



Power Electronic Systems
Laboratory

© 2017 IEEE

Proceedings of the 2017 IEEE International Electron Devices Meeting (IEDM 2017),
San Francisco, CA, USA, December 2-6, 2017

Towards Cube-Sized Compute Nodes: Advanced Packaging Concepts enabling Extreme 3D Integration

T. Brunschwiler,
G. Schlottig,
A. Sridhar,
P. Bezerra,
P. Ruch,
N. Ebejer,
H. Oppermann,
J. Kleff,
W. Steller,
M. Jatlaoui,
F. Voiron,
Z. Pavlovic,
P. McCloskey,
D. Bremner,
P. Parida,
F. Krismer,
J. W. Kolar,
B. Michel

Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Towards Cube-Sized Compute Nodes: Advanced Packaging Concepts enabling Extreme 3D Integration

T. Brunschwiler¹, G. Schlottig¹, A. Sridhar¹, P. Bezerra², P. Ruch¹, N. Ebejer¹, H. Oppermann³, J. Kleff³, W. Steller⁴, M. Jatlaoui⁵, F. Voiron⁵, Z. Pavlovic⁶, P. McCloskey⁶, D. Bremner⁷, P. Parida⁸, F. Krismer², J. Kolar², and B. Michel¹

¹IBM Research – Zurich, Rüschlikon, Switzerland, email: tbr@zurich.ibm.com

²ETH – PES, Zurich, Switzerland, ³FhG – IZM, Berlin, Germany, ⁴FhG – IZM – ASSID, Dresden, Germany,

⁵Murata, Caen, France, ⁶Tyndall National Institute, Cork, Ireland, ⁷Optocap, Livingston, UK,

⁸IBM TJ Watson Research Center, Yorktown, NY, US

Abstract—Novel heat removal and power delivery topologies are required to enable ‘extreme 3D integration’ with cube-sized compute nodes. Therefore, a technology roadmap is presented supporting memory-on-logic and logic-on-logic in the medium and long-term, by (i) dual-side cooling and integrated voltage regulators, and (ii) interlayer cooling and electrochemical power delivery.

I. INTRODUCTION

Vertical and heterogenous integration of integrated circuits are key drivers supporting further system efficiency and performance gains in times of slowing transistor node scaling. Yield issues motivated a first 2.5D integrated product, with FPGA dies mounted on a silicon interposer [1]. Memory stacks are the first 3D products at high volume, outperforming single-die DRAM components in terms of capacity per volume and IO power efficiency [2]. Currently, gaming servers, but also high-performance computers use modules with a GPU integrated on a silicon interposer together with high-bandwidth memory stacks [3]. These first products work with existing heat removal and power delivery methods, either due to the low power dissipation or the 2.5D integration of the high-power dies. Chip stacks with high-power logic dies will require novel cooling and power delivery solutions to mitigate transistor degradation by thermal gradients and electromigration due to excessive current densities in the solder interconnects. In this paper, we present a heat removal and power delivery roadmap towards cube-sized compute nodes through ‘extreme 3D integration’. Modules with memory-on-logic are the next step from today’s products, with a medium term time-frame target. Dual-Side Cooling (DSC) and Integrated Voltage Regulators (IVRs) will serve these type of stacks. Full stacking flexibility is needed in the long-term, to allow stacks of multiple logic and memory dies, with up to tens of layers. Volumetric heat removal and power delivery, such as Interlayer Cooling (ILC) and Electrochemical Power Delivery (EPD) topologies are key to support this family of stacks.

II. DUAL-SIDE COOLING (DSC)

Embedded microcavities in a silicon-interposer enable convective heat extraction from the bottom side of a chip stack, while maintaining electrical functionality using Through-Silicon-Vias (TSVs). DSC results by combining this with a back-side cold plate (Fig. 1). This topology supports high-power dies as the bottom and/or top tier of the chip stack. A liquid coolant is introduced through a manifold, delivering fluid

to the back-side cold plate and in parallel to the silicon-interposer cavity. The flow in the cold plate is branched into subsections, to mitigate pressure drop [4], compared to the cross-flow heat exchange in the interposer. The resulting flow rate depends on the hydraulic diameter, defined by the dual-shell concept and the TSV aspect ratio. The fluid cavity can be populated with microchannels or pin-fin arrays in 2- or 4-port configuration, with fluid delivery and drainage in east and west or east & west and north & south, respectively. Lowest junction temperatures for a given pressure drop are achieved by the 4-port with microchannels [5] (Fig. 2). A DSC thermal demonstrator module was built with one thermal die to proof the heat removal performance. CuSn interconnects were formed between the interposer shells to establish required sealing and electrical interconnects (Fig. 3). Thermal simulations using 3D-ICE [6] predicted the measured performance within the experimental uncertainty. Further simulations using the validated model showed thermal gradients of less than 50K at a pressure drop of 30kPa for a chip stack with a 4cm² footprint, consisting of a GPU, cache and CPU layer (total power 672W) [7].

III. INTERLAYER COOLING (ILC)

Volumetric heat removal can be achieved by the integration of microcavities into the backside of each die in the stack (Fig. 4) [8]. This topology enables ‘extreme 3D integration’. However, smaller TSV pitches need to be accommodated, resulting in reduced hydraulic diameters. Therefore, pressure drops of up to 100kPa are required and the lateral chip stack dimension needs to be ≤ 10 mm, to yield acceptable cooling performance [7]. A three die thermal demonstrator stack, including TSVs with a pitch of 100 μ m was built (Fig. 4). A temperature gradient of 54.7K resulted from an aggregate power of 390W and 2.5W/mm² dissipated in hot-spots (Fig. 5). The corresponding volumetric heat flux was 3.9kW/cm³ [8]. In a thought experiment, a 4cm² die with a uniform power dissipation of 1.5W/mm² (total of 600W) is diced and integrated into a 50 tier stack, resulting in a cube of 2x4x2.55mm³. The aggregate area and volumetric heat flux is 75W/mm² and 29 kW/cm³, respectively. The thermal response considering straight channels with a height and width of 33 μ m and 16 μ m was modeled. For a pressure drop of 100kPa, a thermal gradient of 60K could be maintained even for an uniform TSV pitch of 23 μ m. The cube form-factor stack would have a maximal wiring length of 8.5mm, compared to 40mm of the single die implementation [7].

IV. INTEGRATED VOLTAGE REGULATORS (IVR)

IVRs are used to provide granular voltage domains in multi-core CPUs and rapid Dynamic Voltage Scaling (DVS) in response to changing workloads for a more energy-efficient operation. The multiphase inductor-based buck converter is the most suitable topology due to its easy regulation and high achievable efficiencies at high power densities. A fully embedded IVR (2D-IVR) on a microprocessor die is expensive because of the large size of the inductors and lack of magnetic components in current CMOS processes. A hybrid 2.5D-IVR approach, investigated in this work, allows technology separation where the active low-loss CMOS power switches are directly embedded in the microprocessor die and the passive components are fabricated in the low-cost interposer (Fig. 6). The converters' specifications are based on requirements of microprocessors and is demonstrated for a low, but scalable output power of 0.8W at 0.8V nominal voltage. The high input voltage (twice the transistors' breakdown voltage) of 1.6V is feasible using the stacked half-bridge configuration of [9]. A Power Management IC (PMIC) was fabricated using the 14 nm deep-submicron CMOS process. It contains 4-half-bridges, open-loop control circuitry and a programmable load (Fig. 9). Racetrack coupled inductors (Fig. 7) were implemented consisting of electroplated copper winding and Ni₄₅Fe₅₅ magnetic core. Deep trench mosaic PICS capacitors with low loss at >100MHz switching frequencies [12] were implemented in the silicon interposer. The PMICs were flip chip attached to the silicon interposer by thermocompression bonding, while the inductor chips were assembled using SnBi solder reflow (Fig. 6). The functionality of the PMIC was verified without the passive components at frequencies >100 MHz (Fig. 8). Multi-objective optimizations, similar to [13] were performed to optimally size the transistors and passives, in order to maximize the IVRs' power density at ~90% efficiency and switching frequencies >100 MHz. For the given specifications, a maximum efficiency of 88.4% was achieved at a power density of 0.06W/mm² (Fig. 10).

V. ELECTROCHEMICAL POWER DELIVERY (EPD)

A vertically scalable technology for EPD was proposed by enhancing the coolant with soluble redox-active compounds as a means to use the coolant also as an energy carrier and allowing power distribution close to V_{dd} with minimal conversion steps [15] (Fig. 11). Redox flow cells were implemented on silicon by fabricating 1x1cm² arrays of microchannels by deep reactive ion etching and aligning two such half-cells using a semi-automatic pick-and-place tool. A stack of electrochemically functional materials comprising a pair of macroporous carbon papers sandwiching a nanoporous separator was placed in between the two half-cells prior to compression and underfilling (Fig. 12). EPD was demonstrated using vanadium-based electrolytes in sulfuric acid (1.6 M vanadium ions in 4 M total sulfate) comprising predominantly V^{II} in the negolyte and V^V in the posolyte as active species, respectively (state-of-charge ~90%). A maximum power of 741mW at 0.72V was measured in the 1x1cm² silicon-based flow cell (Fig. 13). The heat sink performance was quantified

under steady-state heat flux conditions, yielding a thermal resistance of 41Kmm²/W. Power densities of 2.6W/cm² have been reported in optimized flow cell configurations [16] while the ohmic limit (U²/4R) for power density is estimated at 9.4W/cm², indicating further opportunities for power density improvements in integrated redox flow cells.

VI. SUMMARY & CONCLUSION

A quantitative benchmarking study was performed, on the presented technologies (Fig. 14). The supported die power density for cooling and power delivery is depicted in the y-axis. In general, power delivery constrains the power density more than the cooling solutions. The current feed through external IVR (ext-IVR) is limited to 0.72W/mm² due to electromigration of the solder joints from the stack to the board. A 2D-IVR demonstration with 1W/mm², a voltage conversion ratio of 0.77 and efficiency of 84% [14] does improve the density only by 9%. However, DVS improves the system efficiency for dynamic workloads. A 58% improvement results with our 2.5D-IVR for stacks with >4 dies, due to the 1:2 voltage conversion ratio, considering the integration of inductors and capacitors on each of the dies and not only the interposer. For stacks with up to 4 dies, its poor inductor density is limiting and defines the die power density to 0.3W/mm². Larger conversion ratios would be of benefit and could be accomplished by the integration of high-bandgap semiconductors onto the CMOS-die. Back-side cooling (BSC) is sufficient to match the performance of the all IVR technologies up to 5 dies. From 6 dies on, DSC is required. ILC outperforms DSC for stacks with more than 6 or 2 dies for 4cm² and 1cm² stacks, respectively. Single-phase operation outperforms the use of a dielectric refrigerant in two-phase mode. The EPD power densities demonstrated in a silicon [15] or flow cell [16] implementation outperform the IVR cases for stacks of more than 50 to 100 dies. However, with further advances, based on improved redox chemistries, one can expect higher power densities up to the theoretical limit of 9.4 W/cm², outperforming IVRs at >10 dies.

ACKNOWLEDGMENT

This work was funded in part by the European Commission (CarrICool, FP7-ICT-619488) and the Swiss Confederation (Nano-Tera CMOSAI, 20NAN1-123618 and SNSF REPCOOL, 147661).

REFERENCES

- [1] C. Erdmann et al., *Solid-State Circuits*, vol. 50, pp. 258-269, 2015.
- [2] J. Jeddloh, B. Keeth, *Symp. on VLSI Technology*, pp. 87-88, 2012.
- [3] C. Lee et al., *66th IEEE ECTC*, pp. 1439-1444, 2016.
- [4] G. Schlottig et al., *Journal of Electronic Packaging*, 138 (1), 2016.
- [5] O. Ozsun et al., *16th IEEE ITherm*, Orlando, FL, pp. 473-479, 2017.
- [6] A. Sridhar et al., *IEEE ICCAD*, San Jose, CA, 2010.
- [7] T. Brunswiler et al., *Journal of Electronic Packaging*, 138 (1), 2016.
- [8] T. Brunswiler et al., *IEEE Conference on 3DIC*, 2009.
- [9] P. Bezerra et al., *18th IEEE COMPEL*, Stanford, California, USA, 2017.
- [10] M. Jatlaoui et al., *10th EuMIC*, Paris, France, pp. 203-206, 2015.
- [11] F. Voiron et al., *IEEE IWIPP*, Chicago, US, pp. 48-51, 2015.
- [12] F. Neveu et al., *IEEE T. on Power Electronics*, vol. 31, pp. 3985, 2016.
- [13] P. Bezerra et al., *13th COBEP/SPEC*, Fortaleza, Brazil, pp. 1-6, 2015.
- [14] H. Krishnamurthy et al., *IEEE ISSCC*, San Francisco, CA, 2017.
- [15] P. Ruch et al. *IBM J. Res. Dev.* 55, 2011.
- [16] R. Elgammal et al., *Electrochimica Acta* 237, pp. 1-11, 2017.
- [17] P. Parida et al., *IEEE SEMI-THERM*, San Jose, CA, 2017.




Fig. 1: Sketch of the dual-side cooling topology.




Fig. 2: X-ray image of a 4-port interposer with microchannel cavities.




Fig. 3: a) Interposer with sealing-rings and b) embedded TSVs required to prevent electrical shorts in case of water as a coolant. A cavity height of $160\mu\text{m}$ and a pin-fin diameter of $75\mu\text{m}$ was achieved for a TSV pitch, diameter and depth of $225\mu\text{m}$, $15\mu\text{m}$ and $125\mu\text{m}$, respectively. c), d) Microscopic view of the dual-side cooled thermal demonstrator module in 2-port and pin-fin configuration.




Fig. 4: a) Photograph of the interlayer cooled thermal test vehicle. b), c) SEM view of the cross-section depicting the three active dies with microchannels and TSVs.

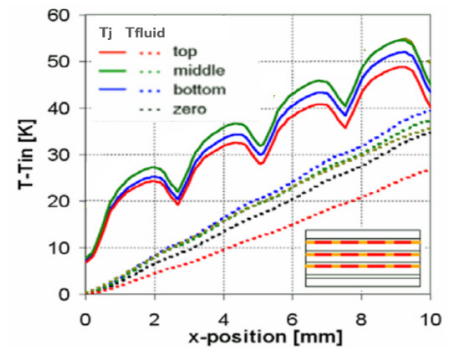


Fig. 5: Modeled junction and fluid temperature of the ILC chip stack with 3 tiers and 4 cavities, populated with pin-fins ($100\mu\text{m}$ height, $50\mu\text{m}$ diameter) across the 1cm^2 chip stack footprint.




Fig. 6: Photograph of the assembled 2.5D-IVR corresponding to the schematic of Fig. 9. The PMIC (14 nm Bulk CMOS technology) chip and the coupled inductors are soldered on the silicon interposer which contains the input and output decoupling capacitors. The IVR was designed for a nominal power of 0.8W at frequencies $> 100\text{MHz}$




Fig. 7: Microscopic bottom view of the designed racetrack coupled inductor optimized to operate at 100 MHz. The inductor uses electroplated copper windings and electroplated $\text{Ni}_{45}\text{Fe}_{55}$ magnetic core.




Fig. 8: Measurement of the switching nodes' voltage, to verify the functionality of the PMICs' control and power stage.

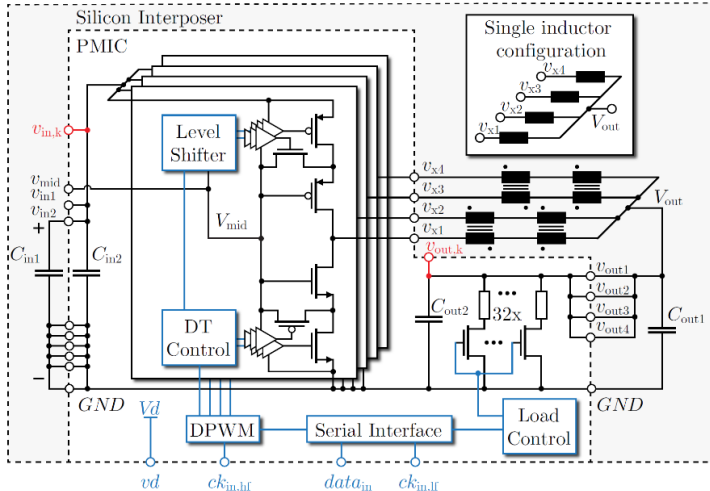


Fig. 9: Schematic representation of the complete buck converter showing the boundary between the PMIC chip and the interposer including the pads' representation. The PMIC chip contains the power switches, the load, and the circuitry to operate the IVR in open loop mode. The power half-bridges can be operated with phase-shift allowing the use of coupled and single inductors. Dead-time, duty-cycle and enabling of the phases are controlled via serial communication.

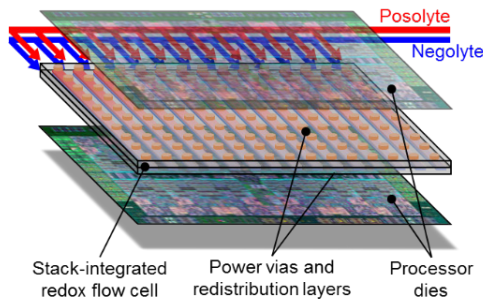


Fig. 11: Concept of combined power delivery and cooling for integrated circuits using liquid electrolytes in redox flow cells.




Fig. 12: Cross-section of a redox flow cell consisting of two half-cells with fluid distribution channels etched in silicon. Highly p- and n-doped silicon current collectors sputter-coated with Ni / C or Au were used to sandwich carbon paper electrodes and a nanoporous PP-based separator. The channels were 10 mm long, 700 μm wide and 250 μm deep with 75 μm thick walls.




Fig. 13: Polarization curve at a flow rate of 500 mL/min for both posolyte and negolyte, which were composed of 1.6 M vanadium in 4 M total sulfate electrolyte charged to a state-of-charge (SOC) of ~90%. The three boxes represent three different operating points with voltage efficiencies of 90%, 70% and 49%.

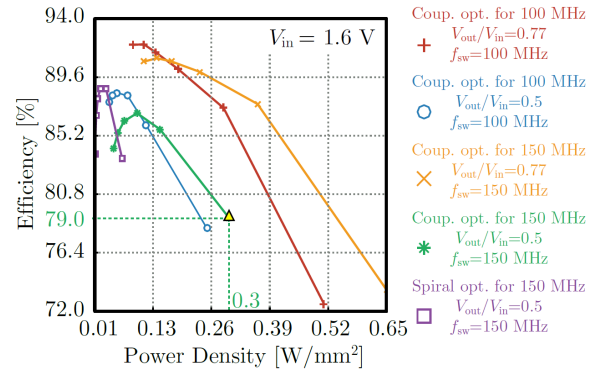


Fig. 10: Simulation results of efficiency and power density of the assembled four-phase IVRs (using racetrack coupled inductors with magnetic core and spiral air-core inductors) operating at different switching frequencies and conversion ratios. The simulations consider the transistors' losses as modeled in Cadence, measured DC inductance and ESR of the inductors, and measured capacitance, ESR and ESL of the capacitors. The optimized IVR using coupled inductors operating at 100 MHz at nominal conditions (0.8W, 0.5 conversion ratio) achieves a peak efficiency of 88.4% at 0.06 W/mm². Higher power densities are achievable by increasing the switching frequency and the conversion ratio. For the summary graph of Fig. 14, the operation point labeled with 'Δ' was chosen.




Fig. 14: Benchmarking of power delivery and heat removal technologies: the maximal power density per die is presented for the number of dies in a stack. (ext-IVR: external IVR; 2.5D-IVR: inductor and capacitor integrated on interposer or on every die back-side for $n > 1$ (79% efficiency, 0.5 conversion ratio, 0.3W/mm² output power density); 2D-IVR: inductor integrated into back-end-of-line layers [14] (84%, 0.77, 1W/mm²); EPD: EPD in silicon; EPD-FB: EPD in flow cell; EPD-Th: theoretical limit of EPD; ILC: 1p, 10: ILC single-phase, 10mm chip size; ILC: 2p, 20: ILC two-phase, 10mm chip size). R1234ze was considered as a refrigerant for 2p cooling with a channel width of 75μm (no need for fluid sealing) instead of 50μm for single-phase water cooling [17].