Thermal Scaffolding for Ultra-Dense 3D Integrated Circuits

Dennis Rich
Electrical Engineering
Stanford University
Stanford, CA
denrich@stanford.edu

Anna Kasperovich
Electrical Engineering
Stanford University
Stanford, CA
akaspero@stanford.edu

Mohamadali Malakoutian
Electrical Engineering
Stanford University
Stanford, CA
malakout@stanford.edu

Robert M. Radway
Electrical Engineering
Stanford University
Stanford, CA
radway@stanford.edu

Shiho Hagiwara
Fujitsu Research
Kawasaki, Kanagawa, Japan
hagiwara.shiho@fujitsu.com

Takahide Yoshikawa
Fujitsu Research
Kawasaki, Kanagawa, Japan
yoshikawa.takah@fujitsu.com

Srabanti Chowdhury
Electrical Engineering
Stanford University
Stanford, CA
srabanti@stanford.edu

Subhasish Mitra
Electrical Engineering and Computer Science
Stanford University
Stanford, CA
subh@stanford.edu

2023 60th ACM/IEEE Design Automation Conference (DAC) | 979-8-3503-2348-1/23/$31.00 ©2023 IEEE | DOI: 10.1109/DAC56929.2023.10247815

***Abstract*— We address the thermal challenge of ultra-dense 3D (e.g., monolithic 3D) integrated circuits with multiple high-speed computing engines in the 3D stack. We present a new *thermal scaffolding* approach achieved through a combination of (1) new Back-End-of-Line (BEOL)-compatible dielectric materials for simultaneous high thermal conductivity and low dielectric constant, (2) new 3D physical co-design of BEOL dielectrics with thermal metal structures for uniform heat conduction with minimal metal insertion overhead, and (3) previous, experimentally demonstrated heatsink advances. Physical designs of thermal scaffolding enable 12-tier 7nm ultra-dense 3D IC with max temperatures ≤125 degrees Celsius: an iso-footprint, iso-delay, 4x improvement in stacked tiers.**

Keywords—thermal management, 3D integrated circuits, back-end-of-line thermal conductivity

I. Introduction

Ultra-dense 3D integrated circuits (*3D ICs*), e.g., monolithic 3D ICs (Fig. 1), can yield large Energy-Delay Product (*EDP*) benefits for data-intensive applications when compared to traditional 2D computing systems [1,2]. To achieve these benefits, multiple tiers of logic and memory (e.g., thin layers of logic and/or memory devices, with associated signal/global metal routing) are integrated in 3D with ultra-dense (e.g., ≤ 100 nm pitch) vertical connections made using back-end-of-line (*BEOL*) inter-layer vias (*ILVs*) with limited aspect ratios [3]. Existing BEOL routing structures already use such nanoscale ILVs. 3D ICs become critical as fundamental limits to process technology miniaturization make traditional scaling paths more difficult. However, major thermal challenges must be overcome to enable high-speed and high-power computing engines on multiple 3D layers [4-5]. Without new techniques, upper-tier maximum temperatures in future 3D ICs will greatly exceed the upper limit required for reliable operation (e.g., 125°C in [6]).

We use the monolithic 3D IC in Fig. 1 to understand temperature rise and thermal dissipation in 3D layers (detailed analysis in Sec. III). Each of the N tiers in Fig. 1 contains one layer of high-speed and high-power silicon logic devices (e.g., computing engines) and BEOL layers (e.g., for signal routing) consisting of copper routing and ultra-low-κ inter-layer dielectric (ILD). Layers are electrically connected with ultra-dense ILVs. In some designs, silicon memory, memory access devices, and additional BEOL are also present on each tier. The 3D IC is externally cooled by an attached heatsink, which dissipates all heat generated to the ambient with a heat sink-specific heat transfer coefficient – $\boldsymbol{h}$ (W/m^2/K).

The maximum temperature $\boldsymbol{T_j}$ is determined by the heatsink, the ambient temperature, and the thermal characteristics of the N tiers. Heatsink innovations such as [7] can remove 1000 W/cm^2 with just 10°C heat rise across the heatsink (i.e., h=10^6 W/m^2/K), although the inlet water (ambient temperature) must be 100°C. However, even with such advanced heatsinks, $\boldsymbol{T_j}$ in the 3D IC will remain high if the effective thermal resistance across the 3D tiers is high. From our analysis (Sec. III), the thermal resistance across the tiers is the dominant (85%) contributor to $\boldsymbol{T_j}$ in a 3-tier 3D IC with such advanced heatsink. Reducing the sensitivity of the 3D IC to the thermal resistance across the tiers is therefore critical in enabling many-tier ultra-dense 3D ICs.

Existing approaches for addressing 3D thermal challenges include: **(1)** increasing BEOL thermal conductivity by exploiting dense high-conductivity metal in the power delivery network or adding dummy metal thermal vias (e.g., thermal-aware metallization [5, 8-9]), **(2)** changing the power distribution of the design by modifying its floorplan (e.g., thermal-aware floorplanning [10]), or **(3)** scheduling task execution to control temporal power profiles (e.g., thermal-aware scheduling [4]). Unfortunately, these techniques do not enable many tiers in 3D ICs, providing too little reduction in $\boldsymbol{T_j}$ (up to 56%) or requiring massive footprint overhead (78%) and reducing operating frequency (17%) to do so (Sec. IV).

We instead architect new ***thermal scaffolds*** (Fig 2) into 3D IC physical designs by **(1)** Fabricating the BEOL using both existing (ultra-low-κ) dielectric and new *thermal dielectric* (with 500× thermal conductivity and just 2× higher dielectric constant vs. existing BEOL dielectric) to achieve lateral thermal benefits without severely increasing signal capacitance, **(2)** placing *'pillar'* vertical heat conduction structures (densely packed BEOL metal vias) near

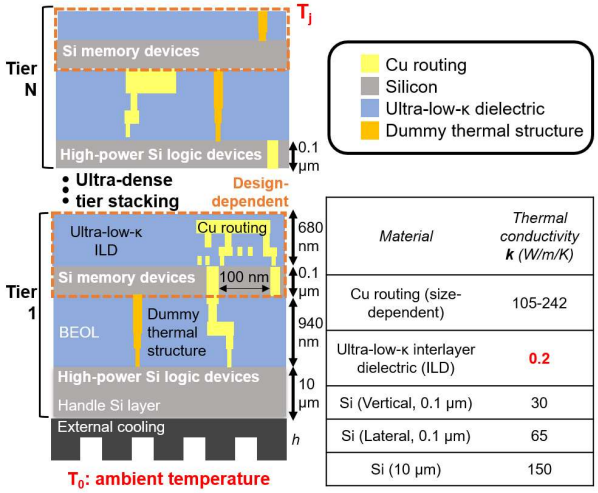

Figure 1. Ultra-dense monolithic 3D chip. The porous ultra-low-capacitance interlayer dielectric's thermal conductivity is estimated to be notably low based on [18] and [19] (details in Sec. II). Handle Si dimension from [12]. 3D Si layer thickness from [13]. 3D Si thermal conductivity from [14].

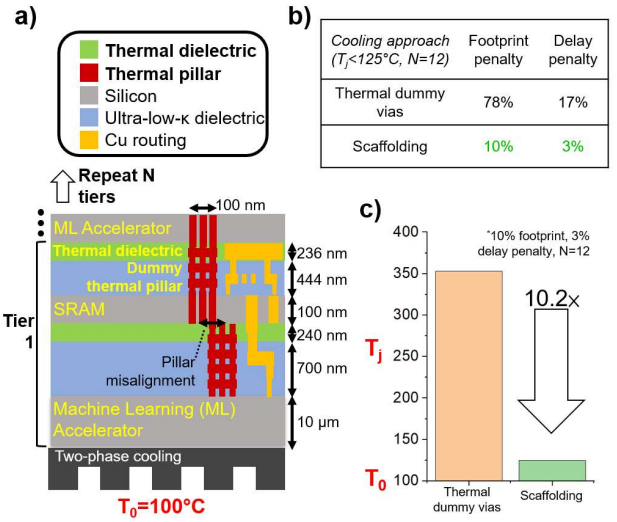

Figure 2. Our new thermal scaffolding approach to reduce the BEOL thermal resistance. (a) Lateral heat conduction using new thermal dielectric materials and vertical heat conduction with constellations of 'pillars' form the basis for scaffolding on a machine learning accelerator ultra-densely interleaved with SRAM cache (simulations in Sec. III, system parameters in Fig. 8). (b) Comparison to thermal dummy vias at 12 tiers and 125°C. Scaffolding achieves the same thermal benefits with 7.8× less footprint and 5.6× less delay penalty. (c) Scaffolding achieves a 10.2× reduction in T_j-T_0 compared to thermal dummy vias with the same 10% footprint and 3% delay penalties.

high-power floorplan elements between floorplanning and detailed place-and-route to improve vertical thermal conduction, and **(3)** integrating these structures into the power delivery network and routing the BEOL (Fig. 2b). Crucially, the *combination* of the thermal dielectric and pillars creates a *thermal scaffold* which efficiently conducts heat to the heatsink. We present experimentally verified models of a diamond-based thermal dielectric, use the ASAP7 7nm PDK [11] to implement a physical design flow with scheduling, floorplanning, and thermal dummy fill, and integrate the thermal dielectric in the PDK and place thermal pillars during floorplanning to realize the thermally scaffolded physical design (Sec. III). We find that (Sec. IV):

(1) Our *thermal scaffolding* approach enables thermal 3D scaling at low overhead: ***up to 4***× 3D IC tiers (with 4× power density) with just 10% footprint and 3% delay penalty before reaching T_j>125°C. Baseline techniques either have 78% footprint and 17% delay penalty to achieve the same 3D tiers or yield T_j>353°C at iso-footprint and iso-delay. The thermal benefits are *scalable* (demonstrated in a 100× scaled preliminary Fujitsu Research design) and *generalizable* (demonstrated in both a RISC-V core [15] and systolic array accelerator [16]) across designs.

(2) While more advanced heatsinks enhance its impact, thermal scaffolding provides benefits (between 4× and 1.25× tier scaling) across multiple heat sink technologies and maximum temperature constraints (85°C to 125°C).

(3) *Scaffolding-aware* software and circuit techniques (e.g., intelligent task scheduling and power gating) could further reduce the footprint penalty of scaffolding.

(4) The lateral thermal conductivity of the thermal dielectric increases the area cooled by a single pillar (Fig. 3), reducing the footprint penalty due to pillars from 34% (with no thermal dielectric) to 10% (with both thermal dielectric and ultra-low-κ dielectric).

(5) The new thermal dielectric opens a wide design space that enables effective scaffolding with macros (without thermal dielectric, a single 25μm×25μm hard macro contributes 15% of the maximum allowed temperature rise, reduced to 5% with the thermal dielectric) and heterogeneous tier designs that require pillar misalignment (Fig. 2a) (thermal dielectric increases tolerable

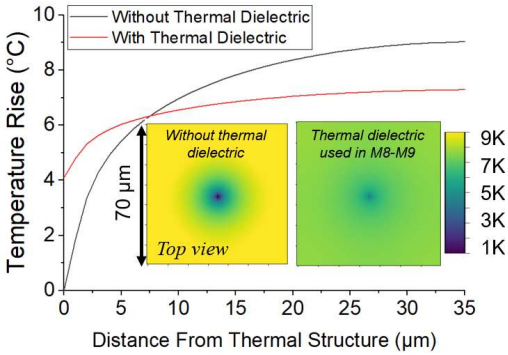

Figure 3. Pillar and 240 nm thick thermal dielectric (M8-M9 in ASAP7 PDK) reducing temperatures tens of μm away compared to pillar without thermal dielectric. Constant peak Gemmini systolic array power (95 W/cm^2) is dissipated throughout.

misalignment per tier from 300 nm to 1 μm within only 3°C of fully aligned temperature).

II. New Thermal Dielectrics

A *thermal dielectric* as described in Sec. I must have: (1) a very high thermal conductivity (**'high-k'**) to reduce BEOL thermal resistance and enable many ultra-dense 3D tiers, (2) a low relative permittivity (low dielectric constant or **'low-κ'**, today **κ**≈2 [17]) to reduce BEOL routing parasitic capacitance, and (3) fabrication compatibility with existing BEOL and ILV processes—in particular, low-temperature fabrication (≤400 °C) [1] in thin layers (~100-200 nm) with small feature sizes (~10's of nm) [7]. As today's ultra-low-κ ILDs are proprietary, we estimate thermal conductivity and dielectric constant from published trends: dielectric constant of **κ**≈2 [18] with thermal conductivity of **k**≈0.2 W/m/K (extracted from a meta-analysis of porous materials [19], since porosity is a dominant strategy to reduce dielectric constant **κ** [18]).

Polycrystalline diamond exhibits thermal conductivity between 100 and 1000 W/m/K [20]. The conventional deposition temperature for high-quality polycrystalline diamond through chemical vapor deposition (CVD) ranges from 600°C to 1000°C, far above the BEOL thermal budget of 400°C. Recently, a growth method has been developed that allows for 400°C deposition of diamond of comparable quality to conventionally grown diamond [21].

We verify the effectiveness of low-temperature-grown diamond as a thermal dielectric by developing a thermal conductivity model based on experimental data [21,22,23]. The film thermal conductivity of diamond grown to thicknesses of ~350nm, 600nm, and 1.9μm is fit to an effective thermal conductivity (*ETC*) model (Eq. 1) [24], by sweeping the grain-boundary thermal resistance (***R*** – m^{2}K/W). Here,

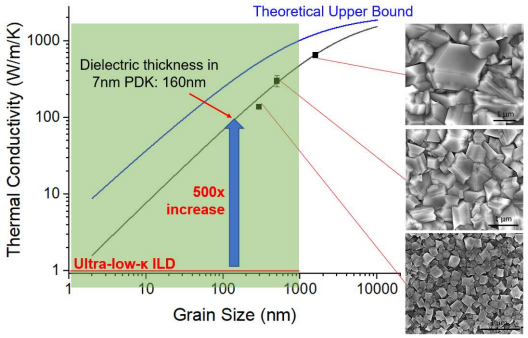

Figure 4. Modeled in-plane thermal conductivity of nanocrystalline diamond by grain size with experimental data overlay and SEM. Dielectric thickness in [11] taken from upper layers. 1900nm grain size is grown at 650°C [21], 650nm at 400°C [22] and 350nm at 500°C [23].

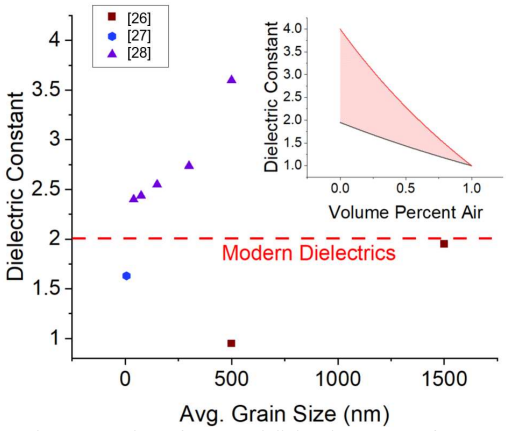

Figure 5. Literature review of measured dielectric constants for polycrystalline diamond with grain sizes comparable to scaffolding layer thickness and width. Inset shows theoretical benefit of added porosity to 500nm grain size diamond in [28] and 1500nm grain size in [25] with upper and lower bounds.

$\mathbf{k_0}$ is the single-crystal thermal conductivity (W/m/K), Λ_0 is the single crystal phonon mean-free-path (nm), and d is the grain size (nm).

$$k = \frac{k_0/(1+\Lambda_0/d^{0.75})}{1+R\cdot[k_0/(1+\Lambda_0/d^{0.75})]/d} \quad (1)$$

Fig. 4 gives our extracted model for the thermal conductivity of polycrystalline diamond. The thermal resistance at grain boundaries was determined to be 1.15 m^2K/GW. In-plane thermal conductivity for the thermal dielectric layer is swept between a minimum of 105.7 W/m/K representing an experimentally derived modeled thermal conductivity of a 160nm grain size film (the size of a single upper layer [7]) to a maximum of 500 W/m/K which represents a conservative experimentally derived estimate for a large-grained (>1um) thin film. Effective through-plane thermal conductivity is estimated as 30-105.7 W/m/K using the ETC approach [25] to account for thin film effects. The range is primarily controlled by the thermal boundary resistance of the film which is swept from an experimentally demonstrated maximum to an ideal minimum.

Nanostructured materials undergo surface bond contraction and have an increased surface to volume ratio which contributes to bandgap expansion at grain boundaries [26]. These combined effects cause nanocrystalline diamond films to have a suppressed dielectric constant compared to single crystal diamond. Certain deposition methods can also be used to introduce air gaps into diamond films further reducing the dielectric constant [27, 28] The effect of these inclusions on dielectric permittivity is modelled with the Maxwell-Garnet equation (f is the volume fraction of air in the porous material, ϵ_2 is diamond's relative dielectric permittivity, and ϵ_1 is the relative dielectric permittivity of free space):

$$\epsilon = \epsilon_2 \cdot \frac{\epsilon_1+2\cdot\epsilon_2+2\cdot f\cdot(\epsilon_1-\epsilon_2)}{\epsilon_1+\epsilon_2+f\cdot(\epsilon_1-\epsilon_2)} \quad (2)$$

Since porosity's impact (size and volume%) on diamond dielectric constant has not been experimentally measured, to estimate the dielectric constant of the diamond film utilized in scaffolding, the range of dielectric values accessible with increased film porosity is modelled with Eqn. 2 using a spread of dielectric constant values from literature of non-porous films (Fig. 4) [26,28]. We thus estimate a pessimistic dielectric constant of 4.

III. 3D IC Synthesis, Physical Design & Thermal Simulation

We implement two VLSI flows (Fig. 6) to demonstrate thermal scaffolding benefits using the derived properties of the thermal dielectric: (1) conventional 3D thermal and (2) scaffolding. Single-tier physical designs are stacked to form a 3D IC. Steps in the flow are described as categorized in Fig. 6: steps unique to thermal scaffolding (Sec. IIIA), steps unique to conventional 3D thermal (Sec. IIIB), and steps common to both flows (Sec. IIIC).

A. Thermal Scaffolding: We divide the design of thermal scaffolds in a 3D IC tier into three categories: thermal dielectric placement, pillar design, and pillar placement (Fig. 6).

Thermal dielectric placement: The thermal dielectric is only integrated in the upper layers of the BEOL (Fig. 2). In our designs, only the uppermost 240nm of the total BEOL thickness—two 80nm metal layers and one 80nm via layer—is fabricated with the thermal dielectric. These layers—M8, V8, and M9 in the ASAP7 PDK—are modified to reflect the dielectric constant of 4 from Sec. II.

Pillar design: The specific geometry of a single pillar is implemented using Cadence Innovus. The *add_stripe* command forms a set of vertically aligned rectangles (whose size we call the 'pillar footprint'), one on each interconnect layer, with maximum density vias between each adjacent pair. These structures are integrated with the power delivery mesh routing. The resulting pillars are exported and thermally simulated using COMSOL finite element analysis (Fig. 7). At 100nm×100nm pillar footprint, this conductivity is 105 W/m/K. This size is chosen to balance thermal conductivity reduction at small sizes [29] and the potential electrical and mechanical impacts of large pillars on surrounding transistors, which require further study.

Pillar placement: Pillars must be placed outside hard macro boundaries (e.g., SRAM blocks). We place pillars and macros during Innovus floorplan initialization, before detailed place-and-route (Fig. 6). The constraints that must be satisfied are (1) placement of macros within their respective floorplan units and (2) maximum temperature limit after thermal simulation with PACT [30]. For each heat source of area $\boldsymbol{A}$, (1) the optimistic case of uniform pillar covering is thermal simulated for increasing pillar count $\boldsymbol{P}$ until $\boldsymbol{T_j<T_{target}}$, giving the minimum thermally required pillar count $\boldsymbol{P_{min}}$ within that heat source, (2) The required pillar pitch is calculated as $\boldsymbol{(A/P_{min})^{0.5}}$, and macros are placed with gaps between them close to this pitch, and (3) $\boldsymbol{P_{min}}$ pillars are placed between the macro gaps and in a grid at the required pitch within the heat source. Depending on uniformity, temperature constraints may not be met: in this case, fill is increased past $\boldsymbol{P_{min}}$. Parameters of the PACT thermal simulation (e.g., number of tiers, $\boldsymbol{T_{target}}$, heatsink $\boldsymbol{h}$) are adjusted to trade off cooling and footprint/delay penalty. In the preliminary scaled Fujitsu Research design, this placement algorithm is run on a single multiply-accumulate, generating a pattern of pillars which is repeated across the MAC array and used in the thermal simulation.

B. Conventional 3D Thermal: The conventional 3D thermal flow incorporates thermal-aware metallization [5,8-9], floorplanning [10], and scheduling [4].

Thermal-aware metallization: We calibrate the Innovus timing-aware dummy fill and dummy via target fill fractions to fill statistics from a TSMC physical design using their confidential dummy fill algorithm. After calibration, our mean metal fill density is found to be within 5% of the results from the TSMC algorithm. To trade off area and delay penalties for more cooling, target placement fill density is decreased. This leads to less dense routing, which is shown to allow more room for thermal fill insertion (Fig. 7b). Finally, the BEOL is thermally simulated with COMSOL (Fig. 7a, details in Sec. IIIC) including dummy fill. Interconnect delay with dummy fill is extracted and used to in Innovus timing calculations (Sec. IIIC).

Thermal-aware floorplanning: A floorplan is generated by (1) duplicating the timing-driven single-tier starting floorplan generated by Innovus to multiple tiers and (2) performing thermal-aware floorplanning using a simulated annealing suite [31]. At each annealing step, the temperature is calculated using expected BEOL thermal properties and dimensions, functional unit areas, and power. The relative weights of peak temperature and total area in the cost function are swept from 100% area to 100% temperature. In a 4-tier core, the area at 100% temperature weighting is 16% more than the area at 100% area weighting. To remain close to the operating

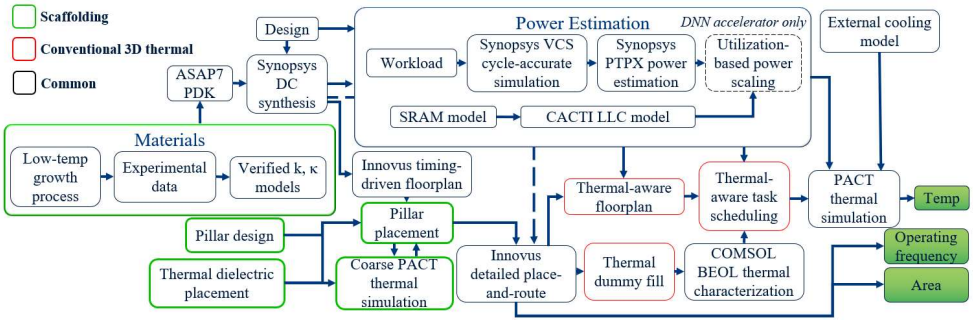

Figure 6. (a) 3D IC synthesis, physical design, and thermal simulation framework.

frequency from the original timing-driven floorplan, total wirelength increase is kept within 5%.

Thermal-aware scheduling: In this work, an N-tier design has N copies of the design. Each of these N copies are ranked by effective thermal resistance by simulating each in PACT with all others turned off: copies with higher peak temperatures have higher thermal resistance. Copies with the lowest thermal resistances are then assigned the highest-power tasks. This mimics thermal-aware task assignment of known workloads performed in real systems. Similar results could be achieved by dynamic swapping [4].

C. Common Flow Steps: Finally, we describe the VLSI flow for tiers of the 3D IC and thermal simulation steps for the full 3D IC.

Designs: We analyze three designs: (1) a Gemmini systolic-array-based design [16] with an interleaved 3D SRAM last-level cache (LLC) (Fig. 2b) to take advantage of large ultra-dense 3D benefits [1], (2) a preliminary design from Fujitsu Research with estimates of systolic array power, LLC and scratchpad memory capacity, footprint, and operating frequency scaled to modern workloads, and (3) a RISC-V multi-core Rocket SoC [15] design.

Synthesis and place-and-route: To derive the power density of the functional units of the Rocket and Gemmini designs, we synthesize using Synopsys Design Compiler (DC). Target period is swept from 0.1 ns to 2 ns in both designs. In our un-optimized designs, synthesis does not complete below 0.7 ns (Rocket) and 0.9 ns (Gemmini). Additional area savings (due to reduced buffer count and cell size) of 10% are realized by further increasing from this minimum period to targets of 0.8 ns and 1 ns. Place-and-route with modifications based on the cooling strategy is performed using Cadence Innovus. The area and delay (sum of target period and worst negative slack) are recorded.

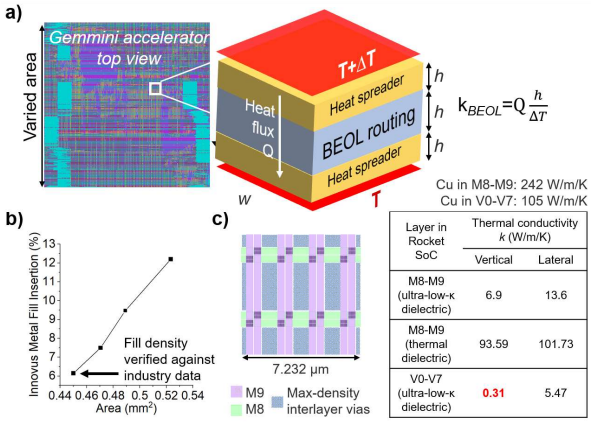

Layer in Rocket SoC	Thermal conductivity k (W/m/K)	
	Vertical	Lateral
M8-M9 (ultra-low-κ dielectric)	6.9	13.6
M8-M9 (thermal dielectric)	93.59	101.73
V0-V7 (ultra-low-κ dielectric)	0.31	5.47

Figure 7. BEOL thermal conductivity estimation. (a) Slice of physical design is chosen to have same density as overall design and simulated in COMSOL to estimate k_{BEOL} as in [5]. This is performed for both routing in lower layers (V0-V7, w=1µm) and upper layers (w=7.232µm, M8-M9). Copper thermal conductivities are a function of dimension [29]. Thermal boundary conductance is found to be negligible at 10^9 W/m²/K [34]. (b) Timing-aware Innovus metal fill insertion increases with area slack, allowing for better cooling. (c) Upper-layer power delivery network with inserted interlayer vias and power routing based on densities in [8].

Power estimation: The memory-bound spmv benchmark [32] is chosen for the Rocket core to be representative of workloads benefiting from ultra-dense 3D IC's high memory-to-compute bandwidth [1]. Matrix multiplication is run on the systolic array. We simulate activity of the design using the Synopsys Verilog Compiler Simulator (VCS) and estimate each functional unit's maximum power with Synopsys PrimePower. Systolic array power is scaled from 72% (the maximum utilization in the simulated workload) to 100% to estimate a worst-case. Finally, the power and bandwidth of the 3D SRAM LLC is calculated using FinCACTI [33]. The power in the preliminary Fujitsu Research design is estimated with proprietary internal simulations.

Thermal simulation of BEOL: We implement standard abstraction methods for chip-scale thermal simulation, lumping BEOL sections into homogenous models with thermal conductivities calculated using COMSOL finite element analysis. The upper metal layers M8-M9 are modeled separately from the rest of the BEOL (V0-V7), which is shown in [5] to be necessary for accuracy within 5%. The upper layers use ultra-low-κ dielectric in the conventional 3D cooling flow and thermal dielectric in the scaffolding flow. We calculate the average metal density of the physical design, select a BEOL slice within 1% of that value, and simulate with the methods and equation in Fig. 7a [5]. Ultra-dense

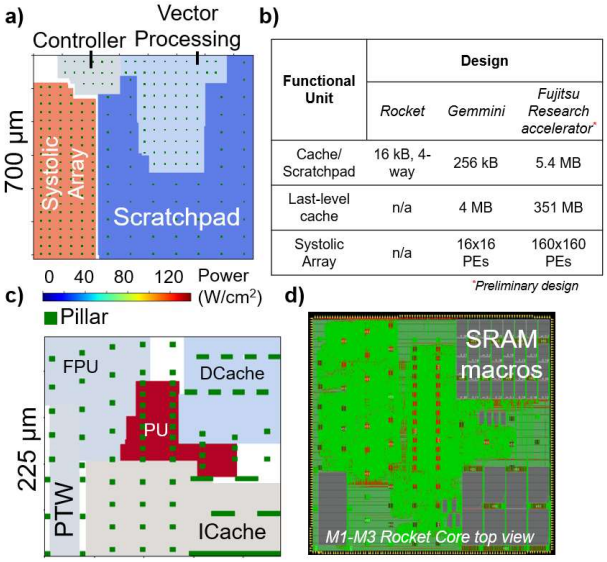

Functional Unit	Design		
	Rocket	*Gemmini*	*Fujitsu Research accelerator**
Cache/ Scratchpad	16 kB, 4-way	256 kB	5.4 MB
Last-level cache	n/a	4 MB	351 MB
Systolic Array	n/a	16x16 PEs	160x160 PEs

*Preliminary design

Figure 8. Overview of designs (a) Floorplan and power density of scaffolded Gemmini accelerator with pillar overlay. (b) Sizes of memories and array of designs. (c) Floorplan and power density of scaffolded Rocket core. Each Rocket core contains a pipelined processing unit (PU), instruction cache (ICache) and data cache (DCache), page table walker (PTW), and floating-point unit (FPU). (d) Physical design of scaffolded Rocket core.

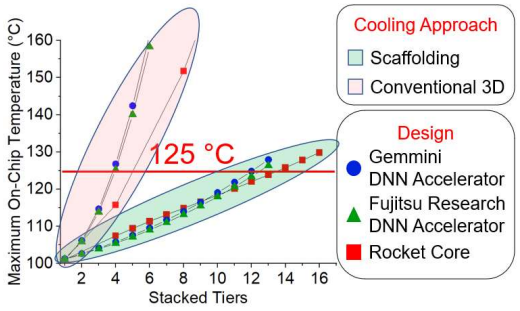

Figure 9. (a) Scaling of peak temperature with number of tiers in three different designs for conventional 3D cooling and scaffolding, both using a porous two-phase heatsink. An example design point at 2.8% delay penalty and 10% area penalty is chosen for fair comparison across cooling strategies.

ILVs are inserted in the upper layers with power delivery network densities based on [8] (Fig. 7c). COMSOL methods are verified by replicating experimentally measured conductivities [35] within 5%.

Thermal simulation of 3D IC: BEOL thermal properties, 3D IC floorplan, estimated power, and heatsink properties are integrated into the final thermal simulation with PACT, with results cross-referenced against COMSOL and Cadence Celsius. Another thermal property abstraction is performed as in Fig. 7 on a section of the 3D IC with pillars. Using these thermal conductivities, we simulate the full 3D IC and extract a temperature map.

IV. Results

We make the following key observations:

Observation 1: *Thermal scaffolding enables 3D IC thermal scaling of up to 3-4× higher tier counts (with corresponding increase in power) across multiple architectures and scalable to much larger designs.* Just 3 tiers (159 W/cm^2) of stacked Gemmini accelerators are possible with conventional 3D thermal, whereas 12 tiers (636 W/cm^2) are possible with thermal scaffolding. This 4× benefit comes at just 10% footprint and 3% delay penalties (Fig. 9). Scaffolding supports tradeoffs between the degree of 3D cooling vs. area and delay penalties (Fig. 10) by altering scaffolding pillar placements as in Sec. IIIA. With scaffolding, Rocket achieves 13 tiers at 10.6% footprint and 2.6% delay penalties. The Fujitsu Research accelerator (Fig.8b) achieves 12 tiers at 9.4% footprint penalty.

Observation 2: *Conventional 3D cooling cannot achieve equivalent cooling to scaffolding without large area and delay penalties.* As the area and delay penalties are increased to 10%, conventional 3D cooling can only achieve 4 tiers (Fig. 10). To achieve 12 tiers with conventional 3D thermal, 78% area and 17% delay penalty is required in the Gemmini design. Similar penalties are observed in other designs (Table I).

Table I. Comparison of Cooling Strategy Penalties Across Designs At Near-Constant Scaffolding Penalty

	Conventional 3D Thermal			*Vertical Conduction Only*			*Scaffolding*		
Design	***A***	***B***	***C***	***A***	***B***	***C***	***A***	***B***	***C***
Footprint penalty (%)	78	69	74	34	25	30	10	10.6	9.4
Delay penalty (%)	17	13	n/a	7	7	n/a	3	2.6	n/a

A: Gemmini, B: Rocket, C: Fujitsu Research (preliminary, no timing).

Observation 3: *Improved heatsinks enhance thermal scaffolding benefits. Thermal scaffolding yields large benefits for a range of emerging heatsinks.* With two-phase cooling (requiring 100°C ambient [7]), 4× tier scaling is achieved in the Gemmini 3D IC. With Si-integrated microfluidics (10× reduced ***h*** [36]) and room-temperature water, scaffolding can achieve 8 tiers with T_j<125°C from 5 tiers with conventional 3D thermal. The conventional 3D thermal tier count is larger because room temperature water, instead of boiling water, can be used. Single-phase cooling also enables thermal limits below 100°C, which may be required by some technologies. With T_j<85°C, a microfluidically cooled and scaffolded Gemmini 3D IC achieves 5 tiers vs. 3 with conventional 3D thermal (Fig. 11). Benefits are smaller than with the two-phase heatsink case because total heat removal is limited by the heatsink.

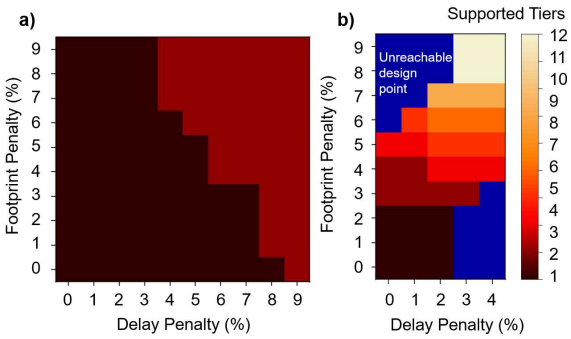

Figure 10. Fine grained area-delay penalties of (a) conventional 3D thermal and (b) scaffolding.

Observation 4: *The thermal dielectric gives flexibility in pillar placement, resulting in (a) reduced penalties due to pillars, (b) insertion of larger hard macros, and (c) lower penalty for pillars on adjacent tiers.* The thermal dielectric mitigates physical design challenges when pillar placement is constrained (e.g., by hard macros). *(a)* The thermal dielectric in the Gemmini 3D IC reduces penalties of achieving 12 tiers from 34% to 10% footprint and 7% to 3% delay compared to the ultra-low-κ dielectric. (Table I). *(b)* Replacing ultra-low-κ dielectric with thermal dielectric underneath a 25µm×25µm macro with four surrounding pillars in the 6-tier Gemmini can reduce temperature rise from 15°C to 5°C. *(c)* Without the thermal dielectric, the nearest pillar on adjacent tiers must be within 300nm to stay within 3°C heat rise per tier. Using the thermal dielectric between tiers increases this alignment tolerance to 1µm.

Observation 5: *Scaffolding creates new opportunities for co-design with software-, architecture- and circuit-level techniques (e.g., power gating, task scheduling) techniques to significantly increase the number of 3D tiers and/or reduce the footprint/delay penalties associated with 3D heat removal.* An analysis of these opportunities in a real system requires physical design techniques for power gating and scaffolding-aware scheduling—beyond the scope of this work. A motivating example of fine-grained heat sources, constrained so that only one is active at a time (Fig. 12a), represents individually scheduled and gated multiply-accumulate

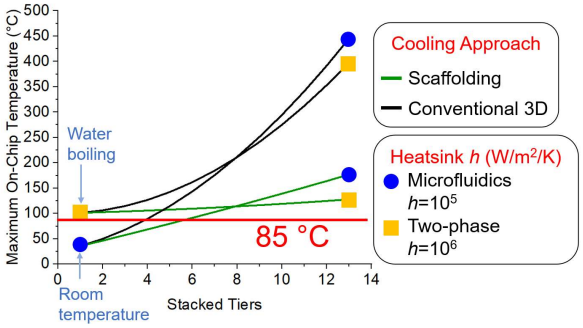

Figure 11. Exploration of Gemmini with microfluidic heatsink [36] and two-phase heatsink [7]. Water boiling requires an ambient temperature of 100°C, but the more efficient heatsink still gives lower peak temperature above 8 tiers. However, if lower peak temperatures (e.g., 85°C) are required, scaffolding still gives benefits in the microfluidic cooling case.

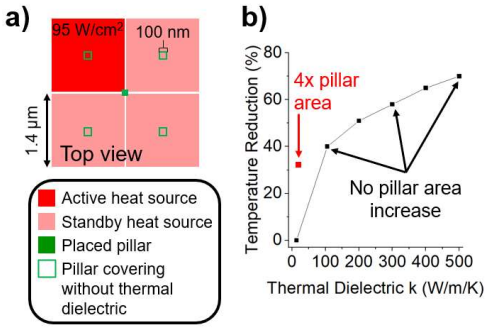

Figure 12. (a) Toy example to demonstrate the benefits of co-design of scaffolding with power gating and task scheduling. Only one of the four heat sources is active at a time: this may be achievable by software advances. (b) Benefit of the thermal dielectric.

units. With the thermal dielectric, a single pillar can cool all four heat sources, reducing the peak temperature more than a floorplan with 4× the pillar count and no thermal dielectric (40% vs 32%, Fig. 12b). Benefits increase above 70% as the thermal dielectric thermal conductivity improves (Fig. 12), all at 75% less area penalty.

V. Conclusion

We simulate our proposed thermal scaffolding, incorporating new hardware-calibrated diamond-based dielectrics and novel thermal pillar physical design techniques, enabling 12 tier ultra-dense 3D 7nm cores and DNN accelerators with just 125°C peak temperature: a 3-4× increase (iso-footprint and iso-delay) over conventional 3D thermal techniques. Scaffolding is the first technique to thermally enable ultra-dense 3D IC designs with numerous tiers of stacked compute. Scaffolding provides many future co-design opportunities, including fine-grained power-gating and scheduling techniques to reduce overheads by 75%.

Acknowledgment

The authors would like to acknowledge the support of TSMC, the DARPA 3DSoC program, Stanford SystemX Alliance, Precourt, the NSF-GRFP under Grant No. DGE-1656518, and the Alfred P. Sloan Research Fellowship.

References

[1] M. M. Sabry Aly *et al.*, "The N3XT approach to energy-efficient abundant-data computing," *Proc. IEEE* vol. 107, pp. 19-48, Jan 2019.
[2] R. M. Radway *et al.*, "The future of hardware technologies for computing: N3XT 3D MOSAIC, Illusion scaleup, co-design," *IEDM 2021*, pp. 25.4.1-25.4.4.
[3] T. Srimani *et al.*, "Heterogeneous integration of BEOL logic and memory in a commercial foundry: multi-tier complementary carbon nanotube logic and resistive RAM at a 130 nm node," *VLSI 2020*, pp. 1-2.
[4] A. K. Coskun, J. L. Ayala, D. Atienza, T. S. Rosing, and Y. Leblebici, "Dynamic thermal management in 3D multicore architectures." *DATE* 2009.
[5] H. Wei, T. F. Wu, D. Sekar, B. Cronquist, R. F. Pease, and S. Mitra, "Cooling three-dimensional integrated circuits using power delivery networks," *IEDM 2012*.
[6] F.P. McCluskey, T. Podlesak, and R. Grzybowski, "High temperature electronics (1st ed.)," *CRC Press*, 1997.
[7] J. W. Palko *et al.*, "High heat flux two-phase cooling of electronics with integrated diamond/porous copper heat sinks and microfluidic coolant supply," *ITherm 2016*, pp. 1511-1517.
[8] S. K. Samal, K. Samadi, P. Kamal, Y. Du, S. K. Lim, "Full chip impact study of power delivery network designs in gate-level monolithic 3-D ICs," *IEEE T Compur. Aid D.* 35(10), 2016.
[9] T. Y. Chiang, K. Banerjee, and K. C. Saraswat, "Compact modeling and SPICE-based simulation for electrothermal analysis of multilevel ULSI interconnects," *ICCAD* 2001.
[10] S. K. Samal, S. Panth, K. Samadi, M. Saedi, Y. Du, S. K. Lim. "Fast and accurate thermal modeling and optimization for monolithic 3D ICs," *DAC* 2014, pp. 1-6.
[11] L.T. Clark *et al.*, "ASAP7: A 7-nm FinFET predictive process design kit," *Microelectronics Journal*, vol. 53, pp. 105-115, July 2016.
[12] J. P. Gambino, "Thin silicon wafer processing and strength characterization," *IPFA,* pp. 199-207, 2013.
[13] P. Batude *et al.*, "3D monolithic integration," *ISCAS 2011*, pp. 2233-2236.
[14] C. Jeong, S. Datta, M. Lundstrom, "Thermal conductivity of bulk and thin-film silicon: A Landauer approach," *J. Appl. Phys.* 111(9), 2012.
[15] K. Asanović *et al.*, "The Rocket Chip generator," *UCB/EECS Tech. Rep.,* 2016.
[16] H. Genc *et al.*, "Gemmini: Enabling systematic deep-learning architecture evaluation via full-stack integration," *DAC 2021*, pp. 769-774.
[17] M. Xie, M. Li, Q. Sun, W. Fan, S. Xia, W. Fu, "Research progress on porous low dielectric constant materials," *Mater. Sci. Semicond.*, 139, 106320, 2022.
[18] M.H. Lee and W.S. Shue, "The overview of current interconnect technology challenges and future opportunities," *IEDM* (2020).
[19] A.M. Evans *et al.*, "Thermally conductive ultra-low-k dielectric layers based on two-dimensional covalent organic frameworks," *Nat. Mater.,* 20, 1142-1148 (2021).
[20] M. Malakoutian et al., "Diamond Integration on GaN for Channel Temperature Reduction," *WiPDA* 2021.
[21] M. Malakoutian *et al.*, "Record-Low Thermal Boundary Resistance between Diamond and GaN-on-SiC for Enabling Radiofrequency Device Cooling" *ACS Appl. Mater. Interfaces*, 13(50), p. 60553-60560, 2021.
[22] M. Malakoutian *et al.*, "Low Thermal Budget Growth of Near-Isotropic Diamond Grains for Heat Spreading in Semiconductor Devices," *Adv. Funct. Mater*., 32, p. 2208997, 2022.
[23] M. Malakoutian, C. Ren, K. Woo, S. Pasayat, S. Chowdhury, "The Impact of Polycrystalline Diamond Growth Nucleation on N-Polar GaN Channel Performance," *MRS Fall Meeting*, Boston, MA, 2020.
[24] H. Dong, B. Wen, and R. Melnik, "Relative importance of grain boundaries and size effects in thermal conductivity of nanocrystalline materials," *Sci Rep*, 4, p. 7037, 2014.
[25] J. Alvarez-Quintana, J. Rodríguez-Viejo, F.X. Alvarez, and D. Jou, "Thermal conductivity of thin single-crystalline germanium-on-insulator structures," *Int. J. Heat Mass Transf.*, 54(9–10), p. 1959-1962, 2011.
[26] H. Ye, C. Q. Sun, and P. Hing, "Control of grain size and size effect on the dielectric constant of diamond films" *J. Phys. D Appl. Phys.,* 33(23), 2000.
[27] H. Sakaue, N. Yoshimura, S. Shingubara, and T. Takahagi , "Low dielectric constant porous diamond films formed by diamond nanoparticles," *Appl. Phys. Lett.,* 83, p. 2226-2228, 2003.
[28] Z. L. Wang et al., "Effect of grain size and pores on the dielectric constant of nanocrystalline diamond films" *Appl. Phys. Lett.,* 90(13), p. 133118, 2007.
[29] J.M Lugo and A.I. Oliva, "Thermal properties of metallic films at room conditions by the heating slope," *J. Thermophys. Heat Transf.*, 30(3), 2016.
[30] Z. Yuan, P. Shukla, S. Chetoui, S. Nemtzow, S. Reda, and A.K. Coskun, "PACT: An extensible parallel thermal simulator for emerging integration and cooling technologies," *IEEE TCAD,* 41.4, pp. 1048-1061, May 2021.
[31] J. Knechtel *et al.*, "Corblivar", 2017. github.com/IFTE-EDA/Corblivar
[32] Y. Lee *et al.*, "riscv-tests", 2017. github.com/riscv-software-src/riscv-tests
[33] A. Shafaei, Y. Wang, X. Lin, and M. Pedram, "FinCACTI: Architectural analysis and modeling of caches with deeply-scaled FinFET devices," *VLSI 2014,* pp. 290-295.
[34] O. D. Restrepo, D. Singh, M. Rabie, P. Paliwoda, and E. C. Silva, "Ab initio electrical, thermal conductance, and lorenz numbers for advanced CMOS interfaces," *IEEE Trans. Electron Devices*, 69(5), 2022.
[35] E. Bury, B. Kaczer, S. Van Beek, and D. Lintern, "Experimental extraction of BEOL composite equivalent thermal conductivities for application in self-heating simulations," *ESSDERC* 2018.
[36] D. B. Tuckerman, R. F. Pease, "Higher-performance heat sinking for VLSI," *IEEE Electron Device Lett.*, 2.5, 1981, pp. 126.