

Integrated Microfluidic Power Generation and Cooling for Bright Silicon MPSoCs

Mohamed M. Sabry Arvind Sridhar David Atienza
Embedded Systems Lab (ESL)
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
{mohamed.sabry,arvind.sridhar,david.atienza}@epfl.ch

Patrick Ruch Bruno Michel
IBM Research – Zurich
Säumerstrasse 4, 8803 Rüschlikon, Switzerland
{ruc,bmi}@zurich.ibm.com

Abstract—The soaring demand for computing power in our digital information age has produced, as an undesirable side-effect, a surge in power consumption and heat density for Multiprocessors Systems-on-Chip (MPSoCs). The resulting temperature rise results in operating conditions that already preclude operating all the cores at maximum performance levels, in order to prevent system overheating and failures. With more power demands, MPSoCs will face a power delivery wall due to the reliability limitations of the underlying power delivery medium. Thus, state-of-the-art power and cooling delivery solutions are reaching their performance limits and it will no longer be possible to power up simultaneously all the available on-chip cores (situation known as *dark silicon*). In this paper we investigate a recently proposed disruptive approach to overcome the prevailing worst-case power and cooling provisioning paradigms for MPSoCs. This proposed approach integrates MPSoC with an on-chip microfluidic fuel cell network for joint cooling and power supply (i.e., localized power generation and delivery). By providing alternative means to power delivery integrated with cooling, MPSoCs are expected to gain in I/O connectivity. Based on this disruptive technology, we can envision the removal of the current limits of power delivery and heat dissipation in MPSoC designs, subsequently avoiding dark silicon and enabling a paradigm shift in future energy-proportional computing architecture designs.

I. INTRODUCTION

Recent processing technology trends have caused a leap from a purely dimension-frequency driven scaling of integrated circuits (ICs) to the new Multiprocessors Systems-on-Chip (MPSoCs). These trends, while continuing to increase the functionality and performance of computers, have undermined three aspects of Dennard scaling [1]: chip-size, energy consumption and heat generation in ICs. These aspects have only exhibited growth in the last two decades as more and more functionality and processing cores are being packed into a single silicon die.

Consequently, energy efficiency of such devices has become a major design concern. There are three major reasons [2, 3] for the reduced energy efficiency in current devices: (1) energy required for communication in large chips vis-a-vis computation, (2) losses in power delivery and (3) energy required

Mohamed M. Sabry and Arvind Sridhar are authors of equal contributions to this work. This work was partly funded by the EC FP7 STREP GreenDataNet project (no. 609000) and the SNF Sinergia project REPCOOL (grant no. 147661). We would like to thank Sani Nassif for his inputs on the target architecture and power figures and the discussions on the power distribution network.

78-3-9815370-2-4/DATE14/©2014 EDAA

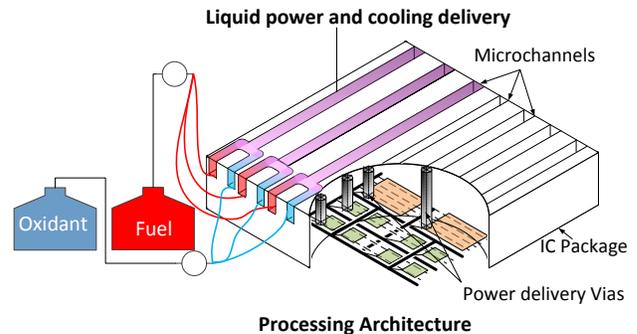


Fig. 1. Schematic diagram of the proposed integrated power and cooling supply MPSoC.

for cooling down ICs to safe operating temperatures. These limitations prevent MPSoCs from operating at maximum performance. Indeed, the upper bound limitations on the power consumption lead to switching off several components during various intervals of operations, which is also known as *dark silicon* [4, 5]. Several research efforts have addressed issue (1) (and implicitly (3)) by increasing on-chip specialized modules and developing heterogeneous architectures [5]. However, this increase in specialization would require additional I/O connectivity for increased bandwidth, and it does not solve the issue of feeding the required power at maximum performance limits. Issue (2) has been addressed by increasing the number of controlled-collapse chip-connection (c4) microbumps dedicated for power delivery in the chip package in order to reduce the effective resistance of on-chip power distribution networks (PDNs) [3]. This decreases the number of pins dedicated for I/O, limiting the off-chip bandwidth of MPSoCs undermining their performance.

The problem of heat removal in devices, which is issue (3), has been recently addressed by integrating active liquid cooling using microchannels directly on the MPSoC die [6]. Single- or two-phase cooling of ICs can dynamically adapt to changes in heat dissipations, reduce energy spent on cooling [6] and enable even denser packaging of devices via 3D stacking of ICs with interlayer cooling of such devices [6–8]. Recent advances in this cooling technology have delivered promising results and given rise to a new idea: to combine

liquid cooling of ICs with on-chip power generation to overcome the drawbacks of the traditional approaches to achieving energy efficiency described above.

Recently, a radically new concept to provide power delivery and cooling in the same fluid medium has been proposed [9, 10]. In this new concept, microfluidic fuel cells (also referred to as redox flow cells) are utilized to generate the required power electrochemically, while the flowing fluid in these cells is used to dissipate the heat generated from the underlying processing platform. Such a disruptive technology can potentially deliver power to *bright silicon* MPSoCs, with little loss, while freeing up the connections to the package fully for I/O links. We propose in this paper, and evaluate for the first time, the concept of an MPSoC with an integrated microfluidic fuel cell array that also functions as a micro-scale heat sink. This concept is illustrated in Fig. 1 where several microchannels etched on top of the silicon die (or multiple stacked dies), that are electrically connected in parallel, constitute a microfluidic fuel cell array. The main contributions and organization of this paper are as follows:

- 1) A brief theory and fundamentals of electrochemical power generation in microfluidic fuel cells is presented. In addition, a numerical model based in COMSOL is developed to simulate the power generation in these microchannels and its accuracy has been validated against existing implementations in the literature (Section II).
- 2) A case study is performed using the concept technology implemented on a realistic MPSoC: the 8-core IBM POWER7+ processor. Electrical simulations are performed to demonstrate the capability of the proposed microfluidic fuel cell array to power the memory units of the processor, while thermal simulations are performed to demonstrate their heat-removal capability (Section III). We show that we can provide up to 6 W of electrical power to feed the memory modules of the target MPSoC, while cooling the whole architecture down to a peak temperature of 41 °C and requiring 4.4 W liquid pumping power.
- 3) The results are discussed and a brief outlook on the future development of the proposed technology is presented (Section IV).

II. OVERVIEW AND FUNDAMENTALS OF INTEGRATED MICROFLUIDIC FUEL CELLS

The proposed technology integrates power and coolant delivery through the same medium, namely a solution containing electrochemically active species flowing through etched microchannels. Unlike the previously-proposed microchannels used solely for cooling [8], these microchannels would carry two different electrolytic fluids that are used in power generation. This kind of structure is commonly referred to as redox (reduction-oxidation) flow cell [11, 12]. Colaminar flow of the fluids in a single microchannel is illustrated in Fig. 2. The two redox fluids are named fuel and oxidant, where the fuel (or reductant) describes the fluid-containing species that transfer electrons to an electrode (anode) and the

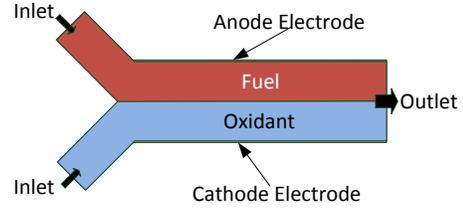
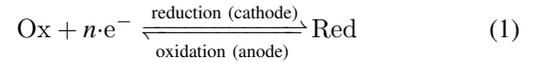


Fig. 2. Planar schematic view of the targeted microchannel structure, highlighting the fluid inlet/outlet ports, electrodes location and the corresponding electrolytes location [12].

oxidant is the fluid containing species that accepts electrons from another electrode (cathode). At each electrode, the redox species involved are converted from an oxidized form (Ox) to a reduced (Red) form or vice versa:



where n is the number of electrons needed to change between the oxidized and reduced states of one redox species.

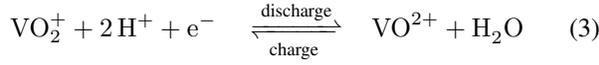
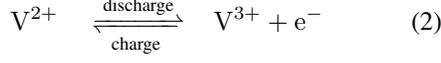
Redox flow cells are a type of secondary battery which stores energy in the electrolytes instead of the electrodes [11]. The independent dimensioning of energy storage capacity (size of electrolyte reservoir) and power density (design of electrochemical cell) have in the past spurred research and applications in the field of grid-scale energy storage. Flow cells in which both the reactants and products are soluble are particularly attractive as a continuous flow of electrolyte ensures a steady energy supply without the build-up of solid deposits or morphological changes of the electrodes. In conventional redox flow cells, a semi-permeable membrane is used to separate the two half-cell compartments to avoid mixing of the two electrolytes. However, for microchannels with small hydraulic diameter D_h , 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 [13]. Thus, no membrane is needed which implies simpler fabrication of flow cells.

The power density per electrode unit area depends on several parameters related to the electrochemical system as outlined in the ensuing subsections. However, all values reported for redox flow cells are typically below 1 W/cm², which is 10-50x lower than the power demand of high-performance processing architectures. A maximum power density of 0.77 W/cm² was reported for a membrane-based all-vanadium redox flow battery with [14]. The highest power density for a membrane-less all-vanadium redox flow cell with flow-through electrodes was reported as 0.3 W/cm² [15]. So far, no attempts have been made to study the provisioning of power to an MPSoC by an integrated electrochemical flow cell.

A. Redox Flow Cell Fundamentals

Various redox species have been investigated for redox flow batteries, with the all-vanadium redox flow battery being the most commercially relevant [11]. This all-vanadium battery was selected as a model system in the present study. In this

electrochemical system, vanadium is used in both half-cells, thereby preventing irreversible capacity loss due to cross-contamination between the half-cells. The electrochemical reactions for charge and discharge for each half-cell are nominally written as:



Reaction (2) takes place at the negative electrode of the electrochemical cell and corresponds to an oxidation of V^{II} (V^{2+} , Red) to V^{III} (V^{3+} , Ox) during discharge. The standard electrochemical potential for this reaction is $E_{neg}^0 = -0.26$ V vs. standard hydrogen electrode (SHE). Reaction (3) occurs at the positive electrode and corresponds to a reduction of V^{V} (VO_2^+ , Ox) to V^{IV} (VO^{2+} , Red) with a standard electrochemical potential of $E_{pos}^0 = +0.99$ V vs. SHE. The equilibrium electrode potentials E for the positive and negative electrodes depend on the actual concentration of vanadium species ($C_i, i \in \{\text{V}^{\text{II}}, \text{III}, \text{IV}, \text{V}\}$) in solution according to the Nernst equation [16, 17]:

$$E_{neg} = E_{neg}^0 + \frac{RT}{F} \ln \frac{C_{\text{V}^{\text{III}}}}{C_{\text{V}^{\text{II}}}} \quad (4)$$

$$E_{pos} = E_{pos}^0 + \frac{RT}{F} \ln \frac{C_{\text{V}^{\text{V}}}}{C_{\text{V}^{\text{IV}}}} \quad (5)$$

with the universal gas constant R , Faraday constant F and temperature T . The open-circuit voltage (OCV) of the full cell is given by $U = E_{pos} - E_{neg}$. The standard OCV of the all-vanadium cell is $U^0 = E_{pos}^0 - E_{neg}^0 = 1.25$ V. During flow of a current I , three main kinds of voltage losses occur within the cell. These losses are referred to as polarization losses or overvoltages [16, 17]. The ohmic overvoltage is given by $\eta_{\Omega} = RI$, where R is the total resistance of the cell comprising both electronic and ionic resistances. The charge-transfer overvoltage η_{ct} is due to the limited rate of the electron transfer reaction at the electrode/electrolyte interface and is described by the Butler-Volmer model of reaction kinetics [16, 17]:

$$I = I_0 \left[\frac{C_{red}}{C_{red}^*} e^{\alpha RT \eta_{ct}/F} - \frac{C_{ox}}{C_{ox}^*} e^{-(1-\alpha) RT \eta_{ct}/F} \right] \quad (6)$$

where the subscripts ox and red denote the oxidized and reduced form of the redox species in one half-cell, C denotes the surface concentration of these species and C^* indicates the redox species concentration in the electrolyte bulk which is unaffected by the reaction at the electrode surface. The exchange current density (I_0) represents the ease of electron transfer and is related to the standard kinetic rate constant k^0 as $I_0 = k^0 F A C_{ox}^{*(\alpha)} C_{red}^{*(1-\alpha)}$ with the electrode area A and transfer coefficient α . Expression (6) also implicitly includes the mass transfer overvoltage η_{mt} , which results from the

depletion of reactants at the electrode surface. According to Reactions (4) and (5), a change in surface concentration of the redox species must result in a change in electrode potential by η_{mt} for the negative and positive electrode, respectively [16, 17]:

$$\eta_{mt,neg} = \frac{RT}{\alpha F} \ln \frac{C_{red}^*}{C_{red}} \quad (7)$$

$$\eta_{mt,pos} = -\frac{RT}{(1-\alpha)F} \ln \frac{C_{ox}^*}{C_{ox}} \quad (8)$$

The overvoltages are additive to give the total voltage loss $\eta = \eta_{\Omega} + \eta_{ct} + \eta_{mt}$. The above equations governing the electrochemical energy conversion can be coupled to the equations for momentum, charge and mass conservation in order to set up a physical model of the redox flow power delivery system. Momentum conservation is provided by the Navier-Stokes and continuity equations for steady-state flow:

$$\rho(\mathbf{v} \cdot \nabla \mathbf{v}) = -\nabla p + \mu \nabla^2 \mathbf{v} \quad (9)$$

$$\nabla \cdot \mathbf{v} = 0 \quad (10)$$

where p is the pressure of the fluid. Charge conservation is imposed for each phase j which may be electronically conductive (electrodes) or ionically conductive (electrolytes):

$$-\nabla \cdot (\sigma_j \nabla \phi_j) = 0 \quad (11)$$

Finally, conservation of mass implies that all species i generated or consumed by the electrochemical reactions at the electrode surface must be transported by diffusion or convection:

$$\nabla \cdot (-D_i \nabla C_i + C_i \mathbf{v}) = S_i \quad (12)$$

with the source term $S_i = I/F$. It is noted that several of the parameters describing the electrochemical and fluid properties are temperature-dependent, such as the kinetic rate constant k^0 and the diffusion coefficients D_i . Therefore, it is expected that the electrochemical performance of an on-chip redox flow system will depend on the heat dissipation of the chip and cooling performance of the fluid.

B. Modeling and Validation

To incorporate the redox flow cell, which provides integrated power delivery and cooling to the target MPSoC, we have created a physical model to extract the corresponding behavior and validated it with experimental data extracted from literature [18]. We have used COMSOL Multiphysics [19] to simulate the underlying physical dynamics described in the previous subsection. To assure the validity of the model, we have created the same microchannel layout of the experimental setup as described in [18]. This microchannel is of 33 mm length, 2 mm width and 150 μm height. The reaction and flow specific parameters are shown in Table I.

Fig. 3 shows the polarization (cell voltage vs. current density) curve that is obtained from the literature and our

TABLE I

LIST OF PARAMETERS [18, 20] USED IN THE APPLIED REDOX FLOW CELL.

Parameter	Anode	Cathode	Unit
Volumetric Flow rate	[2.5, 10, 60, 300]		$\mu\text{L}/\text{min}$
Density	1260		kg/m^3
Dynamic Viscosity	2.53		$\text{mPa}\cdot\text{s}$
Standard Potential (E^0)	-0.255	0.991	V
Oxidant Inlet Concentration (C_O^*)	80	992	mol/m^3
Reductant Inlet Concentration (C_R^*)	920	8	mol/m^3
Diffusion Coefficient (D)	1.7	1.3	$10^{-10} \text{ m}^2/\text{s}$
Kinetic Rate Constant (K^0)	2	1	10^{-5} m/s

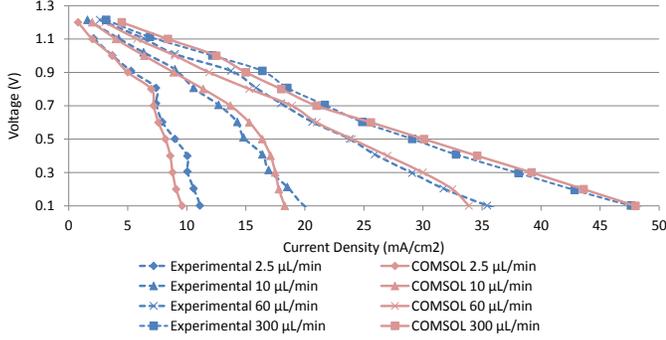


Fig. 3. Polarization curve of the flow cell with the experimental measurements [18] and our developed model on COMSOL [19].

COMSOL model to match these data. The model agrees well with the measurements for different flow rates. We have also measured the accuracy of this model against other flow cells with planar electrodes and we find that the maximum error, in all validation cases, is within 10%. Thus, we conclude that the model accurately simulates the electrochemical performance of the flow cells.

III. TARGET MPSOC INTEGRATION AND IMPACT

To examine the impact of integrating redox flow cells for power delivery and cooling, we use the IBM POWER7+ [21] 8-core MPSoC for a case-study in this work. This $21.34 \text{ mm} \times 26.55 \text{ mm}$ MPSoC (Fig. 4) has a peak power consumption density of $26.7 \text{ W}/\text{cm}^2$.

To integrate redox flow cell to the target MPSoC, we need to investigate two aspects. First, we should explore the connectivity of the redox flow cells with the MPSoC through the power distribution network (PDN) in subsection III-A. Next, we explore the heat removal capability of the redox flow cells, and study the impact of heat dissipation of the MPSoC and increased temperature on the power generation in the flow cells in subsection III-B.

A. Connectivity with Power Delivery Network

An important challenge is interfacing the redox flow cells with the MPSoC power distribution network (PDN). Since the electrode potentials generated in the electrochemical cells depend upon the thermodynamic equilibria of the specific oxidation-reduction reactions taking place inside the cell, voltage regulator modules (VRMs) would be necessary to translate these cell potentials to the voltage-levels required

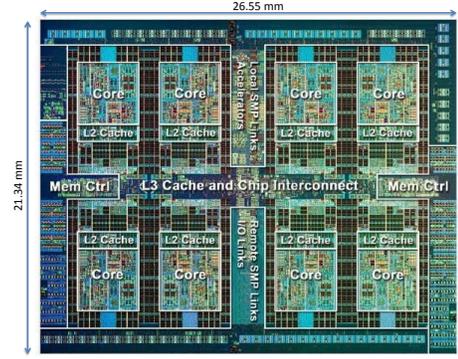


Fig. 4. Floorplan of the targeted IBM POWER7+ MPSoC.

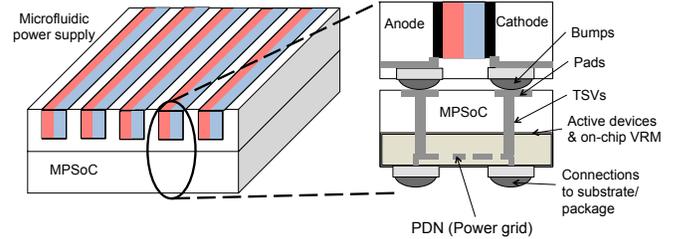


Fig. 5. Combined microfluidic power supply and silicon interposer-based voltage regulator module (VRM) architecture for an MPSoC.

to power the different modules inside the MPSoC. As the power generation in the proposed technology happens inside the package of the chip, these voltage regulators must also be placed within the package to minimize power losses incurred in conversion and transport.

Through-silicon vias (TSVs) and wafer bonding techniques traditionally used in 3D integration of ICs can be used to connect the flow cells electrodes to the on-chip PDNs of the substrate facing MPSoC as illustrated in Fig. 5. Advances made in the recent years in developing on-chip voltage regulators using switched capacitors [22] or buck converters [23] can enable the implementation of integrated voltage regulators in the proposed technology. In case of buck converters, intermediate silicon interposer layers may be needed for the implementation of passive components such as inductors. But such implementations would bring with them an added penalty of increased thermal resistance between the MPSoC and the microchannel heat sinks, thereby limiting this particular benefit of the proposed technology. The output of the VRMs can then be connected to the on-chip power grid of the MPSoC. Multiple such VRMs can provide different voltage levels to the different logic blocks in the MPSoC. This concept is illustrated in Fig. 5 and Fig. 6.

We perform power grid simulations, based on the concept described above, for the targeted IBM POWER7+ MPSoC [21]. Since the current state-of-the-art microfluidic power generation is limited to power densities of the order of $0.7 \text{ W}/\text{cm}^2$ (cf. Section II), only the L2 and L3 cache memory units in the processor were powered using the microfluidic cells, assuming the rest of the logic blocks in the MPSoC were powered using conventional supplies external to the package

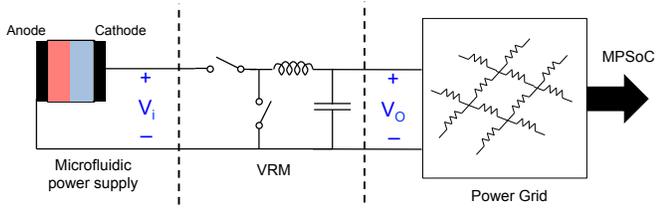


Fig. 6. Schematic diagram of the proposed power distribution network for an MPSoC using microfluidic power supply and silicon interposer-based VRM.

TABLE II

LIST OF PARAMETERS [20, 24] USED IN THE MICROFLUIDIC REDOX CELL ARRAY CONNECTED TO THE IBM POWER7+ CHIP.

Parameter	Anode	Cathode	Unit
Number of channels	88		
Channel width	200		μm
Channel height	400		μm
Channel pitch	300		μm
Channel length	22		mm
Volumetric flow rate (total) \dot{V}	676		ml/min
Thermal conductivity	0.67		W/(m · K)
Thermal capacitance	4.187		$10^6 \text{ J}/(\text{m}^3 \cdot \text{K})$
Inlet temperature	300(27)		K($^{\circ}\text{C}$)
Density ρ	1260		kg/m ³
Dynamic viscosity μ	2.53		mPa · s
Standard potential E^0	-0.255	1.0	V
Oxidant inlet concentration (C_{Ox}^*)	1	2000	mol/m ³
Reductant inlet concentration (C_{Red}^*)	2000	1	mol/m ³
Diffusion coefficient (D)	4.13	1.26	$10^{-10} \text{ m}^2/\text{s}$
Kinetic rate constant (k^0)	5.33	4.67	10^{-5} m/s

supplied using copper interconnects. The cache memories in this architecture consume an average power density of $1 \text{ W}/\text{cm}^2$, which translates to a total current requirement of 5 A for this chip at a supply voltage of 1 V. A microfluidic flow cell array of 88 channels was assumed to have been laid out on the channel layer for this chip. The various structural, chemical and flow properties of this flow cell array are tabulated in Table II.

COMSOL simulations of this microfluidic cell array provide the relationship between voltage levels and current driving capability of this system, plotted in Fig. 7. As can be seen, at a supply voltage of 1 V, the proposed microfluidic flow cell array can provide a current of 6 A, which is adequate to power up the caches. The voltage distribution in the chip as a result of connecting this microfluidic cell array to the power grid (that feeds only the cache regions) is shown in Fig. 8. Future advances in bridging the performance gap between microfluidic cells and the MPSoCs can result in the entire chip being powered by the proposed technology.

B. Heat Dissipation Potential and Impact on Redox Flow Cell

In addition to power generation, the redox fluids also serve as the heat dissipation medium. Thus, it is important to explore the cooling potential of the flow cell array and the corresponding impact on the power generation pattern. In this subsection, we explore the flow cell operation in non-isothermal operating conditions, where the temperature affects the polarization curve of the flow cell (cf. Fig. 7). In this analysis, we use both COMSOL [19] and the compact thermal

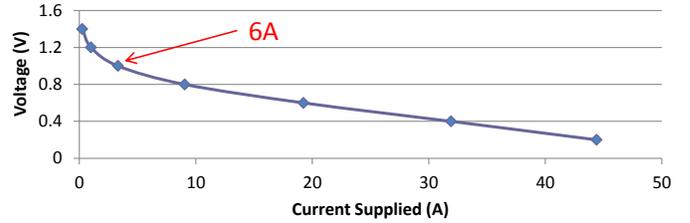


Fig. 7. Voltage-current characteristic of the microfluidic flow cell array consisting of 88 channels (described in Table II).

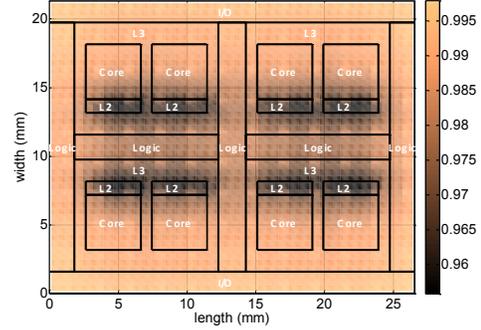


Fig. 8. Voltage distribution (in V) in the power grid circuit supplying power to the cache memory units of the IBM POWER7+ processor using the proposed microfluidic cell array.

model 3D-ICE [7] for system-level evaluation.

Fig. 9 shows the thermal map of the targeted MPSoC when the electrolytes are injected at the flow rate mentioned in Table II and the MPSoC is operated at full load. With the applied flow rate, which is translated to an average flow velocity of 1.4 m/s, the target MPSoC is kept at relatively low temperatures, with 41°C peak value. It is important to mention however that this flow rate corresponds to a pressure drop of 1.5 bar/cm, which is comparable with the pressure drops used for microchannel liquid cooling in the literature [7].

To assess the energy efficiency of the target system, we calculate the pumping power needed to inject the electrolytes at the indicated flow rates (cf. Table II). Based on Darcy Weisbach pressure drop equation and Bernoulli's pumping power equation (assuming a $\eta_p = 50\%$ efficiency pump [6]), we find that the pumping power needed is $P = \frac{\Delta p \cdot \dot{V}}{\eta_p} = 4.4 \text{ W}$. Thus, we demonstrate the dual benefits of power generation and heat removal using the proposed technology. In fact, the results in Subsection III-A shows that the flow cells generate more energy than the value spent in liquid pumping.

In addition to the heat removal capability of the flow cells, it is important to study the effect of the corresponding rise in flow cell temperatures in the generated power. When using the flow rate in II, our simulations indicate that the polarization curve (Fig. 7) does not show significant sensitivity to changes in temperature on the MPSoC. In fact, we observe a maximum 4% increase of the generated current at a fixed potential. Despite the temperature-sensitivity of several parameters (kinetic reaction rate, diffusion coefficient, electrolytic conductivity, density, dynamic viscosity, and the transfer coefficient [20]), the relatively small increase in temperature could not trigger

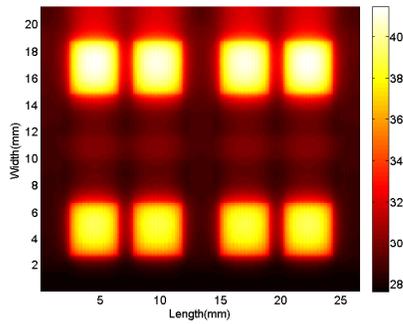


Fig. 9. Thermal response (in °C) of the targeted MPSoC while cooled with a redox flow cell array.

any significant change in the parameters. However, if the flow rates are reduced to enforce significant change in temperature (48 ml/min), or the fluid inlet temperature is increased to 37 °C, the corresponding generated power is increased by up to 23% due to the combined enhancement of diffusion coefficient and reaction rate. This observation actually increases the power generated in the flow cells, opens up a promising research direction where high temperatures, that are normally seen as a detriment in high-performance devices, can actually be used to enable *bright silicon* operation of such devices.

IV. OUTLOOK AND CONCLUSION

In this paper we have demonstrated the potential of addressing cooling and power delivery issues, caused by limitations in Dennard scaling, using a disruptive approach for integrated power generation and cooling. Indeed, this new approach uses the coolant fluid as the means of delivering energy to the chips. Electrochemical systems are still challenged to satisfy the power demand of a microprocessor. In addition, this approach affects the computing efficiency implicitly, due to the consequential I/O and wiring reduction for power delivery to the computational modules. Thus, more wiring area and I/O will be feasible for enhancing the inter-module communication and I/O connectivity.

To quantify the potency of this approach, electrochemical measurements on vanadium redox couples such as V^{II}/V^{III} and V^V/V^{IV} were used to extract fundamental data for mass transport and reaction kinetics to provide input for the modeling effort in this paper. Implementation in analytical and numerical models of flow-through microfluidic cell geometries allowed an assessment of the power density as function of channel dimensions, flow rate and temperature. The results of the simulation show that it is possible to provide power for the memory modules of the IBM POWER7+ MPSoC while efficiently cooling the whole target MPSoC. However, the state-of-the-art is yet not capable of providing power for the logic modules of high-performance processor designs.

In conclusion, to allow a full electrochemical power supply of chip stacks from the same footprint a two pronged approach is needed: (1) The power density of processors has to be reduced by improved architectures that minimize data motion and (2) the power density of electrochemical power delivery has to be massively improved. An example of the first effort is in the increased on-chip specialization to provide less

power hungry compute systems that are based on efficiency improvements and educated compromises in computation and performance. Such devices can now also be used as servers for memory bound applications with better efficiency compared to current servers [25].

For the second effort experimental investigations of electrochemical conversion in single-channel redox flow systems have been triggered to explore approaches to overcome power density limitations. We believe that future computers which additionally could be built around this fluidic power delivery scheme would be less power-intensive due to their reduced communication power demand. The fluidic means of removing heat allows considerable increases in packaging density while the fluidic delivery of power with the same medium saves the space used by hardwired power delivery. The savings in space allow much denser architectures, sharply reduced energy needs and latency as communication paths shrink. Together with the improved power densities of redox flow batteries it should be possible to deliver the complete power demand of high-performance, high-density, processing architectures.

REFERENCES

- [1] R. H. Dennard et al. Design of ion-implanted MOSFETs with very small physical dimensions. *IEEE Journal of Solid-State Circuits*, 9(10), 1974.
- [2] Simon Moore and Daniel Greenfield. The next resource war: computation vs. communication. In *SLIP*, 2008.
- [3] S. L. Wright et al. Characterization of micro-bump c4 interconnects for si-carrier sop applications. In *ECTC*, 2006.
- [4] H. Esmailzadeh et al. Dark silicon and the end of multicore scaling. In *ISCA*, 2011.
- [5] M. B. Taylor. Is dark silicon useful? harnessing the four horsemen of the coming dark silicon apocalypse. In *DAC*, 2012.
- [6] M. M. Sabry et al. Towards thermally-aware design of 3d mpsoCs with inter-tier cooling. In *DATE*, 2011.
- [7] A. Sridhar et al. 3D-ICE: A compact thermal model for early-stage design of liquid-cooled ics. (to appear in) *IEEE Transactions on Computers*, 2014.
- [8] T. Brunschweiler et al. Interlayer cooling potential in vertically integrated packages. *Microsyst. Technol.*, 15(1):57 – 74, 2009.
- [9] P. Ruch et al. Toward five-dimensional scaling: How density improves efficiency in future computers. *IBM Journal of Research and Development*, 55(5):15:1–15:13, 2011.
- [10] P. Ruch et al. Roadmap towards ultimately-efficient zeta-scale datacenters. In *DATE*, 2013.
- [11] M. Skyllas-Kazacos et al. Progress in flow battery research and development. *Journal of the Electrochemical Society*, 158:R55–R79, 2011.
- [12] S. A. M. Shaegh et al. review on membraneless laminar flow-based fuel cells. *International Journal of Hydrogen Energy*, 36:5675–5694, 2011.
- [13] E. Kjeang et al. Microfluidic fuel cells: A review. *Journal of Power Sources*, 186:353–369, 2009.
- [14] Q. Liu et al. High performance vanadium redox flow batteries with optimized electrode configuration and membrane selection. *Journal of the Electrochemical Society*, 159(8):A1246–A1252, 2012.
- [15] J.W. Lee et al. Microfluidic redox battery. *Lab On A Chip*, 13:2504–2507, 2013.
- [16] A.J. Bard and L.R. Faulkner. *Electrochemical methods: Fundamentals and applications*. John Wiley & Sons, 2001.
- [17] C.H. Hamann and W. Vielstich. *Electrochemistry*. Wiley-VCH, 2005.
- [18] E. Kjeang et al. Planar and three-dimensional microfluidic fuel cell architectures based on graphite rod electrodes. *Journal of Power Sources*, 168(2):379–390, 2007.
- [19] COMSOL Multiphysics simulation infrastructure version 4.3.a. <http://www.comsol.com/>.
- [20] T. Rapp. Electrochemical energy conversion for high power densities in microfluidic fuel cells. Master's thesis, ETH Zürich, 2012.
- [21] IBM POWER7+. <http://www-03.ibm.com/press/us/en/pressrelease/29315.wss>, 2010.
- [22] T. M. Andersen et al. A 4.6w/mm² power density 86% efficiency on-chip switched capacitor dc-dc converter in 32nm soi cmos. In *APEC*, 2013.
- [23] K. Onizuka et al. Stacked-chip implementation of on-chip buck converter for distributed power supply system in sips. *Solid-State Circuits, IEEE Journal of*, 42(11):2404–2410, 2007.
- [24] H. Al-Fetlawi et al. Non-isothermal modelling of the all-vanadium redox flow battery. *Electrochimica Acta*, 55:78–89, 2009.
- [25] R. Luijten and A. Doering. The dome embedded 64 bit microserver demonstrator. In *ICIDT*, 2013.