

10.6 3D-MAPS: 3D Massively Parallel Processor with Stacked Memory

Dae Hyun Kim¹, Krit Athikulwongse¹, Michael Healy¹, Mohammad Hossain¹, Moongon Jung¹, Ilya Khorosh¹, Gokul Kumar¹, Young-Joon Lee¹, Dean Lewis¹, Tzu-Wei Lin¹, Chang Liu¹, Shreepad Panth¹, Mohit Pathak¹, Minzhen Ren¹, Guanhao Shen¹, Taigon Song¹, Dong Hyuk Woo¹, Xin Zhao¹, Joungho Kim², Ho Choi³, Gabriel Loh¹, Hsien-Hsin Lee¹, Sung Kyu Lim¹

¹Georgia Institute of Technology, Atlanta, GA

²KAIST, Daejeon, Korea

³Amkor Technology, Seoul, Korea

Several recent works have demonstrated the benefits of through-silicon-via (TSV) based 3D integration [1-4], but none of them involves a fully functioning multicore processor and memory stacking. 3D-MAPS (3D Massively Parallel Processor with Stacked Memory) is a two-tier 3D IC, where the logic die consists of 64 general-purpose processor cores running at 277MHz, and the memory die contains 256KB SRAM (see Fig. 10.6.1). Fabrication is done using 130nm GlobalFoundries device technology and Tezzaron TSV and bonding technology. Packaging is done by Amkor. This processor contains 33M transistors, 50K TSVs, and 50K face-to-face connections in 5x5mm² footprint. The chip runs at 1.5V and consumes up to 4W, resulting in 16W/cm² power density. The core architecture is developed from scratch to benefit from single-cycle access to SRAM.

3D-MAPS consists of 64 custom processor cores, each paired up with 4KB of SRAM scratchpad data memory. Each processor core is a five-stage pipelined, two-way VLIW machine that supports total of 42 instructions. Our core consists of one computation pipeline and one data memory pipeline, designed to sustain one 4B data memory operation per cycle (see Fig. 10.6.1). The goal of our architectural design is to demonstrate the extreme bandwidth between the data memory and the processing cores using high-density 3D interconnect. The theoretical maximum achievable data memory bandwidth running at 277MHz is 70.9GB/s. The core-to-core communication is achieved via dedicated links and buffers by executing explicit point-to-point message-passing instructions within a 2D mesh network. Such a simple 2D mesh eliminates long wires in other alternative interconnect networks, thereby reducing power, routing congestion, and area. Synchronization among all cores is performed using a global barrier instruction that collapses and broadcasts the final reduced result with a hierarchical H-tree on the core layer.

Tezzaron 3D technology stacks two logic dies using face-to-face (F2F) bonding, where the top die is thinned to 12µm and the bottom die is 765µm thick (see Fig. 10.6.2). These F2F pads are used for signal and P/G connections between the two dies. The diameter of a F2F bonding pad is 3.4µm, and their pitch is 5µm. 3D-MAPS uses 235 I/O cells that are placed along the periphery of the core die. Each I/O cell contains 204 redundant TSVs, where each TSV connects between a metal 1 landing pad and a backside metal landing pad deposited on the backside of the silicon substrate (see Figs. 10.6.3 and 10.6.7). Each backside metal landing pad (56x56µm²) is wire bonded to the packaging substrate. The diameter, height, and pitch of a TSV are 1.2µm, 6µm, and 5µm, respectively (see Fig. 10.6.3). The TSV RC parasitics are measured to be 0.6Ω and 3fF. Tezzaron imposes a minimum TSV density rule, where there must be at least one TSV in every 250x250µm² window. In addition to the TSVs used in I/O cells, we manually insert dummy TSVs in the core region to satisfy the rule over the entire chip area. Figure 10.6.7 shows the 35 dummy TSVs used in each core. The entire memory die contains around 4500 dummy TSVs.

Our RTL-to-GDSII tool chain is based on commercial tools and enhanced with our in-house tools to handle TSVs and 3D stacking. With initial design constraints, the entire 3D netlist is synthesized by Design Compiler. The layout of each die is done separately in Encounter. The power distribution network is designed by our in-house tools. In addition, all signal F2F pads are placed by our in-house tool with proper alignment to the grid structure of metal 6 bond pads. Clock tree synthesis is performed with proper boundary conditions at the clock F2F pads to the memory die, followed by signal routing. Metal 6 is dedicated to F2F connections. The netlists and the extracted parasitic files are fed into

PrimeTime for 3D static timing analysis, followed by the timing optimization in Encounter with the timing constraints from the 3D timing results. CeltIC is used for signal-integrity-aware timing optimization. After timing optimization, filler cells, decap cells, and metal fills are inserted to meet the design specifications. With the post-layout netlist and timing annotations, functional verification is performed with ModelSim. We use an in-house tool and Calibre for 3D LVS and DRC. For IR-drop analysis, we generated a 3D interconnect technology file and ran VoltageStorm based on the stimulus vectors of our benchmarks. We use GAMBIT, FLUENT, and in-house tools for material characterization, thermal gridding, and temperature calculations. Figure 10.6.4 shows the core and memory tile layouts. Details are provided in [5].

3D-MAPS is wire bonded to a four-layer, 0.8mm-pitch land grid array (LGA) package (see Fig. 10.6.2). Out of 324 land pads in the LGA, 294 are used for P/G to supply high current (~2.7A). The package is designed to accommodate high power density and hence high temperature (~90°C from simulation). A dummy silicon substrate is inserted in between 3D-MAPS and the package substrate to increase thermal conductivity. In addition, the center region of the LGA is implemented as a single large copper pad and dedicated to VSS to decrease thermal resistance from the VSS plane to outside. A four-layer PCB test board is designed to verify the functionality of the chip. Power circuitry is added to the board to vary V_{DD} from 0.9V to 1.9V. Additional decoupling capacitors are added to improve power integrity. I/Os are connected to an FPGA test board (Xilinx Virtex-6) for verification. The average parasitic values of our signal package routes are R = 377.5mΩ, L = 4.1nH, and C = 1pF.

The scan cells are ordered into several scan chains that can be bypassed at the chain and core granularities. The scan chains are inserted and optimized with Design Compiler. The test architecture is composed of four 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 derived from the IEEE 1149.1 standard is used to manage the configuration, test, and operation of the chip. 3D-MAPS contains 49,408 scan FFs (772 per core) and 16 chains (4 per sector). As there is no traditional off-chip memory interface, the test system serves as the only access mechanism for this chip. An FPGA development kit serves as the external test driver, delivering the bit streams, managing the chip operation, and observing the results.

We wrote and validated eight parallel computing applications (see Fig. 10.6.5). One of our benchmarks, median filter, demonstrates 63.8GB/s memory bandwidth (out of a potential 70.9GB/s). Our power measurement is made using a Watts-Up Pro power meter. Our measured peak power consumption ranges from 3.5W to 4.0W. Power consumption is also reported under various frequency and core voltage values (see Fig. 10.6.6) using the AES encryption benchmark. Our voltage vs. power study shows that the stand-by power consumption is slightly dependant on the supply voltage and reached up to 0.3W. The power consumption generally grows linearly with frequency; the exception is at 300MHz, where we have exceeded the target frequency and begun to experience timing faults. The power consumption grows super-linearly with core voltage; the exception is at 0.9V, where the chip begins to suffer from near-threshold effects. The threshold voltage of our 130nm PMOS devices is around 0.85V.

Acknowledgments:

This work is funded by the US Department of Defense.

References:

- [1] M. Koyanagi, et al. "Neuromorphic Vision Chip Fabricated Using Three-Dimensional Integration Technology", *ISSCC Dig. Tech. Papers*, pp. 270-271, 2001.
- [2] U. Kang, et al. "8Gb 3D DDR3 DRAM Using Through-Silicon-Via Technology", *ISSCC Dig. Tech. Papers*, pp. 130-131, 2009.
- [3] G. Van der Plas, et al., "Design Issues and Considerations for Low-Cost 3D TSV IC Technology", *ISSCC Dig. Tech. Papers*, pp. 148-149, 2010.
- [4] J.-S. Kim, et al. "A 1.2V 12.8GB/s 2Gb Mobile Wide-I/O DRAM with 4x128 I/Os Using TSV-Based Stacking", *ISSCC Dig. Tech. Papers*, pp. 496-498, 2011.
- [5] M. Healy, et al. "Design and Analysis of 3D-MAPS: A Many-Core 3D Processor with Stacked Memory", *IEEE Custom Integrated Circuits Conf.*, pp. 1-4, 2010.

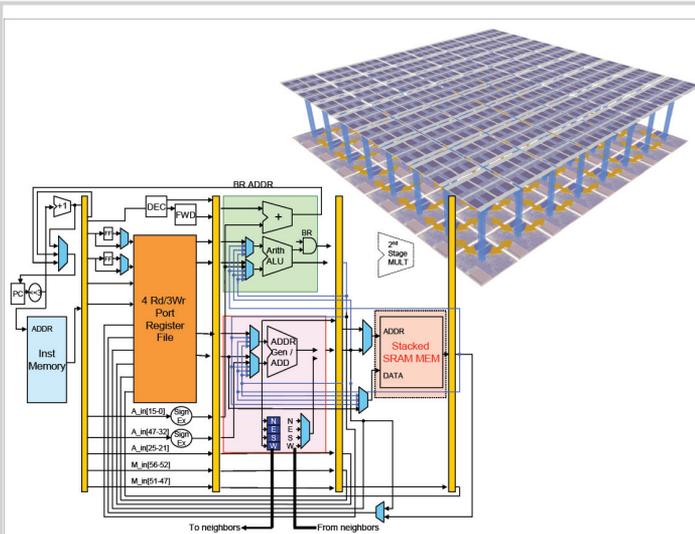


Figure 10.6.1: 3D-MAPS overview and single core architecture.

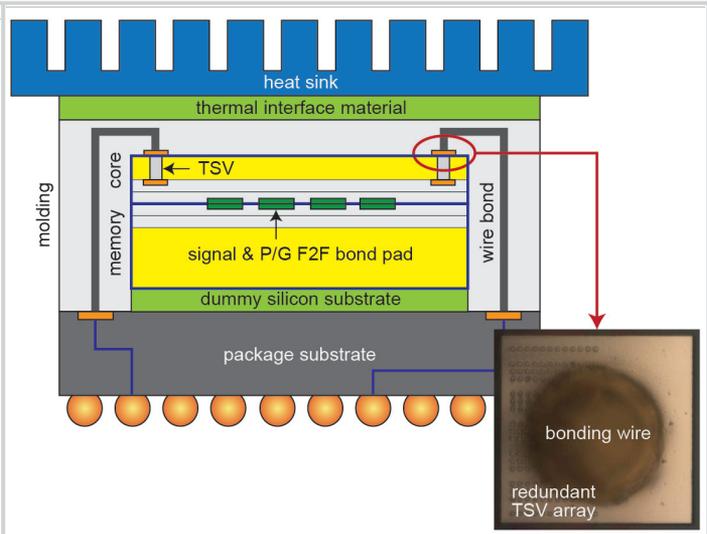


Figure 10.6.2: Stacking, TSV, F2F, and chip-to-package connection.

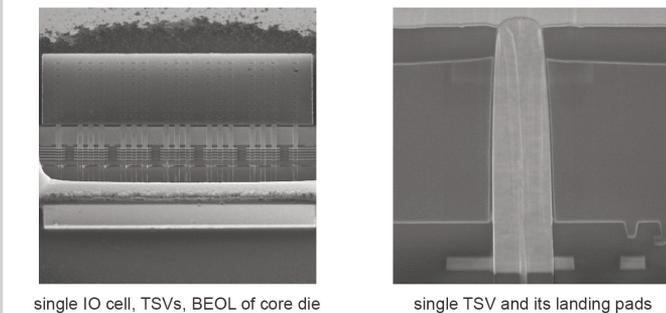
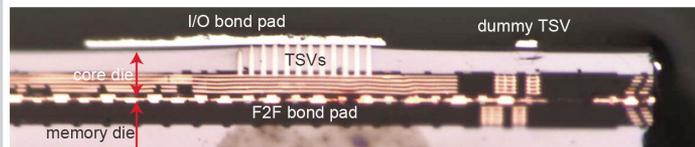


Figure 10.6.3: SEM images of TSVs and face-to-face connections.

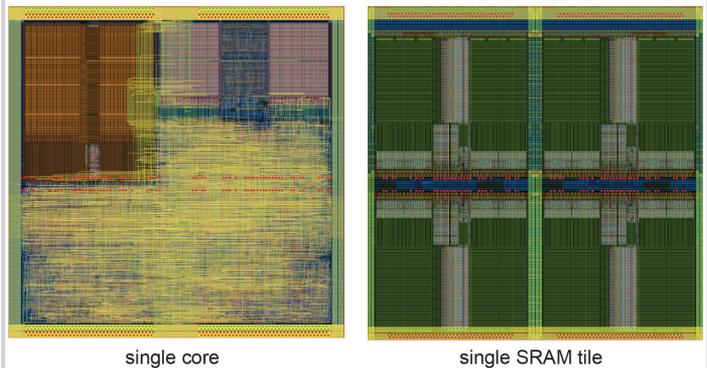


Figure 10.6.4: Single core and memory tile layouts.

benchmark	memory bandwidth	measured power
AES encryption	49.5 GB/s	4.032 W
edge detection	15.6 GB/s	3.768 W
histogram	30.3 GB/s	3.588 W
k-means clustering	40.6 GB/s	4.014 W
matrix multiplication	13.8 GB/s	3.789 W
median filter	63.8 GB/s	4.007 W
motion estimation	24.1 GB/s	3.830 W
string search	8.9 GB/s	3.876 W

Figure 10.6.5: Bandwidth and power measurement results.

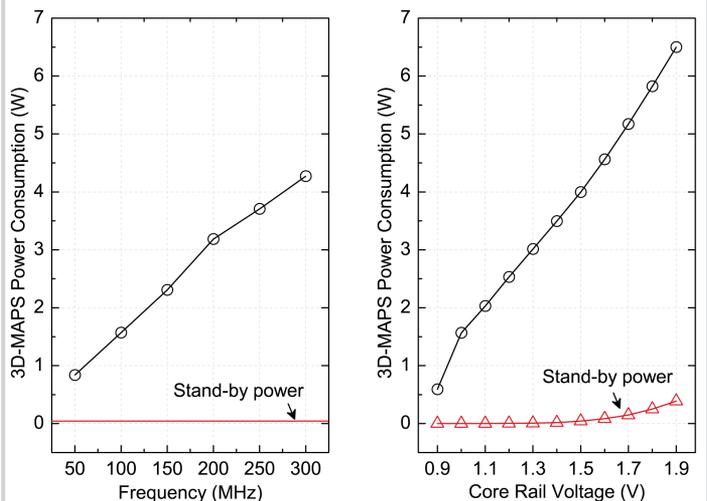


Figure 10.6.6: Frequency vs. power (at 1.5V) and voltage vs. power (at 250MHz) for AES encryption.

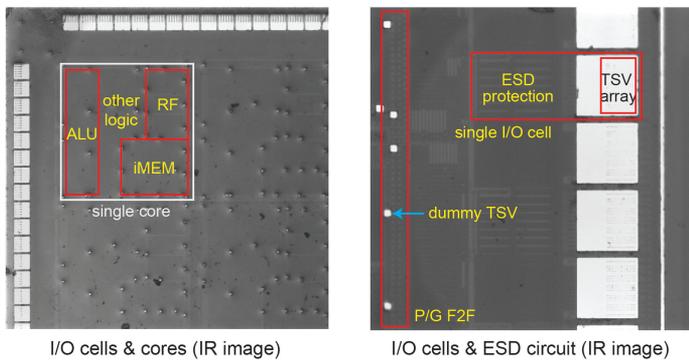


Figure 10.6.7: Die photos.