

Thermal Analysis of a 3D Stacked High-Performance Commercial Microprocessor using Face-to-Face Wafer Bonding Technology

Rahul Mathur^{*‡}, Chien-Ju Chao^{*}, Rossana Liu^{*}, Nikhil Tadepalli[†], Pranavi Chandupatla^{*}, Shawn Hung^{*},
Xiaoqing Xu^{*}, Saurabh Sinha^{*}, and Jaydeep Kulkarni[‡]

^{*}Arm Inc., 5707 Southwest Parkway, Austin, TX, 78735

Email: rahul.mathur@arm.com

[†]Arm Ltd, Cambridge, UK

[‡]University of Texas at Austin, TX, USA

Abstract—3D integration technologies are seeing widespread adoption in the semiconductor industry to offset the limitations and slowdown of two-dimensional scaling. High-density 3D integration techniques such as face-to-face wafer bonding with sub-10 μm pitch can enable new ways of designing SoCs using all 3 dimensions, like folding a microprocessor design across multiple 3D tiers. However, overlapping thermal hotspots can be a challenge in such 3D stacked designs due to a general increase in power density. In this work, we perform a thorough thermal simulation study on sign-off quality physical design implementation of a state-of-the-art, high-performance, out-of-order microprocessor on a 7nm process technology. The physical design of the microprocessor is partitioned and implemented in a 2-tier, 3D stacked configuration with logic blocks and memory instances in separate tiers (logic-over-memory 3D). The thermal simulation model was calibrated to temperature measurement data from a high-performance, CPU-based 2D SoC chip fabricated on the same 7nm process technology. Thermal profiles of different 3D configurations under various workload conditions are simulated and compared. We find that stacking microprocessor designs in 3D without considering thermal implications can result in maximum die temperature up to 12°C higher than their 2D counterparts under the worst-case power-indicative workload. This increase in temperature would reduce the amount of time for which a power-intensive workload can be run before throttling is required. However, logic-over-memory partitioned 3D CPU implementation can mitigate this temperature increase by half, which makes the temperature of the 3D design only 6°C higher than the 2D baseline. We conclude that using thermal-aware design partitioning and improved cooling techniques can overcome the thermal challenges associated with 3D stacking.

Index Terms—3D stacking, face-to-face, thermal

I. INTRODUCTION

With the slowdown of Moore’s Law of two-dimensional scaling [1], the semiconductor industry is at a critical juncture as 2.5D and 3D stacking technologies are being actively explored and adopted by certain end-applications [2]. 3D stacking enables vertical integration of more than one layer of active transistors and interconnects with the goal of increasing compute density. Integrated circuit (IC) designs with natural redundancy and regularity in 2D can be extended or stacked in the third dimension with relative ease. CMOS image sensors [3], DRAM memories [4], and NAND Flash memories [5]

are all examples of ICs with high amounts of regularity and redundancy, and these products have already embraced 3D integration and achieved success in high-volume manufacturing.

However, the adoption of 3D stacking for logic applications has been relatively slow. Functionally complete chips commonly referred to as chiplets, are stacked using packaging technologies. The stacking configuration could be 2.5D, wherein, chiplets are assembled in 2D but interconnected through an underlying substrate, e.g., Silicon interposer or redistribution layer (RDL). Alternatively, the stacking configuration could be 3D, e.g., package-on-package (PoP) wherein DRAM packaged dies are stacked on ASIC die [6] or, two or more compute dies are stacked using through-silicon-via (TSV) and micro-bump technology [7].

Technology trends point towards finer-pitch 3D connectivity in the form of wafer or die-stacking 3D integration, which, we refer to as high-density 3D integration. Hybrid-bonding technologies allow precision alignment of wafers resulting in 3D connection pitches in the range of 10 μm [8] to 1 μm [9]. This technology opens up the possibility of designing systems where functional units are partitioned and co-designed across separate 3D tiers. The advantages of such 3D integration are multi-fold:

- Functional blocks separated across large distances in 2D can be brought closer in 3D, reducing interconnect delay and power.
- Large-die SoCs can be partitioned into smaller dies, improving yield and potentially leading to lower cost.
- Dies from different process technologies can be integrated enabling heterogeneous integration, leading to flexible product migration to advanced nodes for further cost reduction [10].

However, 3D stacking of compute dies can result in higher power density, leading to power delivery and thermal challenges. Performance throttling via dynamic voltage and frequency scaling (DVFS) can combat increasing on-die temperatures, but that would also offset some of the promised performance gains of 3D stacking. In this paper, we provide a comprehensive overview of the thermal impact of high-

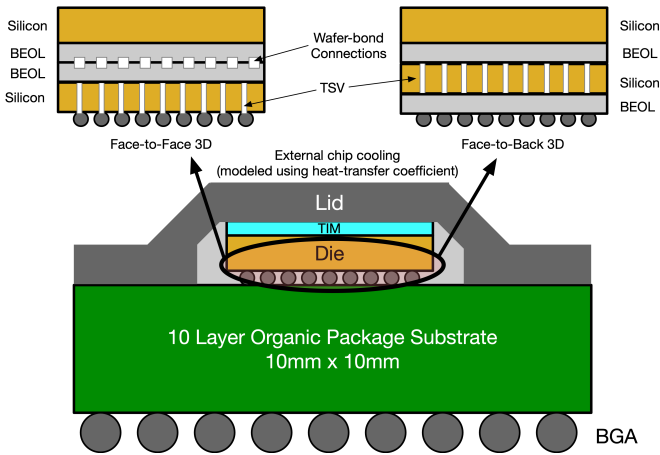


Fig. 1. Cross-section of the die and package model used in this study.

density 3D stacking technologies in the context of a design-aware partitioned and timing-optimized 3D high-performance microprocessor design. The main contributions of this paper are as follows.

- To the best of our knowledge, this paper presents the first comprehensive thermal simulation study on a sign-off quality physical design of a 3D high-performance microprocessor using face-to-face (F2F) wafer-bonding technology.
- The thermal simulation model is calibrated to 2D hardware measurement data from a high-performance CPU-based SoC.
- By evaluating various 3D stacking configurations in our study, we demonstrate that taking thermal implications into account early in the physical design process of a 3D partitioned microprocessor can partially mitigate the thermal impact of 3D stacking.

Figure 1 describes the package and die cross-section model in this study. Face-to-face (F2F) 3D connections are formed using wafer-bond pads at the F2F interface, while TSVs are used to escape the bottom of the die for power delivery and I/O signals. Face-to-back (F2B) 3D connections require TSVs for connection between the two dies. The rest of the paper is organized as follows: Section II explains the implementation flow of 3D microprocessor design. Section III and IV describe the thermal analysis simulation framework and experimental setup. Section V discusses results and provides insights by comparing the thermal characteristics of different 3D configurations. In Section VI, conclusions are drawn for thermal aware design of high-performance microprocessors using 3D integration technologies.

II. 3D CPU IMPLEMENTATION

The 2D reference design is a 64-bit, out-of-order, high-performance Arm[®] CPU. The physical design of the CPU was done in 3D, using a novel physical implementation flow which consists of design-aware RTL partitioning and placement co-optimization across multiple tiers.

TABLE I
COMPARISON OF 2D AND 3D CPU

Metric	2D	3D
Process node	7nm	7nm
Maximum frequency	3 GHz	3 GHz
L2 capacity (normalized)	1X	2X
L2 access latency (normalized)	1X	1X
2D Footprint (normalized)	1X	0.77X

Specifically, using high-density 3D F2F wafer-bonding, the reference design was partitioned into separate memory and logic tiers. The memory tier consists of Level 1 (L1) and Level 2 (L2) caches, while the logic tier consists of the logic blocks including the CPU datapath. Additionally, the 3D CPU was configured to include twice the L2 cache capacity compared to the 2D design. Increasing the L2 cache capacity in a 2D design results in a larger CPU floor-plan, longer wire-delay and hence, has a higher memory access latency. Since 3D stacking enables bringing compute and memory blocks closer together, the 3D CPU design can accommodate a 2X larger L2 size compared to the 2D design with the same memory access latency. Details of the 2D and 3D CPU designs are provided in Table I.

Utilizing a novel 3D implementation flow, which allows cross-tier placement optimization, the 3D CPU implementation was able to achieve the same target frequency (3 GHz) as the 2D design. The critical steps of the flow are described in Figure 2 and more details are provided in [11], [12]. This highlights the importance of design-aware partitioning and cross-tier optimization to achieve high-performance 3D design. Figure 3 shows the placed and routed 2D and 3D CPU designs. Since the 3D CPU consists of a 2X larger L2 cache, it has a larger total area compared to the 2D CPU (1.54X total silicon area relative to the 2D design). However, the 2D footprint of the 3D stacked design is approximately 0.77X of the 2D design.

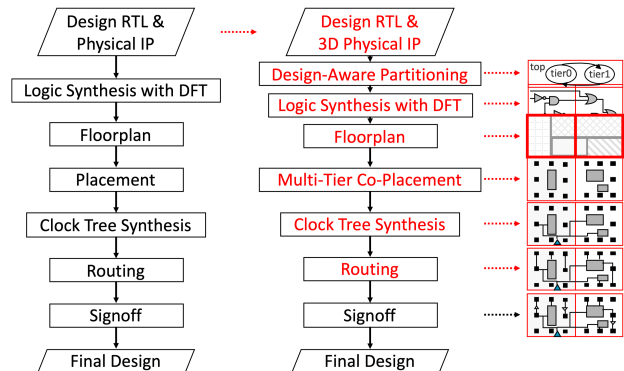


Fig. 2. 2D (left) and 3D (right) physical implementation flow chart. The steps highlighted in red differ from the standard 2D physical implementation flow.

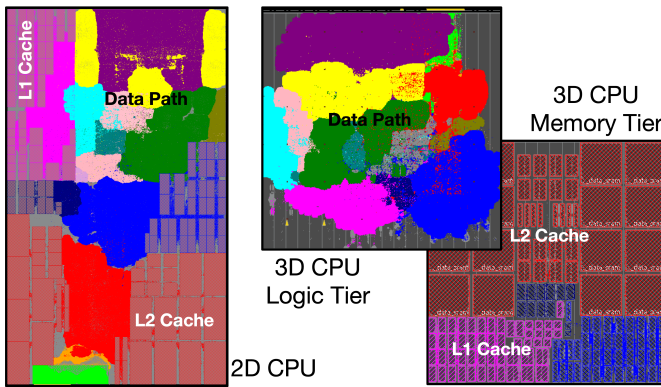


Fig. 3. Layout of 2D CPU, logic and memory tiers of the 3D CPU after physical design implementation sign-off.

III. THERMAL ANALYSIS FRAMEWORK

Figure 4 outlines the flow used to perform package level thermal simulations on the 2D and 3D implementations of the CPU discussed in Section II. *Cadence® Celsius™ Thermal Solver* [13] was used to perform both static and transient simulations. The tool requires a location-based power dissipation map of the die-model as well as a detailed description of the die stack-up, i.e., devices, interconnects and dielectrics along with their thermal conductivity properties. In this section, we describe the methodology to generate the tile-based power map file for the design for thermal analysis and calibration of the thermal boundary conditions with hardware measurement data.

A Vector-based analysis is done using *Synopsys® PrimeTime™ PX* [14], which reports the cell level power consumption of the full design. Two power-indicative workload vectors are utilized to estimate power:

- *dhrystone* is used for characterizing the average power of the CPU. This is indicative of a typical real-world use-case of the CPU.

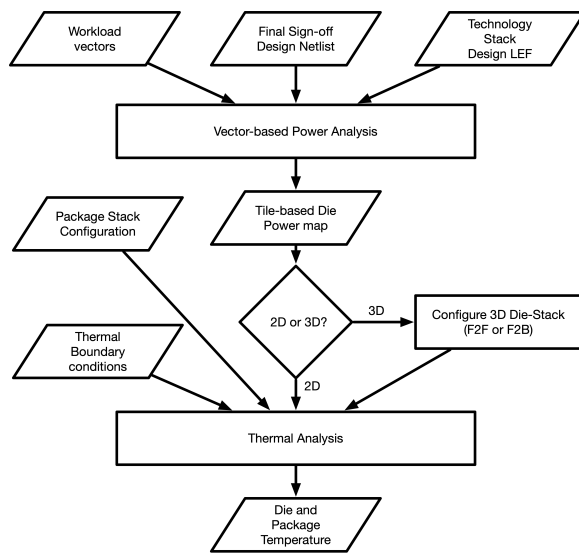


Fig. 4. Flowchart for thermal analysis.

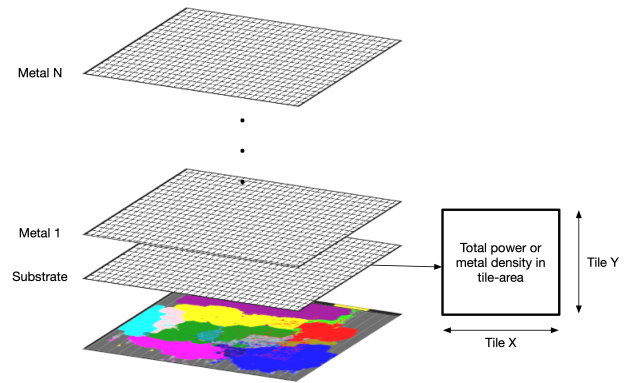


Fig. 5. Tile-based power and metal density map used for thermal analysis.

- *maxpower* is used to characterize the worst-case power dissipation of the CPU. It is important to note that real-world usages typically only have *maxpower* characteristics for a short duration of time.

Based on the results of the vector-based power analysis, a tile-based power map of the die is generated. As depicted in Figure 5, the entire die is divided into equal-sized tiles. The power of each tile is the sum of the internal power, switching power and leakage power associated with the cells within the tile. The number of tiles for the die was chosen to ensure high resolution of the power density within different modules as well as maintain reasonable computation runtime of the tool. The tile-based power map file also contains metal density and thermal properties of all the layers in the back end of line (BEOL) stack as well as the substrate (Silicon). *Cadence Celsius Thermal Solver* uses the power map file along with a complete description of the package stack-up, bumps, die-stack, molding compound, lid, thermal-interface material (TIM) and boundary conditions. Setting up realistic boundary conditions for thermal analysis is critical for getting accurate results and is described in the next subsection.

A. Model Calibration

The thermal simulation boundary conditions are calibrated to hardware measurements on a 4-core SoC test chip fabricated on a 7nm process technology [15]. The *maxpower* workload was run for a fixed number of CPU cycles on all four cores and temperature measurements from on-die temperature sensors were collected. The power dissipation from the four cores was also measured. The package stack-up and die power map matching the SoC was generated and set up for thermal simulations. The hardware measurement setup consists of a heat sink and a fan on top of the package lid. The heat transfer coefficient (HTC) in transient simulations that provided the best match to the temperature measurement data at the different operating frequencies of the chip was finalized as the boundary condition for all subsequent experiments. Figure 6 shows the calibration of our thermal simulation boundary conditions to the temperature sensor measurements.

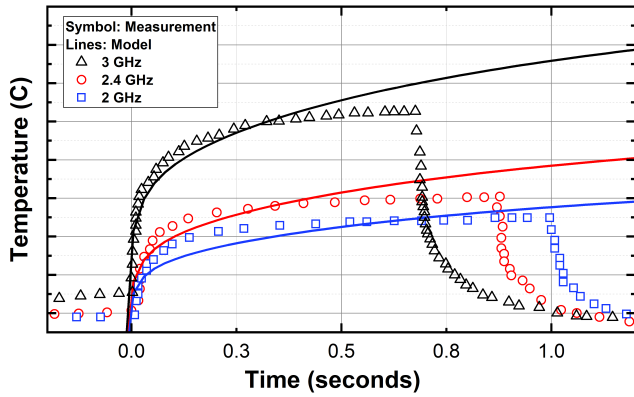


Fig. 6. Thermal simulation boundary condition calibration with on-die measurements. The simulation package, die model, workload, and power dissipated were matched with the measurement setup.

IV. EXPERIMENTAL SETUP

A dual-core CPU cluster was used for our thermal simulation study with the same workload running on both CPUs. It is important to note that the physical implementation of the 3D CPU requires high-density 3D connections in a F2F configuration. However, abstracting the CPU power in terms of tile-based power maps allows us to explore other different 3D configurations such as F2B 3D as well. TSVs are modeled at appropriate locations in the die stack-up, for example, between the bottom-die and the package for F2F configuration; and between bottom-die and top-die in F2B configuration (see Figure 1). For each 3D configuration, logic and memory dies were swapped to study the impact of the proximity of die to external cooling on top of the package. Additionally, a CPU-on-CPU 3D stack is modeled, where each die consisted of a single-core 2D CPU. This represents the current state-of-the-art where functionally complete ‘chipllets’ are stacked in 3D. The analysis of all these configurations is done under the similar die and package size assumptions. The two CPUs are placed $200\ \mu\text{m}$ apart with a margin of 1 mm to the die edges. The static power from a representative system-level cache is allocated to the margin area to model idle caches placed around the CPUs. A $10\times 10\ \text{mm}^2$ package comprising of 10 build-up layers was used for all configurations. All temperature values are reported relative to the dual-core 2D die. Table

TABLE II
3D SIMULATION CONFIGURATIONS

Configuration	Description
# of CPU in cluster	2
F2F	Face-to-Face 3D
F2B	Face-to-Back 3D
Logic-on-Mem	Logic on top and memory on bottom tier
Mem-on-Logic	Memory on top and logic on bottom tier
CPU-on-CPU	2D CPU on top and bottom tier
Spacing between CPUs	$200\ \mu\text{m}$
Margin around CPU cluster	1 mm
Package dimensions	10mm x 10mm

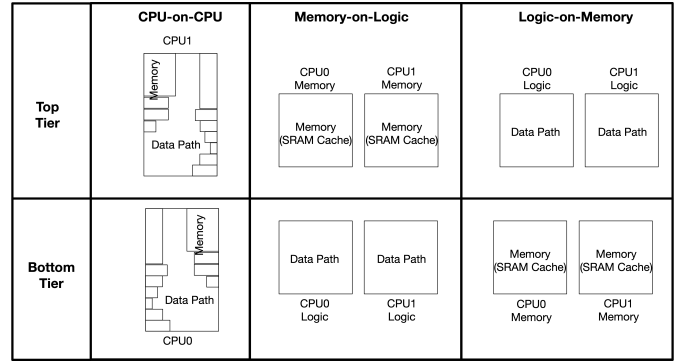


Fig. 7. Pictorial view of the top and bottom tiers of the 3D simulation configurations.

II lists the nomenclature for the different 3D configurations used in our experiments and Figure 7 provides a pictorial representation of how they are organized on the two tiers.

V. RESULTS

The steady-state temperature heat map of a baseline 2D dual-core CPU configuration and the two tiers of the 3D CPU in F2F Logic-on-Memory configuration are shown in Figure 8. The temperature values are relative to the coolest point on the 2D die. Each CPU core is running the *maxpower* workload. The heat map clearly emphasizes that in the 2D CPU, the data-path runs hotter ($\Delta T \approx 6^\circ\text{C}$) than the L1 and L2 cache blocks. This observation directly correlates to the 3D logic-over-memory heat maps, where the logic die is hotter than the memory die despite being in closer proximity to the heat sink.

To gain further insights into the temperature profiles, the maximum and average power density of each die in the 2D and 3D configurations are plotted in Figure 9. Even though the maximum power density on the die is similar between the 2D and the 3D logic-over-memory case, the average power density of the 3D stack is higher, owing to a similar total power of the two designs in a 0.77X smaller footprint. This results in a higher temperatures in the 3D stacked designs.

Figure 10 plots the maximum steady-state temperature (relative to the 2D baseline) for all the 3D configurations running *dhrystone* and *maxpower* workloads. For the 3D logic and memory partitioned CPU, the logic tier is always hotter due to higher power density. For the CPU-on-CPU case, the CPU on the bottom die shows a higher temperature profile since it is not in proximity to the heat sink.

Among the different 3D configurations, the logic-on-memory 3D design has a temperature rise of around 5°C while the memory-on-logic 3D design has a temperature rise of 9°C compared to the 2D baseline. This is because the higher power logic die in the memory-on-logic configuration sits away from the heat sink. CPU-on-CPU has the worst thermal characteristics because of overlapping hotspots from the two tiers, where the maximum and average power density doubles in the CPU-on-CPU configuration compared to a 2D CPU as shown in Figure 9.

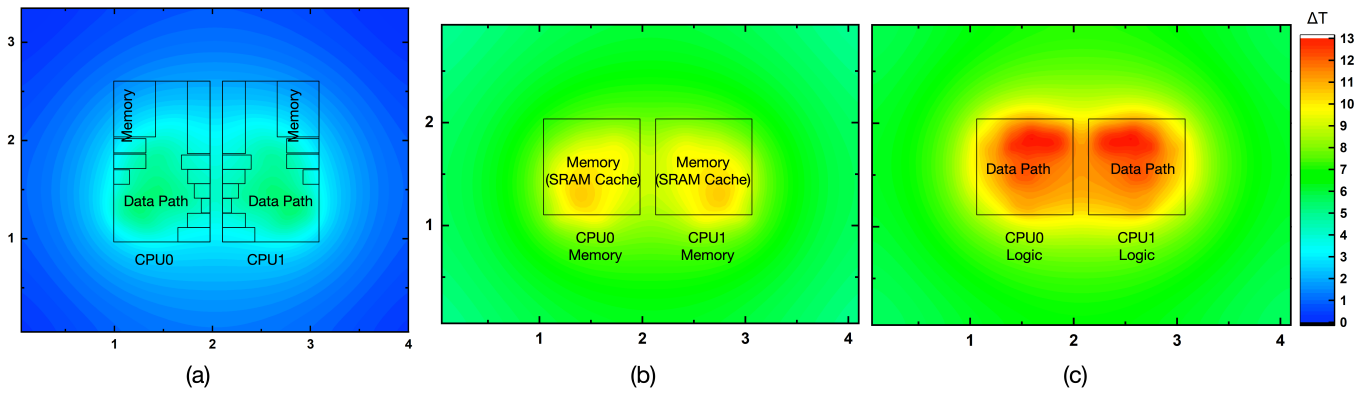


Fig. 8. Heat-maps under the *maxpower* workload for (a) a 2D dual CPU configuration, (b) a 3D dual CPU memory tier and (c) a 3D dual CPU logic tier. The die dimensions are in mm. The temperature values are relative to the coolest point on the 2D die. In this scenario, the 3D stack has the logic tier on top (close to the heat sink) and the memory tier at the bottom (close to the package).

F2F and F2B 3D configurations have a minor impact on the temperature profile, primarily because in each case the bottom die is thinned down to process TSVs, providing a lower effective thermal resistance to the package. To the first order, the temperature rise in 3D stacking is mainly proportional to the effective power density of the design. Figure 11 shows the maximum temperatures on both dies as well as the package across different 3D configurations under the *maxpower* workload relative to the 2D package temperature, which highlights the value of logic-over-memory 3D design in minimizing the temperature rise in the package.

Self-heating of CPUs is a relatively slow process (seconds) compared to the operating frequency of the design (subnanoseconds). Hence, the steady-state temperature profiles provide a limited view of the design or performance trade-offs of 2D and 3D stacked systems. Depending on the time required to reach maximum allowed operating temperature, the power profile of the application and the application runtime, the performance of 3D systems may need to be throttled to stay within the thermal constraints. Figure 12 shows the transient temperature rise for the 2D and 3D stacked configurations under the *maxpower* workload. Assuming the dashed horizontal

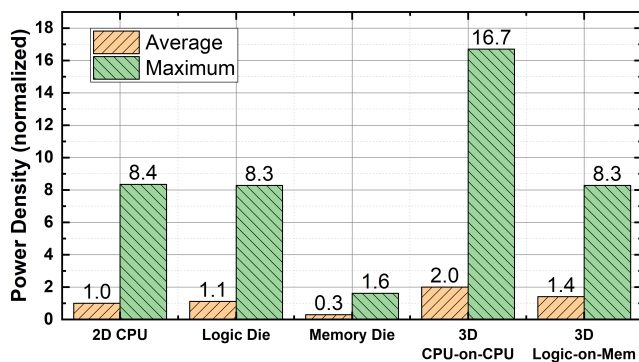


Fig. 9. Maximum and average power density in the 2D and 3D dies for different configurations running the *maxpower* vector.

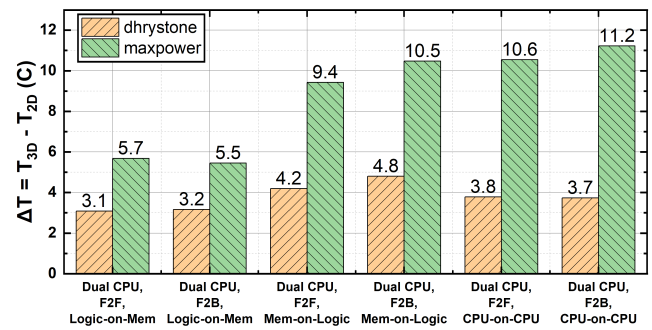


Fig. 10. Increase in maximum temperature for different 3D configurations under the *dhrystone* and *maxpower* workloads relative to 2D baseline.

line corresponds to the maximum allowed temperature for the chip, the 2D CPU can run for 35 seconds before performance needs to be throttled to allow the system to cool down. However, the 3D logic-on-memory design and 3D CPU-on-CPU design can run only for 20 and 15 seconds respectively (a 0.57X and 0.43X reduction). It is important to highlight that this describes a worst-case scenario for power consumption.

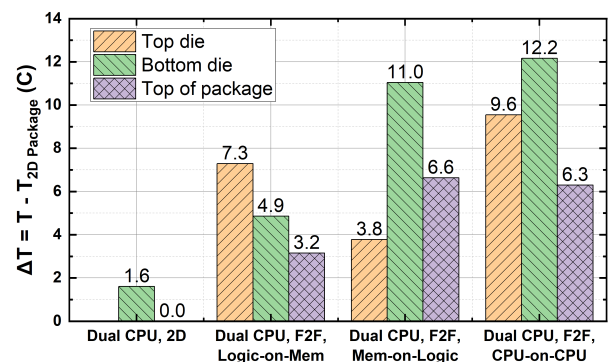


Fig. 11. Package and die temperatures for different 3D configurations under the *maxpower* workload compared to 2D package temperature as baseline.

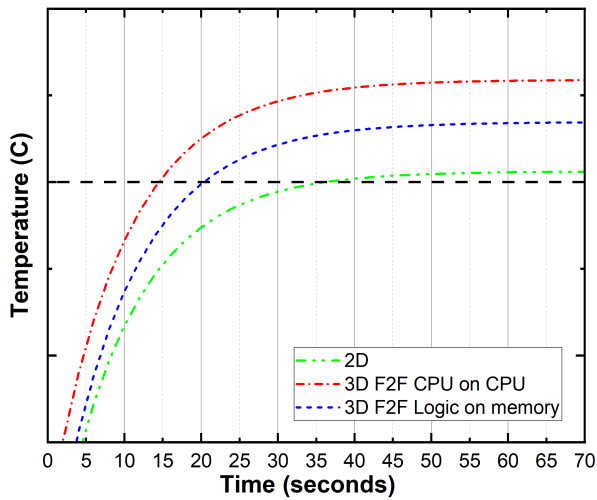


Fig. 12. Transient profile of 2D and 3D configurations under the *maxpower* workload.

We expect that more mainstream workloads featuring average power dissipation can run sustainably with appropriate cooling techniques in 3D. We quantify two known approaches to reduce maximum on-die temperature using the worst-case 3D stacking option, CPU-on-CPU configuration.

The CPU-on-CPU 3D configuration is closer to the current state-of-the-art stacking of functionally complete ‘chiplets’. It does not require design-aware partitioning and timing co-optimization across 3D tiers. However, as seen in the results, this approach has the worst thermal characteristics because the power density of the 3D stack doubles. An approach to mitigate this effect can be to introduce a physical offset between the CPUs on the two tiers and remove overlapping thermal hotspots. Figure 13 shows that the maximum temperature of this configuration can be reduced by around 5°C for *maxpower* and 2°C for *dhrystone* when the offset is swept from 0 to 1 mm.

It should be noted that while the method of introducing an offset works for a CPU-on-CPU configuration, it is not feasible for the 3D configurations with logic and memory partitioning since the design utilizes the 3D proximity to reduce latency, enable a larger cache capacity, and close timing at 3 GHz.

Another approach for reducing die temperatures for 3D could be to use more sophisticated cooling techniques, such as liquid cooling [16]. Figure 14 shows the impact of increasing the heat transfer coefficient (HTC) from the hardware calibrated data, on the maximum temperature of a F2F 3D stack CPU-on-CPU configuration. Increasing the HTC on the top of the package i.e., improved cooling is a more effective way of controlling the temperature rise with 3D. It is worth clarifying that the absolute decrease in temperature is more pronounced for the *maxpower* workload compared to *dhrystone* because the steady-state maximum temperature from the ambient for the *maxpower* workload is much higher than the *dhrystone* workload.

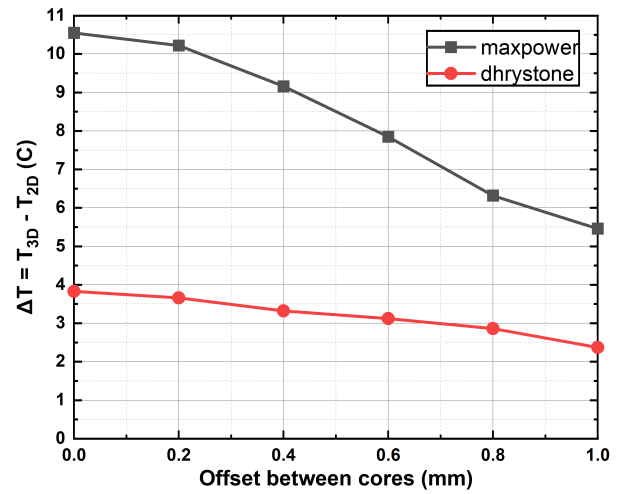


Fig. 13. Effect of offset between two tiers on maximum temperature for dual CPU, F2F, CPU-on-CPU under the *dhrystone* and *maxpower* workloads compared to 2D baseline.

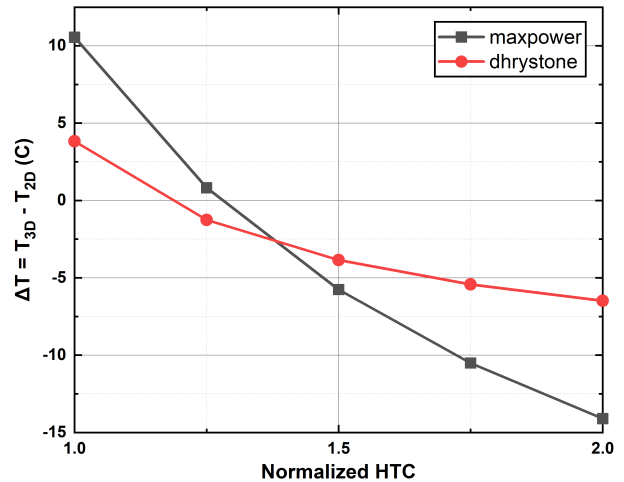


Fig. 14. Effect of increasing HTC on maximum temperature for dual CPU, F2F, CPU-on-CPU under the *dhrystone* and *maxpower* workloads compared to 2D baseline.

It is expected that 3D stacked designs will have a relatively lower thermal headroom compared to 2D designs due to higher power density. Through our analysis, it is clear that thermal-aware design methodologies, especially tightly integrated with the physical design flow, will have to be deployed and carefully used in order to get 3D systems to stay within thermal operating budgets, or else in-situ control techniques like DVFS and throttling may have to be employed more at runtime, compared to 2D. It is shown that logic-over-memory partitioned 3D designs provide the best trade-off to manage the aggravated thermal impact. Using improved cooling techniques can further mitigate this issue and thermal challenges are not a showstopper for the adoption of 3D integration technologies.

VI. CONCLUSION

In this work, a comprehensive thermal simulation study of a 3D stacked high-performance CPU using F2F wafer bonding technology was presented. The physical design of the 3D CPU was partitioned into logic and memory tiers and was implemented using a novel physical implementation flow that allows cross-tier timing and placement optimization. Workload-based power and thermal simulation analysis of multiple 3D configurations was performed and compared. All 3D stacked designs show higher maximum temperature due to higher power density. However, it is found that the logic-over-memory 3D stacked configuration is the least impacted, thermally, resulting in a maximum temperature rise of only 6°C for a worst-case power workload, compared to the 2D baseline. Hence, if thermal-aware design techniques are employed, we emphasize that a relatively small increase in temperature is not a showstopper for 3D integration technologies. This work paves the pathway for future thermal-aware 3D designs.

ACKNOWLEDGMENT

The authors will like to thank Heath Perry and Alex Wasson from Arm for providing CPU temperature measurement data and also thank Robert Christy, Rob Aitken and Brian Cline from Arm for technical feedback on the manuscript.

REFERENCES

- [1] G. Moore. Cramming more components onto integrated circuits. *Electronics*, April 1965.
- [2] G. Yeric. Moore's law at 50: Are we planning for retirement? In *2015 IEEE International Electron Devices Meeting (IEDM)*, pp. 1.1.1–1.1.8, Dec 2015.
- [3] R. Fontaine. The State-of-the-Art of Smartphone Imagers. In *2019 International Image Sensor Workshop (IISW)*, 2019.
- [4] H. Jun et al. HBM (High Bandwidth Memory) DRAM Technology and Architecture. In *2017 IEEE International Memory Workshop (IMW)*, pp. 1–4, May 2017.
- [5] S. Venkatesan et al. Overview of 3D NAND Technologies and Outlook Invited Paper. In *2018 Non-Volatile Memory Technology Symposium (NVMTS)*, pp. 1–5, Oct 2018.
- [6] C. Tseng et al. InFO (Wafer Level Integrated Fan-Out) Technology. In *2016 IEEE 66th Electronic Components and Technology Conference (ECTC)*, pp. 1–6, May 2016.
- [7] W. Gomes et al. Lakefield and Mobility Compute: A 3D Stacked 10nm and 22FFL Hybrid Processor System in 12×12mm², 1mm Package-on-Package. In *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, 2020.
- [8] M. Chen et al. System on Integrated Chips (SoIC(TM) for 3D Heterogeneous Integration. In *2019 IEEE 69th Electronic Components and Technology Conference (ECTC)*, pp. 594–599, May 2019.
- [9] A. Jouve et al. 1m pitch direct hybrid bonding with <300nm wafer-to-wafer overlay accuracy. In *2017 IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S)*, pp. 1–2, Oct 2017.
- [10] L. England et al. Advanced packaging saves the day! — How TSV technology will enable continued scaling. In *2017 IEEE International Electron Devices Meeting (IEDM)*, pp. 3.5.1–3.5.4, Dec 2017.
- [11] X. Xu et al. Enhanced 3D Implementation of an Arm® Cortex®-A Microprocessor. In *2019 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 1–6, July 2019.
- [12] S. Sinha et al. Stack up Your Chips: Betting on 3D Integration to Augment Moore's Law Scaling. In *2019 IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S)*, Oct 2019.
- [13] Celsius thermal solver. https://www.cadence.com/en_US/home/tools/sys-tem-analysis/thermal-solutions/celsius-thermal-solver.html.
- [14] PrimeTime PX. <https://www.synopsys.com/support/training/signoff/prime-timepx-fcd.html>.
- [15] R. Christy et al. A 3GHz ARM Neoverse N1 CPU in 7nm FinFET for Infrastructure Applications. In *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, 2020.
- [16] J. Lau. 3D IC Integration and Packaging. chapter 9. McGraw-Hill Education, 2015.