cost per wafer increases to \$400, a 3D-MAPS chip costs \$9.80, which becomes more expensive than a 2D-MAPS chip.

### 9.2 Applicability of 3D-MAPS

As shown in Section 2, we did not support advanced features such as out-of-order execution and branch prediction in the single core architecture of 3D-MAPS. Assuming we build a

4. We ignore performance degradation caused by increased distance between two horizontally adjacent (east-west) cores.

5. We add the TSV manufacturing cost of a core die only because a memory die has only few dummy TSVs.

TABLE 4
Variables and Constants Used in Our Cost Analysis for 3D-MAPS. Items Marked with Asterisks Are Based on Our Assumptions. $\alpha = 4^*$

| | Description | 2D | 3D |
|---|---|---|---|
| $D_{\text{wafer}}$ | Wafer diameter | 200mm | |
| $P_{\text{wafer}}$ | Single wafer price | \$5,000* | |
| $A_{\text{die}}$ | Die area (mm × mm) | 10×5 | 5×5 (per die) |
| $N_{\text{c}}$ | # Chips per wafer [18] | 565 | 1,168 (per wafer) |
| $D_0$ | Defect density ($/cm^2$) | 0.2* | |
| $Y_{\text{c}}$ | Yield $=(1+(AD_0/\alpha))^{-\alpha}$ | 0.91 | 0.95 (per die) |
| $C_{\text{die}}$ | Cost per die [19] | \$9.72 | \$4.50 (per die) |
| $Y_{\text{stacking}}$ | Stacking yield | - | 1.0* |
| $C_{\text{bonding}}$ | Bonding cost [19] | - | $C_{\text{stacking}}/Y_{\text{stacking}}$ |
| $Y_{\text{TSV}}$ | TSV yield | - | 0.9* |
| $C_{\text{TSV}}$ | TSV cost per wafer [20] | - | $200/$Y_{\text{TSV}}$ |
| $C_{\text{chip}}$ | 3D-MAPS chip cost | \$9.72 | \$9.00+\$0.40+$C_{\text{bonding}}$ |

homogeneous multi-core chip, if larger silicon area is given (e.g., $10\,\text{mm} \times 10\,\text{mm}$), we can either integrate more cores (e.g., 256 cores) or implement more advanced features in a single core. The former chip could increase the memory bandwidth utilization for memory-intensive, well-parallelized applications. Without optimization techniques such as out-of-order execution and branch prediction by compilers aware of the limitation of the simple single core architecture, however, each core might not be able to fully utilize the wide memory bandwidth. On the other hand, the performance of a single core in the latter chip will be increased by more various techniques such as hardware-based speculation, so the throughput of each core will increase. Therefore, the latter chip is likely to outperform the former chip for CPU-intensive applications. In addition, the advanced core might access its memory block more frequently than the simple core because of dynamic scheduling and branch prediction. However, since the total core count does not increase, the total memory bandwidth utilization of the latter chip might be lower than the former chip, in which more cores are integrated. Since predicting and finding optimal combinations of the number of cores and their functions need accurate, realistic simulations, we are also working on these topics as a follow-up research.

## 10 Conclusion

We presented the architecture, design flow and methodologies, analysis, and measurements of 3D-MAPS, 3D massively parallel processor with stacked memory, designed to exploit extremely high memory bandwidth obtained from core-memory stacking. The 3D stacking technology that we use supports numerous inter-die connections through face-to-face bonding pads, thereby enabling very high memory bandwidth in the core-memory stacking structure. 3D-MAPS consists of a core tier containing 64 identical cores in the $8 \times 8$ 2D mesh multi-core architecture and a memory tier containing 64 memory blocks, each of which is dedicated to its core in the core tier. To demonstrate exploiting very wide memory bandwidth that 3D stacking provides, we selected and parallelized eight memory-intensive benchmarks. The operating frequency of 3D-MAPS is 277 MHz and the maximum peak memory bandwidth utilization is 63.8 GB/s while consuming 4 W power.

## Acknowledgment

This research is based upon the work supported by the U.S. Department of Defense.

## References

[1] X. Dong and Y. Xie, "System-level cost analysis and design exploration for three-dimensional integrated circuits (3D ICs)," in *Proc. Asia South Pacific Des. Autom. Conf.*, Jan. 2009, pp. 234–241.
[2] J. Zhao, X. Dong, and Y. Xie, "Cost-aware three-dimensional (3D) many-core multiprocessor design," in *Proc. ACM Des. Autom. Conf.*, Jun. 2010, pp. 126–131.
[3] M. J. Wolf, P. Ramm, A. Klumpp, and H. Reichl, "Technologies for 3D wafer level heterogeneous integration," in *Proc. Symp. Des. Test Integr. Packag. MEMS/MOEMS*, Apr. 2008, pp. 123–126.
[4] K.-W. Lee, A. Noriki, K. Kiyoyama, T. Fukushima, T. Tanaka, and M. Koyanagi, "Three-dimensional hybrid integration technology of CMOS, MEMS, and photonics circuits for optoelectronic heterogeneous integrated systems," in *IEEE Trans. Electron Devices*, Mar. 2011, pp. 748–757.
[5] H. Saito, M. Nakajima, T. Okamoto, Y. Yamada, A. Ohuchi, N. Iguchi, T. Sakamoto, K. Yamaguchi, and M. Mizuno, "A chip-stacked memory for on-chip SRAM-rich SoCs and processors," in *IEEE J. Solid-State Circuits*, vol. 45, no. 1, Jan. 2010, pp. 15–22.
[6] U. Kang, H.-J. Chung, S. Heo, S.-H. Ahn, H. Lee, S.-H. Cha, J. Ahn, D. Kwon, J. H. Kim, J.-W. Lee, H.-S. Joo, W.-S. Kim, H.-K. Kim, E.-M. Lee, S.-R. Kim, K.-H. Ma, D.-H. Jang, N.-S. Kim, M.-S. Choi, S.-J. Oh, J.-B. Lee, T.-K. Jung, J.-H. Yoo, and C. Kim, "8Gb 3D DDR3 DRAM using through-silicon-via technology," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2009, pp. 130–131.
[7] Y. Kikuchi, M. Takahashi, T. Maeda, H. Hara, H. Arakida, H. Yamamoto, Y. Hagiwara, T. Fujita, M. Watanabe, T. Shimazawa, Y. Ohara, T. Miyamori, M. Hamada, M. Takahashi, and Y. Oowaki, "A 222 mW H.264 full-HD decoding application processor with x512b stacked DRAM in 40 nm," in *Proc. IEEE Int. Solid-State Circ. Conf.*, Feb. 2010, pp. 326–327.
[8] G. V. der Plas, P. Limaye, I. Loi, A. Mercha, H. Oprins, C. Torregiani, S. Thijs, D. Linten, M. Stucchi, G. Katti, D. Velenis, V. Cherman, B. Vandevelde, V. Simons, I. D. Wolf, R. Labie, D. Perry, S. Bronckers, N. Minas, M. Cupac, W. Ruythooren, J. V. Olmen, A. Phommahaxay, M. de Potter de ten Broeck, A. Opdebeeck, M. Rakowski, B. D. Wachter, M. Dehan, M. Nelis, R. Agarwal, A. Pullini, F. Angiolini, L. Benini, W. Dehaene, Y. Travaly, E. Beyne, and P. Marchal, "Design issues and considerations for low-cost 3D TSV IC technology," *IEEE J. Solid-State Circuits*, no. 1, pp. 293–307, Jan. 2011.
[9] J.-S. Kim, C. S. Oh, H. Lee, D. Lee, H. R. Hwang, S. Hwang, B. Na, J. Moon, J.-G. Kim, H. Park, J.-W. Ryu, K. Park, S. K. Kang, S.-Y. Kim, H. Kim, J.-M. Bang, H. Cho, M. Jang, C. Han, J.-B. Lee, J. S. Choi, and Y.-H. Jun, "A 1.2 V 12.8 GB/s 2 Gb mobile wide-I/O DRAM with $4 \times 128$ I/Os using TSV based stacking," *IEEE J. Solid-State Circuits*, no. 1, pp. 107–116, Jan. 2012.
[10] JEDEC. (Dec. 2011). *Wide I/O Single Data Rate* [Online]. Available: http://www.jedec.org
[11] C. C. Liu, I. Ganusov, M. Burtscher, and S. Tiwari, "Bridging the processor-memory performance gap with 3D IC technology," *IEEE Des. Test Comput.*, vol. 22, no. 6, pp. 556–564, Nov./Dec. 2005.
[12] G. H. Loh, "3D-stacked memory architectures for multi-core processors," in *Proc. IEEE Int. Symp. Comput. Archit.*, Jun. 2008, pp. 453–464.
[13] D. H. Woo, N. H. Seong, D. L. Lewis, and H.-H. S. Lee, "An optimized 3D-stacked memory architecture by exploiting excessive, high-density TSV bandwidth," in *Proc. IEEE Int. Symp. High-Perform. Comput. Archit.*, Jan. 2010.
[14] B. Hohlt, *Pthread Parallel K-means*. Berkeley, CA, USA: UC Berkeley, 2001.
[15] L. E. Cannon, "A cellular computer to implement the Kalman filter algorithm," Ph.D. dissertation, Montana State Univ., 1969.
[16] Xilinx, *Implementing Median Filters in XC4000E FPGAs* [Online]. Available: http://users.utcluj.ro/baruch/resources/Image/xl23_16.pdf
[17] D. Lewis, M. Healy, M. Hossain, T.-W. Lin, M. Pathak, H. Sane, S. K. Lim, G. Loh, and H.-H. S. Lee, "Design and test of 3D-MAPS, a 3D die-stack many-core processor," in *Proc. IEEE Int. Workshop Testing Three-Dimensional Stacked Integr. Circuits*, Nov. 2010, pp. 1–6.
[18] D. K. de Vries, "Investigation of gross die per wafer formula," *IEEE Trans. Semicond. Manuf.*, vol. 18, no. 1, pp. 136–139, Feb. 2005.

[19] C.-C. Chan, Y.-T. Yu, and I. H.-R. Jiang, "3DICE: 3D IC cost evaluation based on fast tier number estimation," in *Proc. Int. Symp. Quality Electron. Des.*, Mar. 2011, pp. 1–6.
[20] J. H. Lau, "TSV manufacturing yield and hidden costs for 3D IC integration," in *IEEE Electron. Compon. Technol. Conf.*, Jun. 2010, pp. 1031–1042.

**Dae Hyun Kim** received the BS degree in electrical engineering from Seoul National University, Seoul, Korea, in 2002 and the MS and PhD degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2007 and 2012, respectively. He is currently working on physical layout optimization at Cadence Design Systems, Inc. His research interests include electronic design automation and computer-aided design for VLSI, high-performance and/or low-power VLSI and computer systems, and 3D ICs.

**Krit Athikulwongse** received the BEng and MEng degrees from the Department of Electrical Engineering, Chulalongkorn University, Bangkok, Thailand, in 1995 and 1997, respectively, and the MS and PhD degrees from the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, in 2005 and 2012, respectively. Since 2012, he has been a Researcher at the National Electronics and Computer Technology Center, Khlong Luang, Pathum Thani, Thailand. His research interests include embedded systems, physical design, computer architecture, 3D ICs, and VLSI design.

**Michael B. Healy** received the PhD degree from the Georgia Institute of Technology, Georgia, in 2010. He is currently a post-doctoral researcher with the systems group at IBM Research. His work focuses on 3D integration and analyzing the new types of tradeoffs that it enables. Since joining IBM his work has included analyzing the thermal and power-supply tradeoffs of several 3D options for many of IBM's future flagship products, developing design automation tools targeted at early-stage design, and accurately modeling the performance of the memory hierarchy with future memory technologies.

**Mohammad M. Hossain** received the BS degree in computer science and Engineering at Bangladesh University of Engineering and Technology, Dhaka, Bangladesh, in 2005. He is currently working toward the PhD degree in College of Computing at Georgia Institute of Technology, Atlanta. His current research focuses on die-stacked DRAM organization, interference prediction on shared resource in cloud computing, and power and performance optimization in datacenters.

**Moongon Jung** received the BS degree in electrical engineering from Seoul National University, Seoul, Korea, in 2002, and the MS degree in electrical engineering from Stanford University, Stanford, California, in 2009. He is currently working toward the PhD degree in the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, Georgia. His research interests include computer-aided design for VLSI circuits, especially on physical design methods for low power 3D ICs and thermo-mechanical reliability analysis and optimization of TSV-based 3D ICs.

**Ilya Khorosh** is currently a hardware engineer with the Surface tablet group at Microsoft. He received the B.S in computer engineering from the Georgia Institute of Technology in 2013. During his undergraduate work he interned twice with AMD and once at Intel. At AMD he worked on product and test teams for 32 nm processors Llano and Orochi. At Intel he worked on power management validation for the many-core processor Knights Landing.

**Gokul Kumar** received the BE degree in electronics and communications engineering from the Anna University, Chennai, India, in 2007, and the MS degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2010. He is currently working toward the PhD degree with the 3-D Systems Packaging Research Center, the Georgia Institute of Technology. His current research interests include electrical modeling and design of interposers for 3-D system integration.

**Young-Joon Lee** received the MS and BS degrees from Seoul National University, Gwanak-gu, in 2007 and 2002, respectively. He is currently working toward PhD candidate in the School of Electrical and Computer Engineering at Georgia Institute of Technology. His research interests include monolithic 3D IC design automation, low-power design study of TSV-based 3D ICs, timing optimizations for TSV-based 3D ICs, and co-optimization of traditional metrics and reliability metrics on 3D ICs.

**Dean L. Lewis** received his PhD from the School of Electrical and Computer Engineering, Georgia Institute of Technology in 2012. He is currently an engineer with IBM's Microelectronics Division, where he continues to develop 3D stacked-IC technology. He is a member of Tau Beta Pi and the IEEE.

**Tzu-Wei Lin** received the B.S. degree in Electrical and Control Engineering from National Chiao Tung University, Taiwan. Currently, he is working toward the PhD degree in electrical and computer engineering at Georgia Institute of Technology. His research interests include memory system architecture and 3D IC.

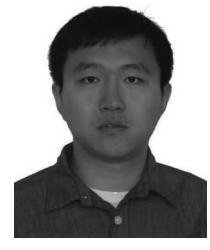

**Chang Liu** received the BS and MS degrees in microelectronics from Peking University, Beijing, China, in 2006 and 2009, respectively, and another MS degree in electrical and computer engineering from Georgia Institute of Technology, Atltanta, in 2011. Since 2012, he has been with Broadcom Corporation, working on multi-Gbps transceivers for optical and backplane/cable applications.

**Shreepad Panth** received the BE degree in electrical and electronics engineering from Anna University, Chennai, India, in 2009. He also received the MS degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2011, where he is working toward the PhD degree. His research interests include physical design and design for test of three dimensional integrated circuits, as well as physical design for monolithic 3D-ICs.

**Mohit Pathak** received the BTech degree in computer science and engineering from the Indian Institute of Technology, Kharagpur, India, in 2004, and the MS degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, USA in 2012. He is currently with Cadence Design Systems, Inc., San Jose, CA. His current research interests include physical design automation for stacked 3-D-integrated circuits, timing optimization, placement and routing algorithms, design for manufacturing techniques, and very large scale integration designs.

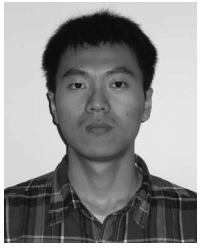

**Minzhen Ren** received the B.S. degree in electrical and computer eningeering from Georgia Institute of Technology, Atlanta, GA in 2011. He is currently working toward the master's degree in electrical engineering at the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta. His research interests include digital VLSI design and DSP algorithm applications for communication systems.

**Guanhao Shen** is a senior design engineer at Advanced Micro Devices. Guanhao received the M.S. degree in computer science from Georgia Institute of Technology, the B.S. degree in computer science from Nanjing Institute of Technology and a second B.E. degree in electronics engineering from Tsinghua University. He is working on architecture, performance modeling and correlation of memory controller and Northbridge for AMD products.

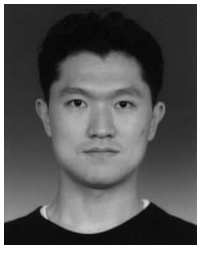

**Taigon Song** received a B.S. degree in Electrical Engineering at Yonsei University in Seoul, Korea, in 2007, and received an M.S. degree in Electrical Engineering at the Korea Advanced Institute of Science and Technology in Daejeon, Korea, in 2009. Currently, he is working toward his PhD degree in the School of Electrical and Computer Engineering at the Georgia Institute of Technology in Atlanta, Georgia. His research interests include silicon interposer design and co-analysis, TSV-to-TSV coupling in 3D ICs, and thermal analysis of 3D ICs with integrated voltage regulators.

**Dong Hyuk Woo** received the B.S. degree in electrical engineering from Seoul National University, Seoul, South Korea, in 2005 and the M.S. and PhD degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2007 and 2010, respectively. He is currently a computer architect at Intel. His current research interests include exascale computing, heterogeneous many-core architecture, general-purpose graphics processing unit, 3-D integration, and emerging memory technologies. He has co-authored a paper that was nominated for the Best Paper Award at HPEC-07 and a paper that was selected as IEEE Micro's top picks from the computer architecture conferences of 2010.

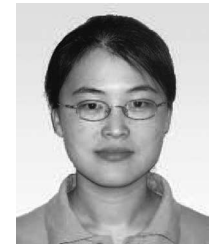

**Xin Zhao** received the BS degree in electronic engineering from Tsinghua University, in 2003, and the MS degree in computer science and technology from Tsinghua University, in 2006, and the PhD degree from the School of Electrical and Computer Engineering, Georgia Institute of Technology, in 2012. In 2013, she joined IBM as an Advisory Engineer and Scientist. Her research interests include circuit and physical design for 3D ICs, design and analysis on low-power circuits, reliability modeling and simulation, and timing analysis. She was the recipients of the Best Paper Award Nominations at the 2009 International Conference on Computer-Aided Design, the 2012 IEEE Transactions on CAD, and the 2012 International Symposium on Low Power Electronics and Design.

**Joungho Kim** received the BS and MS degrees in electrical engineering from Seoul National University, Seoul, Korea, in 1984 and 1986, respectively, and the PhD degree in electrical engineering from the University of Michigan, Ann Arbor, in 1993. In 1994, he joined Memory Division of Samsung Electronics, where he was engaged in Gbit-scale DRAM design. In 1996, he moved to KAIST (Korea Advanced Institute of Science and Technology). He is currently Professor at Electrical Engineering Department of KAIST. Also, he is director of 3DIC Research Center supported by Hynix Inc., and Smart Automotive Electronics Research Center supported by KET Inc. Since joining KAIST, his research centers on EMC modeling, design, and measurement methodologies of 3D IC, TSV, Interposer, System-in-Package (SiP), multi-layer PCB, and wireless power transfer (WPT) technology. He received Outstanding Academic Achievement Faculty Award of KAIST in 2006, Best Faculty Research Award of KAIST in 2008, National 100 Best Project Award in 2009, and KAIST International Collaboration Award in 2010, respectively. He is currently an associated editor of the *IEEE Transactions of Electromagnetic Compatibility*.

**Ho Choi** received the BS degree in electronics. Since 2000, he has been working as a design expert at Amkor Korea Design Center.

**Gabriel H. Loh** received the MS and PhD degrees in computer science from Yale University, and the BE in electrical engineering from the Cooper Union. He is a Fellow Design Engineer at Advanced Micro Devices (AMD). He was also an Associate Professor in the College of Computing at the Georgia Institute of Technology, a visiting researcher at Microsoft Research, and a senior researcher at Intel Corporation. His research interests include computer architecture, processor microarchitecture, emerging technologies and 3D die stacking. He is a senior member of IEEE and the ACM.

**Hsien-Hsin S. Lee** received the PhD degree in computer science and engineering from the University of Michigan, Ann Arbor. He is currenly on leave from the School of Electrical and Computer Engineering at Georgia Tech where he was an Associate Professor and leading the process design kit development at Taiwan Semiconductor Manufacturing Company, Taiwan. His research interests include computer architecture, 3-D IC, and energy-efficient datacenters. Prior to joining academia, he was a senior processor architect with Intel Corp. (1995–2001) and an architecture manager at the StarCore Technology Center of Agere Systems (2001–2002). His doctoral thesis received the Horace H. Rackham School Distinguished Dissertation Award at the University of Michigan. He also received the Department of Energy Early CAREER PI Award, the Georgia Tech ECE Outstanding Jr. Faculty Member Award, the NSF CAREER Award, and an IBM Faculty Award. He has co-authored four conference papers that received the Best Paper Award at MICRO-33, CASES-04, IBM PAC2, and ANCS-11, another four papers nominated for the Best Paper Award, and one paper selected in *IEEE MICRO Top Picks of Computer Architecture Conferences*. He serves on several editorial boards including the *ACM Transactions on Architecture and Code Optimization* (TACO), the *IEEE Transactions on Computer-Aided Design of Integrated Circuits and System (TCAD)*, and the *IEEE Transactions on Computers*. He holds 4 U.S. patents and is a member of Tau Beta Pi and a senior member of the ACM.

**Sung Kyu Lim** received the BS, MS, and PhD degrees from the Computer Science Department, University of California, Los Angeles (UCLA), in 1994, 1997, and 2000, respectively. From 2000 to 2001, he was a Post-Doctoral Scholar at UCLA, and a Senior Engineer at Aplus Design Technologies, Inc. He joined the School of Electrical and Computer Engineering, Georgia Institute of Technology in 2001, where he is currently a Professor. His research focus is on the physical design automation for 3-D ICs, 3-D System-in-Packages, microarchitectural physical planning, and eld-programmable analog arrays. He is the author of Practical Problems in VLSI Physical Design Automation (Springer, 2008). He received the Design Automation Conference (DAC) Graduate Scholarship in 2003 and the National Science Foundation Faculty Early Career Development (CAREER) Award in 2006. He was on the Advisory Board of the ACM Special Interest Group on Design Automation (SIGDA) during 2003–2008 and received the ACM SIGDA Distinguished Service Award in 2008. He is currently an Associate Editor of the *IEEE Transactions On Very Large Scale Integration Systems (TVLSI)* and served as a Guest Editor for the *ACM Transactions on Design Automation of Electronic Systems* (TODAES). He has served the Technical Program Committee of several ACM and IEEE conferences on electronic design automation.