

PowerCool: Simulation of Cooling and Powering of 3D MPSoCs with Integrated Flow Cell Arrays

Artem Andreev, Arvind Sridhar, *Member, IEEE*, Mohamed M. Sabry, *Member, IEEE*, Marina Zapater, *Member, IEEE*, Patrick Ruch, *Member, IEEE*, Bruno Michel, *Member, IEEE* and David Atienza, *IEEE Fellow*

Abstract—Integrated Flow-Cell Arrays (FCAs) represent a combination of integrated liquid cooling and on-chip power generation, converting chemical energy of the flowing electrolyte solutions to electrical energy. The FCA technology provides a promising way to address both heat removal and power delivery issues in 3D Multiprocessor Systems-on-Chips (MPSoCs). In this paper we motivate the benefits of FCA in 3D MPSoCs via a qualitative analysis and explore the capabilities of the proposed technology using our extended PowerCool simulator. PowerCool is a tool that performs combined compact thermal and electrochemical simulation of 3D MPSoCs with inter-tier FCA-based cooling and power generation. We validate our electrochemical model against experimental data obtained using a micro-scale FCA, and extend PowerCool with a compact thermal model (3D-ICE) and subthreshold leakage estimation. We show the sensitivity of the FCA cooling and power generation on the design-time (FCA geometry) and run-time (fluid inlet temperature, flow rate) parameters. Our results show that we can optimize the FCA to keep maximum chip temperature below 95 °C for an average chip power consumption of 50 W/cm² while generating up to 3.6 W per cm² of chip area.

Index Terms—3D MPSoCs, Thermal Modeling, Liquid Cooling, Electrochemical Flow Cell.

1 INTRODUCTION

THE steadily growing computational capacity of microprocessors, as described by Moore’s law [1], had traditionally been accompanied by power scaling, which is also known as Dennard scaling [2]. However, as we penetrate in the deep sub-micron era, we have witnessed a stagnation of Dennard scaling. Approaching the threshold voltage limit, voltages cannot be scaled any more, and leakage becomes significant. To keep increasing overall performance, architectural designs shifted from single- to multi-core processors and Multi-Processor Systems-on-Chips (MPSoCs). Still, this approach suffers from the same fundamental limitations, resulting in limitations on the MPSoC utilization for a fixed power budget, i.e. *dark silicon*.

3D stacking of silicon chips [3] is one of the approaches to achieve increased computing capacity in MPSoCs. Vertical integration of microprocessor units (MPUs) and main memory into 3D MPSoCs can alleviate constraints related to the *memory wall* [4–7]. Area array interconnects using through-silicon vias (TSVs) between vertically integrated processor and memory dies, for instance, provide means for high-speed wide I/O communication and can improve the computational performance by an order of magnitude [8]. It also reduces the energy required for communication (vis-a-vis computation) between devices.

However, stacking multiple core dies is currently constrained mainly due to two issues. The first one is the increased thermal resistance within a 3D stack and the compounded footprint heat fluxes (which may exceed 100 W/cm²). Interlayer liquid-cooling technology for 3D stacks using silicon microchannels carrying single- or two-phase fluids has proved to be a highly efficient solution to this problem [9]. This technology is compatible with TSVs and can be scaled according to the number of dies.

However, interlayer liquid cooling does not address the second major challenge in 3D integration, which is the powering of the vertically stacked IC dies. Reliable power delivery is already a challenge for 2D ICs with 80% of the chip-to-package connections being used for power [10] to minimize power loss and maintain a uniform voltage distribution. In a 3D MPSoC, it is anticipated that an even larger fraction of chip-to-package I/Os as well as TSVs within the stack must be dedicated to power delivery, which inhibits the exploitation of the TSVs as wide I/O communication channels, diminishing the potential bandwidth rewards of 3D integration. Proposed technologies to overcome this “*power wall*” include dual-side power delivery [11] and on-chip voltage conversion [12], but vertical scalability is limited in both cases.

On-chip power generation can address both these challenges simultaneously. A liquid medium providing both effective cooling and power delivery for computing elements, resembling the function of blood in biological systems, can potentially allow even denser integration of computational circuits with greater bandwidth [8]. Such “*electronic blood*” can be a couple of electrolytic solutions flowing through a microscale electrochemical cell and participating in an electrochemical reaction to produce electricity.

The technology of integrated *Flow Cell Arrays* (FCAs) is

- A. Andreev, M. Zapater and D. Atienza are with the Embedded Systems Laboratory at the École Polytechnique Fédérale de Lausanne. E-mail: artem.andreev@epfl.ch
- M. Zapater is with DACYA at Universidad Complutense de Madrid.
- A. Sridhar, P. Ruch and B. Michel are with IBM Research Zurich.
- M. M. Sabry is with the Robust Systems Group at Stanford University.

Manuscript received ; revised .

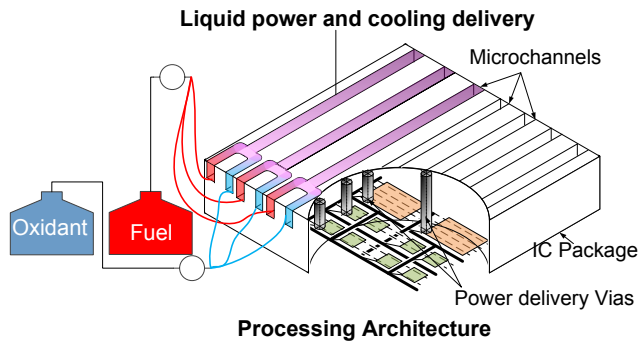


Fig. 1: Concept of integrated microfluidic power supply and cooling for MPSoCs [14].

a realization of the aforementioned concept of combined electrochemical power delivery and integrated liquid cooling of chips [8, 13]. A system-level investigation of this technology was published in [14]. As shown in Fig. 1, several microchannels etched on top of the silicon die are electrically connected in parallel to constitute a microfluidic FCA. The proposed FCA technology can be integrated in the geometry of 3D stacks in the same way as microfluidic cooling systems. PowerCool — a compact model to simulate the electrochemical performance of the FCA — was first advanced in our previous work [15] and validated against fine-grained (and much slower) simulations performed using the commercial COMSOL Multiphysics software. Simulations using PowerCool show that although the power demand of high-performance MPSoCs is significantly higher than the attainable FCA power generation (which were of the order of $\sim 1 \text{ W/cm}^2$), it is possible to power the on-chip higher-level caches.

In this paper, we extend the original PowerCool to achieve a further understanding of flow cell arrays and their benefits for continued 3D integration of MPSoCs, while enabling architectural exploration. Our specific contributions are as follows:

- We present a qualitative analysis of possible 3D stacks configurations to motivate the FCA implementation in 3D MPSoCs.
- We extend the original PowerCool with a compact thermal model (3D-ICE) and subthreshold leakage estimation model to create a new tool for system-level simulations of 3D MPSoCs with interlayer FCAs that enables architectural exploration.
- We validate the compact electrochemical model of FCA against experimental data obtained using a microscale FCA.
- Using our new PowerCool simulator, we analyze the FCA cooling and power generation dependencies on the static (FCA geometry) and dynamic (fluid inlet temperature and flow rate) parameters. We obtain an FCA configuration able to generate 3.6 W of power per cm^2 of chip area while cooling a 3D stack with a power consumption of 50 W/cm^2 below 95°C .
- We discuss the potential and impact of beyond state-of-the-art FCA technologies in power delivery and heat removal.

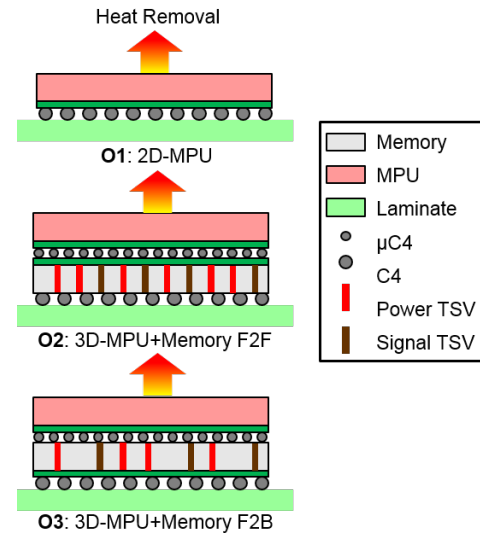


Fig. 2: Conventional 2D and 3D stacks.

2 INTER-TIER FCAs IN 3D MPSoCs

The integrated FCA technology (concept illustrated in Fig. 1) has the potential to provide uniform heat removal and power delivery to all layers of a 3D stack — a key enabler for vertical scalability in 3D MPSoCs. In order to build the PowerCool model to simulate and design systems with FCAs, we must first qualitatively ascertain the impact of this technology on realistic 3D MPSoC architectures. Hence, in this section we:

- 1) perform a qualitative analysis of 3D MPSoCs with integrated FCA to study its impact on power delivery and heat removal. Various stacking configurations/permutations are studied to determine which configuration has the best performance in both these aspects. This configuration will be used later in Section 3 and 4 to build the PowerCool model and perform detailed quantitative study.
- 2) present a realistic 3D MPSoC architecture based on IBM POWER8 microprocessor, which will be used as the power- and thermal-load for the PowerCool model.

2.1 Qualitative analysis of 3D MPSoC configurations with FCAs

In order to perform a qualitative analysis of the impact of FCAs on 3D MPSoC interconnects, we compare them with conventional 2D and 3D chip architectures without microfluidic power delivery (O1, O2 and O3, as shown in Figure 2). In each of these cases, heat is dissipated from the top of the chip (or chip stack) using a conventional air-cooled heat sink or direct-attach cold plate. For the O1 configuration, a single MPU is connected to a laminate via C4 solder connections, which provides the high bandwidth signal I/Os off-chip to a separate memory die (not shown) as well as the very high number of power I/Os required by the MPU.

Then, O2 and O3 represent two possible 3D stack configurations consisting of MPU and a memory die: Face-to-Face

(F2F) and Face-to-Back (F2B). The orientation of each die is indicated using the green line that represents the active region and metalization layers. The MPU is always the top chip (TC) to satisfy the heat-removal requirements of the 3D stack. Both the O2 and the O3 configuration have 3 levels of package interconnects: (1) C4 to the laminate, (2) TSVs in the bottom chip (BC) and (3) micro-C4s between BC and TC. The designs are evaluated based on the following scoring technique.

2.1.1 Qualitative Scoring

Qualitative scores are given to both the heat-removal capability and the number of power connections needed in all the interconnect levels for all topologies. The qualitative scores for heat removal are defined by counting the number of levels/layers heat has to flow from the source to the heat sink (die thicknesses and C4/micro-C4 layers) — a measure of “thermal distance”. Since the MPUs have a much higher heat flux density than memories, the thermal distance between MPUs to the heat sink is given double the value of that of the memory dies to give them a qualitative difference. Adding up the thermal distances in each design gives an indication of the relative difficulty in heat removal in each design. The lower the score, the better the heat removal.

The qualitative scores for the power interconnects are calculated by counting the number of power interconnects in each layer of the 2D/3D chip stack. This is then normalized with respect to the total number of interconnects that is physically possible to be realized in a 3D chip stack. This normalization ensures a fair comparison between designs of different physical sizes. The number of power interconnects for MPUs are counted double that for memory chips. This is to resolve a qualitative difference between the two. It is assumed that the signal interconnects needed between the MPU and the memory chip, and the signal interconnects needed between the 3D stack and devices outside of the package are determined by the bandwidth requirements, and are fixed. These interconnects are counted as a single unit each. The total *possible* number of interconnects in a 2D/3D chip stack is calculated by assuming that each layer is populated by all possible interconnects, i.e.: 1) power interconnects for MPU (2 units), 2) power interconnects for memory (1 unit), 3) signal TSVs (1 unit), which represent a maximum total of 4 units per layer that carries interconnects. Since in all our configurations, the TC is always facing down, it never contains any TSVs/interconnects in that layer and is not counted. For example, design O1 has one layer of interconnect (C4), hence, a maximum possible total of 4 units; designs O2/O3 have three layers of interconnect (C4, BC, μ C4), hence, a maximum possible total of 12 units, and so on. Similarly, the lower the score, the better the design qualitatively.

The above scoring technique is applied for all designs considered in this paper. The resulting scores are plotted in a design-space graph in Figure 3. Here, the x-axis represents the progression from worse to better heat removal capabilities and the y-axis represents the progression from more to less fraction of the total number of electrical interconnects dedicated to power delivery. Then, favorable designs are

TABLE 1: Qualitative evaluation of Conventional 2D and 3D stacks without FCA. The number score is explained in Section

Configuration		O1: 2D-MPU	3D-MPU+Memory	
			O2: F2F	O3: F2B
Heat removal		1	2	3
TC	Type	MPU	MPU	MPU
μ C4	Signal	-	1	1
	Power	-	2	2
BC	Type	-	Mem	Mem
	Signal TSV	-	1	1
	Power TSV	-	3	2
C4	Signal	1	1	1
	Power	2	3	3

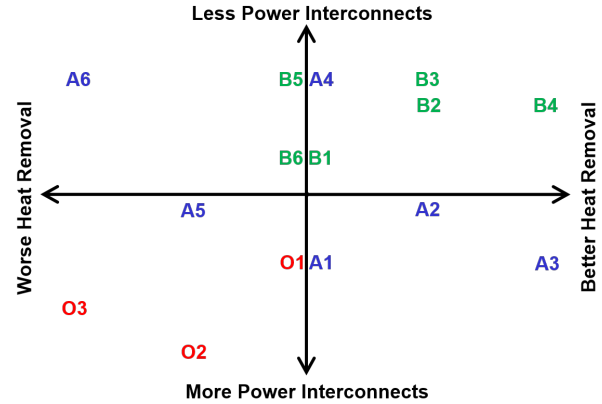


Fig. 3: Qualitative evaluation of the various 3D topologies; **Case O: Conventional 2D/3D MPSoCs** in Fig. 2, **Case A: 3D MPSoCs with SoA FCAs** in Fig. 4 and **Case B: 3D MPSoCs with beyond-SoA FCAs** in Fig. 12.

those which lie towards the top-right corner of this design space.

2.1.2 Case O: Conventional 2D/3D MPSoCs

The scores for the conventional O cases are plotted in red in Fig. 3. We consider the O1 case, which is the most conventional 2D chip with air-cooled heat sinks, to be the normalized benchmark of our analysis in the middle of the x-axis. The back side heat removal performs poorly in the case of O2 and O3 because of the compounded heat flux due to the addition of the memory die and its increased distance to the heat sink. With regard to interconnects, since the bandwidth needed between MPU and memory is typically much higher than that needed between memory and laminate, the O2 configuration requires less signal TSVs in the BC. However, the BC needs to support very high number of power TSVs to supply both the MPU and the memory. The situation is the opposite for O3. The exact scores are also provided in Table 1 for reference.

2.1.3 Case A: 3D MPSoCs with State-of-the-Art FCAs

Next, an FCA die to serve both as a power source and as a heat sink is introduced and the resulting possible 3D stack topologies are explored. For the moment, we only consider the current state-of-the-art FCAs that can deliver power current densities of the order of $1\text{W}/\text{cm}^2$ — sufficient only to power current memory devices. We call these architectures

TABLE 2: Qualitative evaluation of 3D stacks with FCA. The number score is explained in Section

Configuration		A1	A2	A3	A4	A5	A6	B1	B2	B3	B4	B5	B6
Heat Removal		2	1	0	2	3	4	2	1	1	0	2	2
TC	Type	MPU	Mem	MPU	Mem	FCA	FCA	MPU	Mem	MPU	Mem	FCA	FCA
μ C4 top	Signal	1	1	1	1	0	0	1	1	1	1	0	0
	Power	2	1	2	1	1	1	2	1	2	1	2	2
MC	Type	Mem	MPU	FCA	FCA	MPU	Mem	Mem	MPU	FCA	FCA	MPU	Mem
	Signal TSV	1	1	1	1	0	1	1	1	1	1	1	1
μ C4 bottom	Signal	1	1	1	1	1	1	1	1	1	1	1	1
	Power	2	2	2	0	2	0	2	2	0	2	0	2
BC	Type	FCA	FCA	Mem	MPU	Mem	MPU	FCA	FCA	Mem	MPU	Mem	MPU
	Signal TSV	1	1	1	1	1	1	1	1	1	1	1	1
C4	Power TSV	2	2	2	0	2	0	0	0	0	0	0	0
	Signal	1	1	1	1	1	1	1	1	1	1	1	1
C4	Power	2	2	2	2	2	2	0	0	1	0	1	0

2.1.1

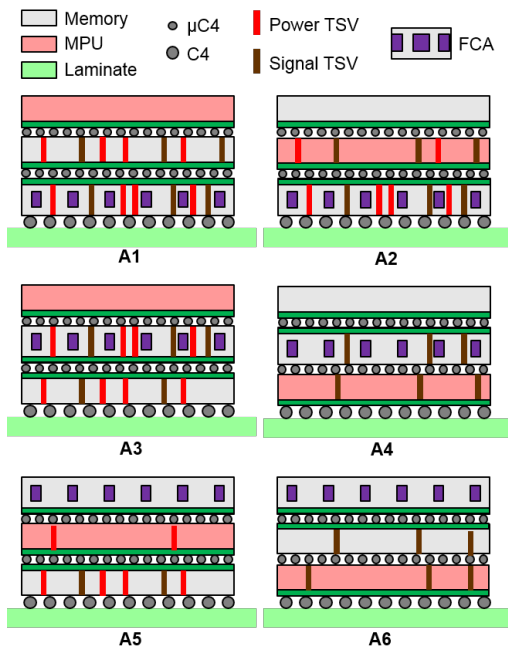


Fig. 4: Case A: FCA provides power for the memory layer only.

“Case A”. By rearranging the positions of the MPU, memory and the FCA as the top chip (TC), middle chip (MC) and bottom chip (BC), six different permutations can be obtained which are illustrated in Fig. 4 (A1-A6). With three dies, there are five levels of interconnects (from bottom): C4s, BC TSVs, bottom micro-C4s (connecting BC and MC), MC TSVs and top micro-C4s (connecting MC and TC). By changing the orientation of the three dies (*i.e.*, facing up or down) many more permutations can be obtained making the total number of possible designs too large to perform any meaningful analysis. Hence, for each permutation of die positions, only those die orientations that minimized the number of TSVs in the design were selected. In case of a conflict between two different die orientation-permutations with similar number of total interconnects, the one that minimizes power interconnects was chosen.

As for the Case O, the designs A1-A6 are evaluated for their heat-removal performance and amount of power interconnects in package, and the results are plotted in Fig. 3 in

blue. In general, Case A outperforms Case O. Within Case A, A4 and A6 have the least number of power interconnects: since the MPU is the BC, its power interconnects are fed from the laminate via C4s, while the power requirements of the memory die is satisfied by the FCA which is connected to F2F. Hence, all power TSVs are eliminated in the chip stack. However, A4 is superior to A6 in heat-removal since the FCA is the MC and is effectively closer to both MPU and the memory. The exact qualitative scores are also provided in Table 2 for interested readers.

Based on Fig. 3, design A4 was chosen as the most interesting implementation for further quantitative study since it requires the least number of interconnects for power delivery and has comparable heat dissipation to the benchmark 2D design O1.

2.2 3D MPSoC architecture and floorplan

In our system architecture simulations, we use a 3D MPSoC, which contains a 12-core 22 nm MPU floorplan with a total power dissipation of 190 W (335 W in “turbo mode”), chip size of 28.9 mm × 22.3 mm and thickness of 50 μm [16], combined with the second generation of High Bandwidth Memory (HBM) [17] chips. Note that the MPU characteristics are similar to the server-class IBM POWER8 microprocessor [18]. The HBM stacks simulated in the present work contain four memory chips of 7 μm height each.

The height of the FCA layer was varied between 50 and 400 μm while the width of the individual microchannels was varied between 50 and 200 μm [14, 15]. Channel dimensions below 50 μm result in a severe rise of hydraulic resistance, therefore limiting the flow rate below the level needed to achieve an effective cooling of the MPSoC. The wall thickness between channels was fixed at 50 μm to maximize the limited chip area usage but still have enough volume to place electrodes. The characteristics for all layers are summarized in Table 3.

	Width, mm	Length, mm	Height, μm	Nominal power, W
MPU	28.9	22.3	50	190, 335
Memory	11.9	7.8	4×7	15
FCA channels	0.05–0.2	22.3	50–400	Section 4

TABLE 3: Definition of layers in design A4 for simulations.

The edge lengths of the MPU and memory allow placing 4 HBM stacks on top of one MPU (Figure 5). Thus, each

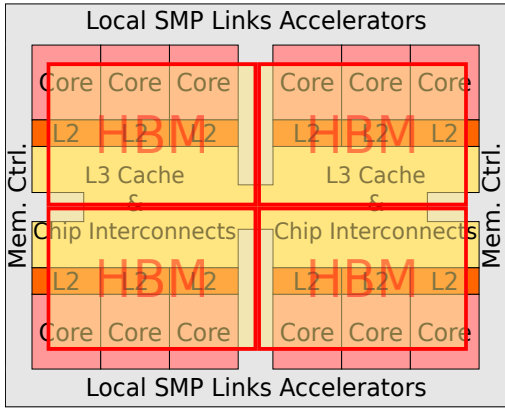


Fig. 5: Placement of HBM stacks on top of MPU

HBM stack is connected to 3 cores of the MPU, allowing 1.33 GB of memory per core.

3 POWERCOOL: COMPACT MODELING OF MICROFLUIDIC POWER DELIVERY

In this section we present our new PowerCool simulator, a comprehensive experimentally validated tool for the combined electro-thermal simulations of microfluidic power delivery in a 3D MPSoCs. Moreover, we describe the extensions with respect to our initial version [15], as follows:

- 1) We briefly describe the fundamental electrochemical concepts of the FCA and the basic equations governing its behavior, that are used in our FCA model.
- 2) We validate the accuracy of the FCA compact electrochemical model against measurements performed on an actual FCA test vehicle.
- 3) We extend the original PowerCool by coupling it with 3D-ICE [19] — a compact thermal model for liquid-cooled 3D MPSoCs. This demonstrates the cooling capabilities of the FCA in the context of a 3D MPSoC, and helps us analyze the impact of the fluid temperature on the electrochemical behavior of the FCA.
- 4) Finally, we further extend PowerCool to enable architectural exploration, by adding a leakage power estimation model. Leakage power significantly depends on temperature, and it is crucial to study its impact on the overall 3D MPSoC with FCA performance and efficiency.

3.1 Introduction to redox flow cells and PowerCool electrochemical model

The proposed FCA concept [14] comprises two electrode terminals, each in contact with a different electrolyte solution (Figure 6). The conversion of chemical energy stored in the liquid to electrical energy at the electrodes occurs via electrochemical redox reactions. Electrical power can be supplied to a load connected across the two electrode terminals, and the power delivered depends on the rate at which the redox reactions take place.

Flow cells with soluble reactants and products in both electrolyte streams are particularly attractive, as a continuous flow of electrolytes ensures a steady energy supply






Fig. 6: Structure of the PowerCool compact model for a discretized flow cell [15].

without the build-up of solid deposits or morphological changes of the electrodes, which leads to flow cell degradation due to increasing hydraulic and electrical resistances and electrodes aging. In conventional redox flow cells, a semi-permeable membrane is used to separate the two half-cell compartments [20] to avoid mixing of the two electrolytes and their participation in the opposite half-cell reaction, which would drastically decrease the cell's output voltage.

However, for microchannels with small hydraulic diameter, D_h , and low flow rates, the Reynolds number ($Re = \rho v D_h / \mu$ with density ρ , velocity v and viscosity μ) is sufficiently small to result in co-laminar flow of the two electrolyte streams, which prevents convective mixing [21]. Thus, no membrane is needed, which implies simpler fabrication of the targeted micro-scale flow cells.

On the other hand, diffusive mixing in membrane-less flow cells may reduce the round-trip energy efficiency which includes recharging the electrolyte solutions to their original state-of-charge. For simplicity, the present work only considers co-laminar flow based membrane-less FCAs, and the key findings of the present study are not affected by the hypothetical presence of a membrane.

As an illustration, consider the all-vanadium redox flow cell, which is currently the most mature and commercially relevant chemistry deployed for grid-scale installations [20]. All-vanadium flow cells utilize two different aqueous vanadium redox couples in sulphuric acid (V^{2+} as fuel and VO_2^+ as oxidant). The corresponding half-cell redox reactions taking place at the electrodes are:



During discharge, Reaction (1) is an oxidation of the fuel

V^{2+} and the standard electrochemical potential for this reaction vs. standard hydrogen electrode (SHE) is $E_0^- = 0.26 V$. Reaction (2) is a reduction of the oxidant VO_2^+ with the standard electrochemical potential $E_0^+ = +0.99 V$ vs SHE. The Nernst equation is then used to calculate an open circuit potential (OCP) of each half-cell as follows:

$$E_{OCP}^{+/-} = E_0^{+/-} + \frac{RT}{nF} \log \frac{a_{ox}^{+/-}}{a_{red}^{+/-}}, \quad (3)$$

where R is the universal gas constant, F is the Faraday constant, n is the number of electrons exchanged in redox reaction, T is the temperature of solution in K and a_{ox} and a_{red} are the activities of the oxidized and reduced species, respectively, at the surfaces of the electrodes, which are approximated as the surface concentrations for simplicity.

The electrochemical reaction at each electrode is characterized by a highly non-linear current-voltage relationship, which is known as the Butler-Volmer equation:

$$j^{+/-} = f_{\eta}^{+/-} = \frac{e^{-(1-\alpha)\eta\gamma} - e^{\alpha\eta\gamma}}{\frac{1}{j_0} + \frac{1}{j_{lim}} (e^{-(1-\alpha)\eta\gamma} + e^{\alpha\eta\gamma})}, \quad (4)$$

where η is the Butler-Volmer overpotential, α is the charge transfer coefficient, $\gamma = \frac{nF}{RT}$ and j_0 and j_{lim} are the exchange and limiting current densities respectively.

Knowing the concentrations of reactants and products at the surfaces of both electrodes, one can use Equations 3 and 4 to calculate the potential difference between the electrodes and the current flowing through the cell. Subtracting from this the voltage drop $\Delta U = jR_{int}$ due to the internal resistance R_{int} of the cell allows to determine the voltage and current supplied to an external circuit. The voltages, currents and the losses are all functions of the local concentrations of the reactants *vis-a-vis* the products. Hence, they change along the length of the flow cell. This situation affects the current flow distribution out of the cell, which in turn influences the concentration distribution, resulting in a highly non-linear system. The reactants are supplied to the surfaces of the electrodes via diffusion from the bulk of the fluids across their flow. The limited rate of this mass transfer of electrolyte species impose the upper limit for the current (j_{lim}), which the flow cell can possibly supply.

In our initial PowerCool's electrochemical model presented in [15], a flow cell is discretized into small sections, and for each section an equivalent electrical circuit consisting of the sources and losses is constructed as shown in Fig. 6 (using the Nernst equation (3) and the Butler-Volmer equation (4)). The temperature and concentration of the fluid in each section is assumed to be uniform in these calculations. Simultaneously, the concentrations are solved for each section along the channel. The resulting model is iteratively solved until convergence. Once all the circuit parameters are computed, modified nodal analysis is applied to each channel section to derive the appropriate circuit equations. Combining them together, it is possible to write the global circuit equations for the entire microchannel flow cell in the following form:

$$\mathbf{G}_{el}\mathbf{X}(t) + \mathbf{C}_{el}\dot{\mathbf{X}}(t) + \mathbf{F}(\mathbf{X}) = \mathbf{U}_{el}(t), \quad (5)$$

where $\mathbf{X}(t)$ is the vector of voltages and currents for all the sections, \mathbf{G}_{el} and \mathbf{C}_{el} are the linear conductance and

capacitance matrices, $\mathbf{F}(\mathbf{X})$ contain the non-linear functions of Butler-Volmer resistances and $\mathbf{U}_{el}(t)$ is a vector of Nernst voltage sources $E_{OCP}^{+/-}$ from all the sections. Solving this system, one finds voltages supplied by each small fuel cell and currents floating through them. From here, the total current and voltage supplied by the whole FCA are determined. Moreover, the mathematical model may easily account for the presence of a membrane, which adds an additional resistive term while reducing the amount of mixing.

3.2 Validation of PowerCool electrochemical model

For experimental validation of the electrochemical model, a single-channel microfluidic membrane-less flow cell was employed [22]. A 400 μm wide and 100 μm deep channel was dry-etched in silicon. Platinum thin-film electrodes were deposited via sputtering, and the channel/electrode assembly was sealed via anodic bonding of the silicon chip with a glass cover. Two electrolyte streams were fed into the channel at defined flow rates using syringe pumps. The redox electrolytes used were 10^{-4} mol/m^3 $\text{FeSO}_4 \cdot 7\text{H}_2\text{O}$ as oxidant and $5 \cdot 10^{-5} \text{ mol/m}^3$ anthraquinone disulphonic acid as fuel, both dissolved in 1 M (*i.e.* 1000 mol/m^3) H_2SO_4 . These electrolytes were chosen because their redox potentials were found to lie within the stability window of the 1 M H_2SO_4 supporting electrolyte on platinum. The dilute concentration of the active species was chosen in order to experimentally resolve the mass-transport limited regime which should be accurately replicated by the electrochemical model.

During our experiments, the flow velocity of the solutions was varied between 0.01 and 1 m/s, and the solution temperature was varied from 298 K to 336 K. Voltage-current profiles were recorded using a potentiostat by linearly sweeping the cell voltage from the open-circuit voltage (OCV) to 0.05 V.

Experimental polarization curves of voltage vs. current density at electrodes surfaces are compared to PowerCool simulations in Figures 7a and 7b. All polarization curves exhibit a non-linear voltage drop close to the OCV (Region A) due to the kinetic overpotential defined by Equation 4. At higher currents, the voltage drop is linear due to the internal resistance of the cell, which exhibits ohmic behavior (Region B). At even higher currents, the rate of reaction is limited by diffusion of the reactants to the electrode surface, resulting in a steep drop in voltage when the current approaches a certain value (Region C). This diffusion-limited current depends on multiple parameters, such as temperature, concentrations of electrolyte species, electrode geometry, flow velocity profile, *etc.* At higher flow velocities (Fig. 7a), the limiting current increases until the ohmic region (Region B) extends down to 0 V, and Region C is no longer discernible. Similarly, higher temperatures (Fig. 7b) lead to an increase in the limiting current due to higher diffusion rates. Exchange current density j_0 also grows with the temperature, which reduces the activation losses (Region A).

The experimental data was fitted with the model described in Section 3.1 using as fitting parameters j_0 and j_{lim} . Fitting j_0 allows to accurately reproduce experimental data in Region A, and fitting j_{lim} allows to achieve better correspondence in Region C. Together, they allow us to achieve an accuracy of 4%.

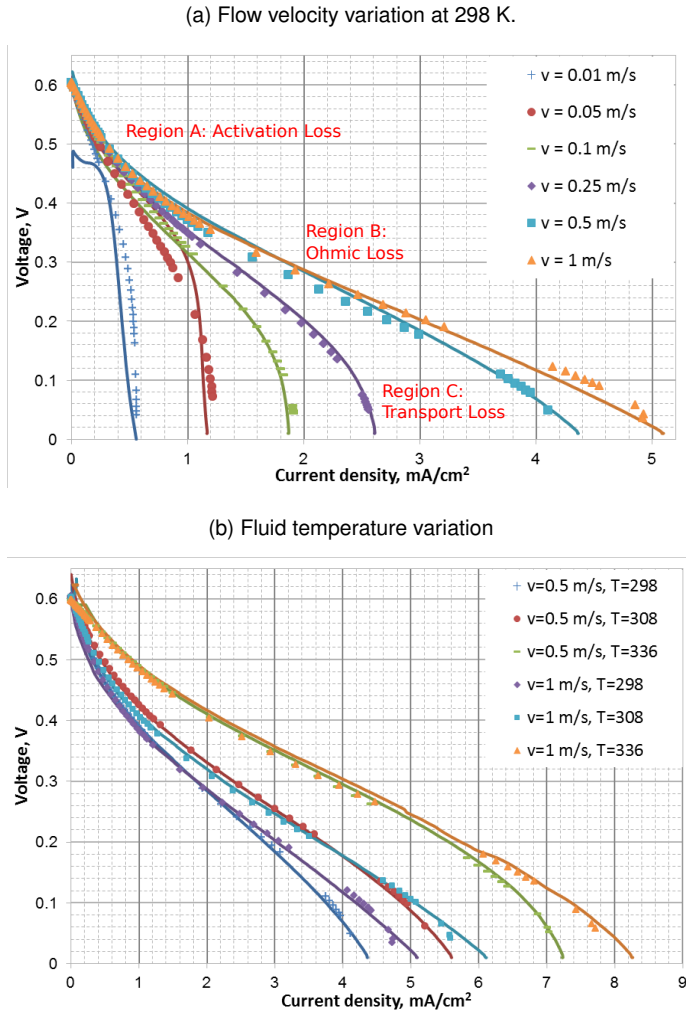


Fig. 7: Experiments (lines) vs. simulation (dots)

In summary, the PowerCool electrochemical model employed in this work is capable of accurate resolving all three regions of the voltage-current characteristics of dilute electrochemical redox flow cells with planar electrodes. This is more challenging than for concentrated solutions in which the kinetic and concentration overpotentials are typically much less pronounced. In the following sections, we describe how the electrochemical model is integrated into the new PowerCool simulation tool.

3.3 Thermal model inclusion

Thermal modeling of an FCA is needed not only for evaluating the cooling capabilities of the FCA, but also to study the impact of fluid temperature on the electrochemical behavior of the FCA. For this purpose we incorporate the 3D-ICE model [19] in PowerCool. 3D-ICE is a compact thermal modeling tool, which is used to calculate volumetric temperature distribution in a given 3D MPSoC with integrated microchannel cooling. The fundamental heat transfer equations of the system are solved using finite-difference discretization and the equivalent electrical circuit representation, as follows:

$$\mathbf{G}_{th}\mathbf{T}(t) + \mathbf{C}_{th}\dot{\mathbf{T}}(t) = \mathbf{U}_{th}(t), \quad (6)$$

where $\mathbf{T}(t)$ is the vector of all node temperatures (as a function of time), \mathbf{G}_{th} is a symmetric seven-diagonal conductance matrix, \mathbf{C}_{th} is a diagonal cell heat capacitance matrix and $\mathbf{U}_{th}(t)$ is a vector of heat sources, wherever they exist, as a function of time.

The input parameters required to build these matrices and the matrices for PowerCool's electrochemical model in Eq. (5) are listed in Table 4. 3D-ICE was included into PowerCool to work in one flow, sharing common variables (such as geometrical and flow parameters of FCA) and simulation results. That is, the FCA generated power is fed into 3D-ICE model to account for it in thermal simulation, and the computed temperatures from 3D-ICE are used to update the electrochemical parameters.

Matrix*	Thermal model	Electrochemical model**
$\mathbf{G}_{th,el}$	Thermal conductivities of each used material, fluid velocity	Electrical conductivities of fluid and electrodes
$\mathbf{C}_{th,el}$	Specific heat capacities of all the materials	Electrode double-layer capacity
\mathbf{F}	-	Diffusion coefficient, a_{ox} , a_{red} , α , T , j_0
$\mathbf{U}_{th,el}$	Power map	a_{ox} , a_{red} , E_0 , T

*To calculate the elements of these matrices it is necessary to know geometrical parameters of every elementary cell
 **All the chemical parameters are required for both anodic and cathodic half-cell reactions

TABLE 4: Main input parameters

3.4 MPSoC leakage model inclusion

Leakage represents a significant fraction of total power consumption in high-performance MPSoC architectures and is affected greatly by the overall system temperature distribution [23]. To improve the accuracy of the simulations, it is therefore important to take leakage power as well as its dependence on temperature into consideration. Leakage power can be represented as a sum of subthreshold leakage, which has an exponential dependence on the temperature, and gate leakage, which is temperature-independent [24]:

$$P_l(T) = P_{gate} + P_{sub}(T) = P_{gate} + P_{ref} \cdot e^{\kappa(T-T_{ref})} \quad (7)$$

where κ is an exponential factor related to the technology of the processor (estimated to be 0.013 for a 22 nm MPU) and P_{ref} is the value of subthreshold leakage, measured at the chip temperature T_{ref} .

We assumed values for P_{ref} and T_{ref} that result in subthreshold leakage being 20% of total MPU power at the reference temperature of 80 °C. The remaining 80% account for the active power and the temperature-independent gate leakage. Equation 7 is used to calculate the subthreshold leakage power for a given temperature distribution on the MPU, thereby enabling an estimation of the actual 3D stack power consumption as a function of temperature. We extended PowerCool to include the above leakage model to the power consumption of the MPSoC.

3.5 Architecture of the new PowerCool simulator

We extended the original PowerCool with 3D-ICE thermal model and the MPSoC leakage power model described

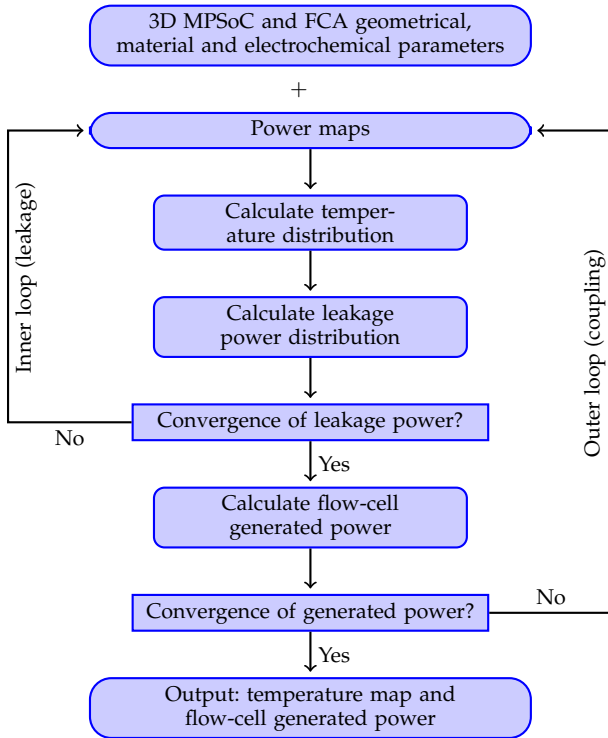


Fig. 8: PowerCool’s simulation algorithm

above to construct a comprehensive modeling and design exploration tool that analyses both cooling effects of single-phase cooling technology and the power generation capabilities of the micro-scale FCA, while estimating the leakage power in the chip. The final extended PowerCool simulation flow is shown in Fig. 8.

Our simulator takes as input the 3D MPSoC and FCA geometrical, material and electrochemical parameters, together with the power consumption (power maps) for each layer of the MPSoC. We assume a scenario in which the power delivery through TSVs without an FCA is not enough to power-up the whole chip, and some parts of the chip need to stay dark (*i.e.*, the chip is either underutilized, or some components operate at a reduced frequency and draw less power). Within this scenario, the FCA generated power is added to the power supplied through the PCB or power TSVs to “brighten” the dark silicon, thus altering the power maps — essentially, the input of the thermal simulation. Therefore, it is required to couple the thermal and electrochemical models. To avoid possible high computational overhead solving the fully coupled on the model-level system, we choose the iterative approach, applying separate thermal and electrochemical models one after another in the loop.

The simulation starts by calculating the temperature distribution from the initial assumptions, then incorporating temperature-dependent leakage model in inner loop shown at Fig. 8. Temperature-dependent leakage power correction changes the total power, in the first iteration, by less than 10%. Together with a low condition number of the thermal model matrix \mathbf{G}_{th} (which is approximately equal to $0.05\|\mathbf{T}\|$) and monotonical dependence of temperature distribution on dissipated power, it ensures fast convergence of

the inner loop.

Electrochemical model is applied next in the outer loop to calculate FCA generated power, which in turn updates the total power map, and the simulation reiterates the inner loop execution. At every iteration of the outer loop, thermal distribution and FCA power are compared with the results of the previous iteration to quantify the converge of the full simulation. FCA increases power maps by less than 35% after the first iteration, which yields fast thermal solution convergence. This results in a change in the thermal distribution, which is then used to calculate the updated FCA-generated power. We observe that the change in the generated power is less than the change in updated thermal distribution with a sensitivity less than 0.7 (for each 10% change in $\|\mathbf{T}\|$, FCA generated power is changed by less than 7%), which ensures convergence of the outer loop. We have observed that, in practice, convergence of the inner loop requires 4 to 5 iterations, and the outer loop converges after 2 to 3 iterations.

The computational complexity of the thermal and electrochemical simulations are $O(n^{1.7})$ [19] and $O(m^{1.2})$ [15], respectively, where n and m are the dimensions of the thermal and electrochemical problems. Since the number of iterations of the outer loop is very low compared to n and m , the total complexity of our iterative approach is $O(n^{1.7})$, since n and m are of the same order of magnitude. On the other hand, solving the fully coupled system has a computational complexity not lower than $O((n+m)^{1.7})$, and it can be as high as $O((n+m)^3)$, in the general case, due to additional couplings.

4 SIMULATIONS OF 3D-MPSOCs WITH INTER-TIER FCA

4.1 System configuration

All numerical experiments in this paper are run for the case of static workloads. Nominal power supplied to the processor (the power without taking into account the FCA generated power and the temperature dependence of leakage) was considered to be 190 W in normal mode and 335 W for a turbo mode:

$$P_{nom} = P_{active} + P_l(T_{ref}) = P_{active} + P_{gate} + P_{ref} \quad (8)$$

In this work, power consumption is evenly distributed among the 12 cores of the MPU. For the memory, a uniform power distribution is assumed for an HBM stack with total power being equal to 15 W (as reported for the first generation of HBM [25]). Leakage power was accounted for as described in section 3.4.

The 3D stack A4 in Figure 4 was modelled without an air-cooled heat sink, therefore, an adiabatic boundary condition was applied to the top of the stack. This is justified because the heat exchange between a 3D stack and surrounding air is estimated to be 2–3 orders of magnitude lower than the heat removal by FCA. We found that this arrangement gave equivalent results to the simulation of complex symmetrical stacks and stacks consisting of repeating identical blocks.

The electrolyte solutions assumed to correspond to an all-vanadium redox flow system, comprising 2.5 M VO_2^+

Study	Height, μm	Width, μm	Flow vel., m/s	Inlet temp., $^{\circ}\text{C}$
Height	50 \rightarrow 400	50	2.5	50
Width	100	50 \rightarrow 200	2.5	50
Flow velocity	100	50	0.8 \rightarrow 2.5	50
Inlet temperature	100	200	2.5	25 \rightarrow 60

TABLE 5: Values of main FCA parameters in power generation studies

as oxidant and 2.5 M V^{2+} as fuel in separate streams. The supporting electrolyte providing ionic conductivity was the mix of 2.5 M H_2SO_4 and 6 M HCl [26]. The electrochemical parameters of all-vanadium flow cell were taken from previous work [15]. The microchannels were assumed to be straight with rectangular cross-section and aligned parallel to the shortest side of the die (*i.e.*, from top to bottom in Figure 5). Flow rates and geometrical parameters of channels were varied in a wide range, as shown in Table 5. The lower limits of the geometrical parameters were set to not violate the thermal constraints and ensure effective cooling. The upper limits were chosen heuristically based on the achieved results. The upper limit for the inlet temperature was chosen according to the temperature range of electrolyte stability [26]. Since fluid temperatures lower than the ambient temperature need an active chiller-based cooling mechanisms to bring the refrigerant back to its initial temperature, we set the lower limit for the inlet temperature to 25 $^{\circ}\text{C}$, which ensures that the solution can be cooled using a chiller-less solution (*i.e.*, using only heat exchangers).

We studied five different workload allocations with respect to their influence on FCA operation by varying the number and position of active cores as follows: (1) all 12 cores are active, 6 active cores are placed in (2) north side of the MPU, (3) south side, (4) west side and (5) checkerboard pattern across the chip. For a fixed FCA (289 channels of 50 μm width, flow velocity 2.5 m/s, inlet temperature 50 $^{\circ}\text{C}$), the in-stack temperature distribution is affected greatly by the workload allocation. In fact, the maximum temperature was found to vary from 75 $^{\circ}\text{C}$ to 101 $^{\circ}\text{C}$ depending on the workload. However, FCA power generation was similar in all cases, ranging from 20.2 W to 20.8 W. Thus, the workload allocation has only a minor effect on the FCA power generation, but a high impact on temperature distribution. Therefore, to understand the characteristics of FCA power generation, only one case of the workload allocation can be considered. The case of uniform workload distribution between all 12 cores of the processor was chosen for further detailed study due to its symmetry and relatively low temperatures.

4.2 FCA power generation exploration

4.2.1 Static parameters

Four parametric studies were conducted to characterize the system in terms of FCA power generation and maximum chip temperature (Figure 9). In these studies we vary one parameter, keeping the other constant (see Table 5). Assuming that the electrode terminals fully cover the vertical walls of the microchannels, an increase in channel height leads to

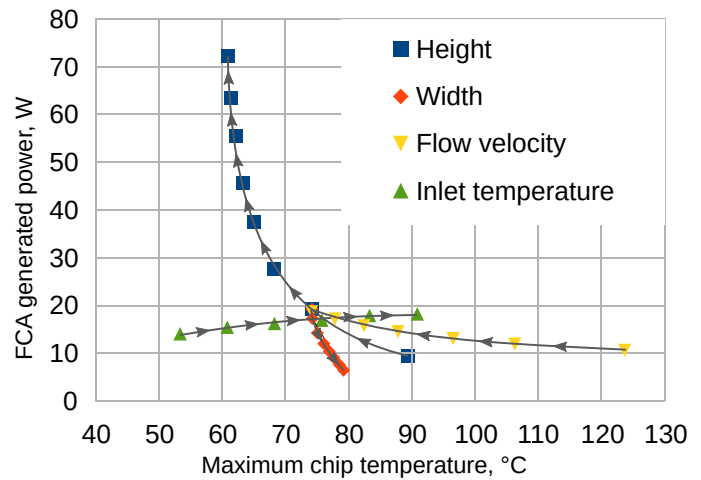


Fig. 9: Generated power and maximum chip temperature as a function of FCA parameter variations in Table 5.

a proportional increase in electrode area and therefore also in the delivered power. Moreover, an increase in channel height requires a larger volumetric flow rate to maintain the same linear flow velocity. Therefore, an increase in channel height also results in improved heat dissipation but increases the pumping power (it was not bigger than 3.56 W in all our simulations, which is significantly lower than FCA generated power or leakage power, and therefore was not considered as a constraint).

The second primary parameter is channel width. If it is decreased, a larger amount of channels can be placed in a single FCA layer, which increases the overall electrode area and also provides better heat transfer from the silicon to the fluids. This results in higher power generation and better cooling capacity, however the hydraulic resistance of the channels rises. Restricting the permissible fluid pressure drop to 1.5 bar/cm along the channels while having high enough flow rate to cool the stack, we set the minimum channel width to 50 μm .

4.2.2 Dynamic parameters

While channel height and width are geometric parameters which are fixed during fabrication of the FCA, the flow velocity and fluid inlet temperatures are two parameters that may be adjusted during operation. Increasing flow velocity was found to be an effective way to both lower chip temperature and increase FCA power generation. However due to a stronger dependence on the flow velocity, the related increase in pumping power outweighs the power generation gain when flow velocity reaches a certain critical value. The maximum flow velocity in our simulations was 2.5 m/s for channels 100 μm high and 50 μm wide, restricted by the aforementioned pressure drop constraint, which is still lower than the critical velocity. For channels with larger cross-sections, hydraulic resistance is lower, and flow velocity can exceed either the critical value or the laminar flow constraints before the corresponded pressure drop reaches its limit.

Finally, inlet temperature can also be increased to maximize FCA power generation, accepting a trade-off with the maximum chip temperature. However, the leakage power

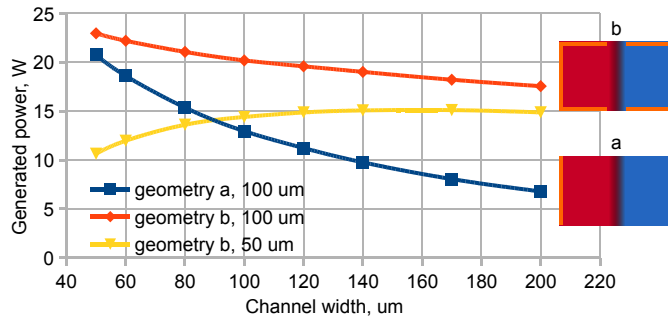


Fig. 10: Comparison of FCA power generation dependence on the channel width for the cases of *a* and *b* electrode placements

rise due to the higher MPU temperature outweighs the FCA power generation increase. Therefore, given the current 22 nm MPU architecture under analysis, it seems reasonable to keep the inlet temperature low for this specific simulation scenario. On the other hand, note that elevated fluid temperatures are beneficial in terms of heat rejection to ambient, and by allowing free cooling at the facility level, an overall energetic benefit in the case of higher inlet temperatures is still likely to be incurred.

4.2.3 Power generation maximizing

The parameter with the most important impact on power generation while also having a beneficial effect on maximum chip temperature was found to be the electrode area (Figure 9). Thus, 100 μm thick FCA layer compatible with inter-tier implementations is able to generate up to 19 W of power (which correspond to 3.0 W/cm^2 of chip area), while a 400 μm thick FCA with the area of electrodes being greater by 4 times can produce 72 W (11 W/cm^2). An increase in electrode area may be achieved by introducing porous electrodes, which is state-of-the-art in large-scale redox flow batteries [20] but not compatible with CMOS process integration. A simple way to enlarge electrode area is to use the top and the bottom sides of the channels in addition to the channel sidewalls. For this configuration, a comparison between the generated power dependencies on the channel width is shown on Figure 10. The case of 100 μm high FCA with planar parallel electrodes is chosen as a baseline (geometry *a*). The power generation in the case of 50 μm high FCA with planar electrodes is roughly a half of the baseline case, and it is not shown on the figure for clarity. Extended electrodes (geometry *b*) allow to increase the best achievable power generation in case of 100 μm high FCA by 20% (up to 3.6 W/cm^2) and to decrease the power generation drop due to increase of the channel width. However, for the case of a 50 μm high FCA, the extended electrodes change the width dependence of power generation from decreasing to increasing with channel width, which reaches its peak at a channel width of *ca.* 170 μm . This happens because of a more complex geometrical dependence of the overall electrode area from the height and the width of microchannels, which also involves the distance between the electrodes and channel walls thickness.

The achieved FCA power generation density allows to generate up to 23 W per FCA layer (12% of the chip

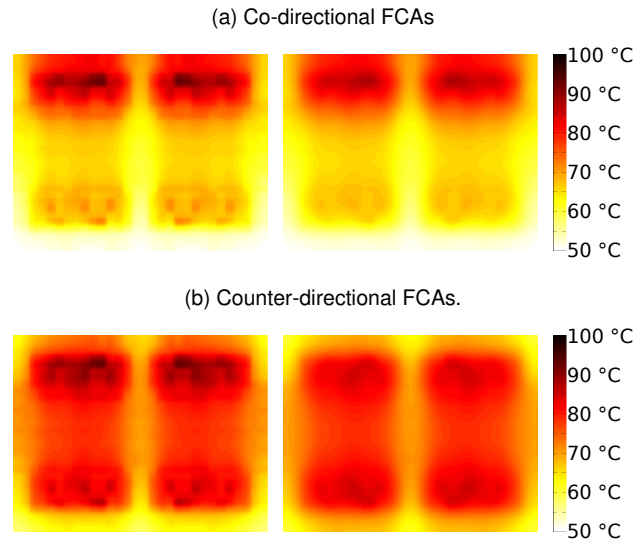


Fig. 11: Simulated heat map of the bottom MPU (on the left) and memory (right) layers with different orientations of baseline FCAs.

power consumption). This value is comparable to the power consumption of the simulated chip elements, i.e., 10-19 W per core (depending on the workload), 17 W for L2 cache, 14 W for L3 cache and 15 W for an HBM stack. Therefore, it is possible to power up one or two components of the 3D MPSoC entirely by the FCA, alleviating the power wall and increasing the fraction of TSVs that are dedicated to I/O communication.

4.3 On-chip thermal gradients

In the configurations studied above, the high MPU heat dissipation results in large temperature differences from inlet to outlet of the order of 20–40 $^{\circ}\text{C}$ depending on the workload and the FCA characteristics. This raises concerns about reliability, as many failure mechanisms have been reported to be accelerated by spatial thermal variations [27, 28].

In the case of more complex stacks with several FCA layers, the temperature rise may be reduced by applying opposite directions of the flow in alternating layers. We considered a stack, which contains one extra FCA and one extra MPU layers on top of A4, resulting in a symmetrical 3D stack with respect to the memory layer. Each MPU is supplied with the nominal power of 335 W with the workload evenly distributed among 12 cores. This resembles a “turbo mode” of the modern processors, when each cores operates at its limit. Heat maps of the bottom MPU and the memory layer in case of co-directional flows in both FCAs are shown on Figure 11a. Temperature of the north cores (the outlet side of the stack) is about 20 $^{\circ}\text{C}$ higher than the south cores (the inlet side). Heat maps for the same stack with the counter-directional flows in different FCAs (Figure 11b) show the temperature difference between south and north cores of only 5 $^{\circ}\text{C}$. These simulations assumed an extreme case of MPU utilization with a nominal power consumption of 335 W.

However, in the second case the average temperature of the stack is higher, which results in higher leakage: 67.4 W

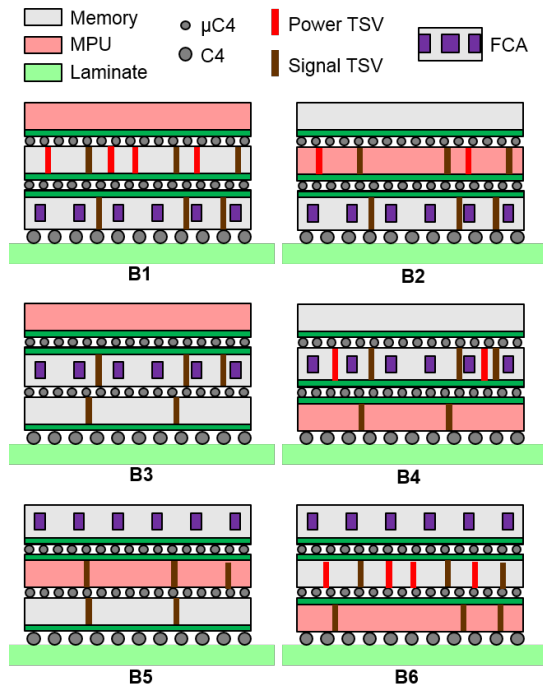


Fig. 12: Case B: FCA provides power for the entire chip stack.

of subthreshold leakage out of 735 W of supplied power (besides 38.3 W of FCA generated power) in the case of the opposing flows against 60.9 W out of 722 W (besides 37.6 W of FCA generated power) in the case of co-directional flows. In conclusion, counter-directional flow reduces thermal gradients, whereas co-directional flow reduces leakage.

5 IMPACT OF BEYOND STATE-OF-THE-ART FCAS ON 3D MPSoCs

So far in this paper, we have performed the qualitative and quantitative analysis, modeling and simulation of state-of-the-art (SoA) FCAs. This technology provides power densities that are at least an order of magnitude less than what is typically needed to power an MPSoC, i.e. Case A in Section 2.1. However, it is possible to theorize the impact of beyond SoA FCAs (Case B) that can provide much higher power densities, comparable to the power consumption of MPSoCs, using a similar qualitative analysis. Such an analysis motivates the development of these technologies and provides some perspectives on the 3D integration of the MPSoCs of the future.

As for Case A in Section 2.1, six different permutations of die positions can be obtained for Case B as illustrated in Fig 12 (B1-B6). Note that the die orientations here (selected using the same criteria as that used for Case A) are slightly different from Case A. The FCA here also delivers power to the whole 3D stack and, given a choice between eliminating power TSVs to the MPU (high) or to the memory die (low), the former was preferred. Also note that in designs B3 and B5, there are no power TSVs drawn from the FCA to the memory die. This is because we assumed that for these cases, as the memory dies are the BCs facing down, a (relatively) low number of power C4s can be easily afforded to

power the memory die from external (conventional) power sources via the laminate. By doing this, it can be seen that both B3 as well as B5 eliminate all power TSVs in the chip stack.

The qualitative scores for Case B are plotted in Fig. 3 in green (the exact scores are also provided in Table 2). As expected, Case B designs outperform Case A overall, as they are superior in terms of power delivery. Within Case B, the designs B2-B5 are equivalent in terms of total reduction of power interconnects. If eliminating power TSVs is a priority, then the designs B3 and B5 will be favored. Between these two designs, B3 would be chosen for its superior heat-removal capability. It is interesting to note from this graph that designs A4 and B5 have very similar quality. This indicates that using judicious design, even with SoA FCA technology it is possible to match the power-thermal performance of beyond SoA technology.

Thus, major advances in FCA power generation would potentially result into a total elimination of power TSVs in 3D MPSoCs, allowing unprecedented vertical chip scalability. Our results show two major directions to increase FCA power generation: i) developing new electrolyte compounds with higher energy densities, and ii) improving FCA design to increase electrodes area. More elaborate electrode configurations than the ones illustrated on Figure 10 with larger surface area are not yet supported by PowerCool’s geometry model, however our preliminary simulations show that it is likely to achieve at least 60 % boost in power generation for interlayer FCAs (5.7 W/cm²) by implementing more complex electrodes.

It has previously been shown that direct chip-level cooling is the only solution to reliably operate high power density single-chip microprocessors [29] and 3D multi-chip MPSoCs [9]. At the same time, trends to incorporate electrochemical storage and conversion devices in greater proximity to servers have been reported in literature [30, 31]. Apart from improved energy efficiency and operating cost savings, the above approaches are in many respects also considered favorable in terms of reliability. The proposed FCA technology reinforces the above trends, and, similarly to conventional battery and fuel cell technologies, may benefit from the experience from large-scale field installations in terms of scalability [20].

6 CONCLUSION

In this paper we have shown a qualitative analysis that describes the benefits of integrated FCA in 3D MPSoCs. We presented our new PowerCool simulator, a compact modeling tool to simulate cooling and power generation performance of FCAs in 3D stacks. We validated the FCA electrochemical model against experimental data and presented a parametric analysis of FCA performance, defining the most important static and dynamic FCA parameters.

The proposed FCA-based interlayer microchannel cooling and power generation technology has proven to be a promising solution to tackle the thermal and power challenges of future 3D MPSoCs. Our research demonstrates that, with the current technology, it is possible to generate up to 3.6 W of power per cm² of the chip area for the interlayer FCAs and up to 11 W/cm² for the topmost one,

while meeting thermal constraints under the worst case scenario of sustained full load. The generation capabilities achieved would be enough to power the on-chip higher-level caches, potentially reducing the amount of power TSVs needed, and alleviating the limitations imposed by the *power wall*.

The technological limits of integrated FCAs are yet to be discovered. However, given the trade-offs encountered in this work, in the future we envision to be able to drastically increase generated power by optimizing FCA geometry to increase overall electrodes area and fuel utilization, which will result in higher FCA power generation, reducing the number of power TSVs in 3D stacks and increasing vertical scalability. Moreover, further advances in investigation of new, more effective electrolytes will significantly help to achieve FCAs, capable of powering up the whole MPSoCs.

ACKNOWLEDGMENTS

The authors would like to thank Quentin Cabrol for carrying out experiments, used for PowerCool validation, as a part of his Master project at ESL-EPFL.

This work has been partially supported by the YINS RTD project (No. 20NA21 150939), funded by Nano-Tera.ch with Swiss Confederation Financing and scientifically evaluated by SNSF, the REPCOOL project (Grant No. 147661), funded by the Swiss National Science, the EC H2020 MANGO project (Agreement No. 671668), and the ERC Consolidator Grant COMPUSAPIEN (Agreement No. 725657).

REFERENCES

- [1] G. E. Moore. Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86(1):82–85, Jan 1998.
- [2] R.H Dennard et al. Design of ion-implanted mosfet's with very small physical dimensions. *Solid-State Circuits, IEEE Journal of*, 9(5):256–268, October 1974.
- [3] John H. Lau. Evolution and outlook of TSV and 3D IC/Si integration. *2010 12th Electronics Packaging Technology Conference, EPTC 2010*, pages 560–570, 2010.
- [4] Wm. A. Wulf and Sally A. McKee. Hitting the memory wall: Implications of the Obvious. *ACM SIGARCH Computer Architecture News*, 23:20–24, 1995.
- [5] C. C. Liu, I. Ganusov, M. Burtscher, and S. Tiwari. Bridging the processor-memory performance gap with 3D IC technology. *IEEE Design and Test of Computers*, 22(6):556–564, 2005.
- [6] G.L. Loi et al. A thermally-aware performance analysis of vertically integrated (3-D) processor-memory hierarchy. *Design Automation Conference (DAC)*, pages 991–996, 2006.
- [7] Gabriel H. Loh. 3D-stacked memory architectures for multi-core processors. *Int. Symp. on Computer Architecture*, pages 453–464, 2008.
- [8] P. Ruch, T. Brunschwiler, W. Escher, S. Paredes, and B. Michel. Toward five-dimensional scaling: How density improves efficiency in future computers. *IBM Journal of Research and Development*, 55(5):15:1–15:13, Sept 2011.
- [9] T. Brunschwiler, B. Michel, et al. Interlayer cooling potential in vertically integrated packages. *Microsystem Technologies*, 15(1):57–74, 2008.
- [10] P. Stanley-Marbell, V.C. Cabezas, and R.P. Luijten. Pinned to the walls — impact of packaging and application properties on the memory and power walls. In *Int. Symp. on Low Power Electronics and Design (ISLPED)*, pages 51–56, Aug 2011.
- [11] T. Brunschwiler et al. Thermal power plane enabling dual-side electrical interconnects for high-performance chip stacks: Concept. In *Electronics System-Integration Technology Conference (ESTC), 2014*, pages 1–6, Sept 2014.
- [12] E. A. Burton and al. FIVR : Fully integrated voltage regulators on 4th generation Intel Core SoCs. In *IEEE Applied Power Electronics Conference (APEC)*, pages 432–439, March 2014.
- [13] Patrick Ruch, Thomas Brunschwiler, Stephan Paredes, Ingmar Meijer, and Bruno Michel. Roadmap towards ultimately-efficient zeta-scale datacenters. *Int. Conf. on High Performance Computing and Simulation, HPCS 2013*, pages 161–163.
- [14] M. M. Sabry, A. Sridhar, D. Atienza, P. Ruch, and B. Michel. Integrated microfluidic power generation and cooling for bright silicon mpsoCs. In *Design, Automation Test in Europe (DATE)*, pages 1–6, March 2014.
- [15] A. Sridhar, M.M. Sabry, P. Ruch, D. Atienza, and B. Michel. Powercool: Simulation of integrated microfluidic power generation in bright silicon mpsoCs. In *Computer-Aided Design (ICCAD)*, pages 527–534, Nov 2014.
- [16] Ibm openpower connect homepage link. http://openpowerfoundation.org/?resource_lib=ibm-openpower-connect-homepage-link, 2016.
- [17] K. Tran and J. Ahn. Hbm: Memory solution for high performance processors. In *MemCon*, Oct 2014.
- [18] E. J. Fluhr et al. 5.1 power8tm: A 12-core server-class processor in 22nm soi with 7.6tb/s off-chip bandwidth. In *International Solid-State Circuits Conference (ISSCC)*, pages 96–97, Feb 2014.
- [19] Arvind Sridhar, Alessandro Vincenzi, David Atienza, and Thomas Brunschwiler. 3D-ICE: A compact thermal model for early-stage design of liquid-cooled ICs. *IEEE Transactions on Computers*, 63(10):2576–2589, 2014.
- [20] M. Skyllas-Kazacos et al. Progress in flow battery research and development. *Journal of The Electrochemical Society*, 158(8):R55–R79, 2011.
- [21] E. Kjeang, N. Djilali, and D. Sinton. Microfluidic fuel cells: A review. *Journal of Power Sources*, 186(2):353 – 369, 2009.
- [22] J. Marschewski et al. Mixing with herringbone-inspired microstructures: overcoming the diffusion limit in co-laminar microfluidic devices. *Lab Chip*, 15:1923–1933, 2015.
- [23] S. G. Narendra. Leakage in nanometer cmos technologies. New York, NY, 2005. Springer.
- [24] S. A. Hall and G. V. Kopcsay. Energy-Efficient Cooling of Liquid-Cooled Electronics Having Temperature-Dependent Leakage. *Journal of Thermal Science and Engineering Applications*, 6(1):011008/1–011008/12, 2014.
- [25] Samsung begins mass producing worlds fastest dram

based on newest high bandwidth memory (hbm) interface. <https://news.samsung.com/global/samsung-begins-mass-producing-worlds-fastest-dram-based-on-newest-high-bandwidth-memory-hbm-interface>, Jan 2016.

- [26] L. Li et al. A stable vanadium redox-flow battery with high energy density for large-scale energy storage. *Advanced Energy Materials*, 1(3):394–400, 2011.
- [27] JEDEC: jep122E. Failure mechanisms and models for semiconductor devices. <http://www.jedec.org/standards-documents/docs/jep-122e>, Oct 2011.
- [28] A. K. Coskun, T. S. Rosing, and K. C. Gross. Utilizing predictors for efficient thermal management in multi-processor socs. *IEEE Trans. on Computer-Aided Design (TCAD)*, 28(10):1503–1516, Oct 2009.
- [29] E. G. Colgan et al. A Practical Implementation of Silicon Microchannel Coolers for High Power Chips. *IEEE Trans. Components Packag. Technol.*, 30:218–225, 2007.
- [30] V. Kontorinis, L. E. Zhang, B. Aksanli, J. Sampson, H. Homayoun, E. Pettis, D. M. Tullsen, and T. Simunic Rosing. Managing distributed ups energy for effective power capping in data centers. In *2012 39th Annual International Symposium on Computer Architecture (ISCA)*, pages 488–499, June 2012.
- [31] Ana Carolina Riekstin, Sean James, Aman Kansal, Jie Liu, and Eric Peterson. No more electrical infrastructure: Towards fuel cell powered data centers. *SIGOPS Oper. Syst. Rev.*, 48(1):39–43, May 2014.



Artem Andreev received the BS (in 2011) and MS (in 2013) degrees in applied physics and mathematics from the Moscow Institute of Physics and Technology, Dolgoprudny, Russia. In 2015 he joined the Embedded Systems Laboratory (ESL) at the École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, to pursue the PhD degree in electrical engineering. His current research interests include steady and transient system modeling and optimization of flow cell implementation in

3D integrated circuits.



Arvind Sridhar is a Post-doctoral Researcher in the Science & Technology department at IBM Research - Zürich. He received his PhD in Electrical Engineering from the Swiss Federal Institute of Technology (EPFL), Lausanne, in 2013. Dr. Sridhar joined IBM Research in 2014, and his current research focus is in electronic packaging, thermal modeling, on-chip power supplies and 3D integration.



Mohamed M. Sabry received the M.Sc. and Ph.D. degrees in electrical and computer engineering from Ain Shams University, Cairo Governorate, Egypt, and from cole Polytechnique Fdrale de Lausanne (EPFL), Lausanne, Switzerland, in 2008 and 2013, respectively. He is currently a postdoctoral research fellow at Stanford University and a visiting scholar in the Embedded Systems Lab, EPFL. His current research interests include system design and resource management methodologies in embedded systems,

and multiprocessor system-on-chips (MPSoCs), especially temperature and reliability management of 2-D and 3-D MPSoCs, with particular emphasis on emerging computing, memory, and cooling technologies. Dr. Sabry was the recipient of the Swiss National Science Foundation Early Post-Doctoral Mobility Fellowship in 2013.



Marina Zapater is currently a Post-Doctoral researcher in the ESL at EPFL, and a Visiting Professor at Universidad Complutense de Madrid, Spain. She received her Ph.D. degree in Electronic Engineering from Universidad Politcnica de Madrid in 2015, an MSc in Telecommunication Engineering and a MSc in Electronic Engineering from the Universitat Politcnica de Catalunya, in 2010. Her research interests include proactive thermal and power optimization of complex heterogeneous systems, energy efficiency in data centers, ultra-low power architectures and embedded systems.

efficiency in data centers, ultra-low power architectures and embedded systems.



Patrick Ruch is Research Staff Member in the Science & Technology department at IBM Research – Zurich. He studied materials science at the Swiss Federal Institute of Technology (ETH) Zurich and received his PhD degree from the same institution for his work performed at the Paul Scherrer Institut (PSI) on electrochemical capacitors. Dr. Ruch joined IBM Research in 2009 and is leading activities in the areas of microfluidic electrochemical energy conversion and thermally driven sorption heat pumps. He

has also contributed to projects on direct liquid-cooled servers and high-concentration photovoltaics. His main research interests are in energy conversion and storage with applications to efficient computing systems and sustainable energy technology.



Bruno Michel received a Ph.D. degree in biochemistry and computer engineering from the University of Zurich, Switzerland in 1988 and subsequently joined the IBM Zurich Research Laboratory to work on applications of scanning probe microscopy to molecules and thin organic films, and later on micro contact printing and large-area soft lithography. Dr. Michel started a thermal and electrical packaging effort on in 2003 to develop improved thermal interfaces and better miniaturized, bio-inspired microfluidic

cooling. He later demonstrated zero-emission datacenters and interlayer cooled chip stacks, and defined a roadmap towards dense and energy efficient sustainable computers. Technologies with reduced thermal resistance also led to He has published over 300 research articles with a Hirsch factor of 65, and holds more than 160 granted patents. Bruno Michel received the IEEE Harvey Rosten award, is a senior member IEEE, a member of the IBM academy of technology, and a foreign member of the US national academy of engineering. Main current research topics of the Zurich group are focused around development of improved technologies for wearables, edge computing, and cognitive computing to enable improved health-care solutions and human centric sensing and computing.



David Atienza (M'05-SM'13-F'16) is associate professor of electrical and computer engineering, and director of the Embedded Systems Laboratory (ESL) at the Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland. He received his PhD in computer science and engineering from UCM, Spain, and IMEC, Belgium, in 2005. His research interests include system-level design and thermal-aware optimization methodologies for 2D/3D high-performance multi-processor system-on-

chip (MPSoC) and ultra-low power system architectures for wireless body sensor nodes. He is a co-author of more than 250 papers in peer-reviewed international journals and conferences, several book chapters, and seven patents. Dr. Atienza received an ERC Consolidator Grant in 2016, the IEEE CEDA Early Career Award in 2013, the ACM SIGDA Outstanding New Faculty Award in 2012, and a Faculty Award from Sun Labs at Oracle in 2011. He served as DATE 2015 Program Chair and DATE 2017 General Chair. He is a Senior Member of ACM and an IEEE Fellow.