Dynamic Power Management for Neuromorphic Many-Core Systems

DOI:

[10.1109/TCSI.2019.2911898](https://doi.org/10.1109/TCSI.2019.2911898)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Hoppner, S., Vogginger, B., Yan, Y., Dixius, A., Scholze, S., Partzsch, J., Neumärker, F., Hartmann, S., Schiefer, S., Ellguth, G., Cederstroem, L., Plana, L. A., Garside, J., Furber, S., & Mayr, C. (2019). Dynamic Power Management for Neuromorphic Many-Core Systems. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(8), 2973-2986. <https://doi.org/10.1109/TCSI.2019.2911898>

Published in:

IEEE Transactions on Circuits and Systems I: Regular Papers

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

Dynamic Power Management for Neuromorphic Many-Core Systems

Sebastian Höppner¹, Bernhard Vogginger¹, Yexin Yan¹, Andreas Dixius, Stefan Scholze, Johannes Partzsch¹, Felix Neumärker, Stephan Hartmann, Stefan Schiefer, Georg Ellguth, Love Cederstroem, *Member, IEEE*, Luis A. Plana², *Senior Member, IEEE*, Jim Garside², Steve Furber², and Christian Mayr

Abstract—This paper presents a dynamic power management architecture for neuromorphic many-core systems, such as SpiNNaker. A fast dynamic voltage and frequency scaling (DVFS) technique is presented which allows the processing elements (PEs) to change their supply voltage and clock frequency individually and autonomously within less than 100 ns. This is employed by the neuromorphic simulation software flow, which defines the performance level (PL) of the PE based on the actual workload within each simulation cycle. A test chip in 28-nm SLP CMOS technology has been implemented. It includes four PEs which can be scaled from 0.7 to 1.0 V with frequencies from 125 to 500 MHz at three distinct PLs. By measurement of three neuromorphic benchmarks, it is shown that the total PE power consumption can be reduced by 75%, with 80% baseline power reduction and a 50% reduction of energy per neuron and synapse computation, all while maintaining temporary peak system performance to achieve biological real-time operation of the system. A numerical model of this power management model is derived which allows DVFS architecture exploration for neuromorphics. The proposed technique is to be used for the second-generation SpiNNaker neuromorphic many-core system.

Index Terms—MPSoC, neuromorphic computing, SpiNNaker2, power management, DVFS, synfire chain.

I. INTRODUCTION

NEUROMORPHIC circuits [1] try to emulate aspects of neurobiological information in semiconductor hardware in order to solve problems that biology excels at, for example robotics control, image processing or data classification. Furthermore, it represents an alternative to general-purpose high-performance computing for large-scale brain simulation [2].

Manuscript received November 8, 2018; revised March 5, 2019 and March 27, 2019; accepted April 13, 2019. This work was supported in part by the European Union Seventh Framework Programme (FP7) under Grant 604102, in part by the EU's Horizon 2020 Research and Innovation Programme through the Human Brain Project (HBP) under Grant 720270 and Grant 785907, and in part by the Center for Advancing Electronics Dresden (CFAED). This paper was recommended by Associate Editor P. K. Meher. (Corresponding author: Sebastian Höppner.)

S. Höppner, B. Vogginger, Y. Yan, A. Dixius, S. Scholze, J. Partzsch, F. Neumärker, S. Hartmann, S. Schiefer, G. Ellguth, L. Cederstroem, and C. Mayr are with the Faculty of Electrical and Computer Engineering, Technische Universität Dresden, 01062 Dresden, Germany (e-mail: sebastian.hoepfner@tu-dresden.de).

L. A. Plana, J. Garside, and S. Furber are with the Advanced Processor Technologies Research Group, The University of Manchester, Manchester M13 9PL, U.K.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSI.2019.2911898

Besides system capacity, energy efficiency is one major scaling target, especially when large scale brain models are to be simulated with reasonable power consumption, which is typically limited by the effort for cooling and power supply. Energy efficiency is mandatory for the application in mobile, battery powered scenarios such as drones or mobile robots.

There exist approaches for low power and energy efficient neuromorphic circuits using analog subthreshold circuits [3]–[6] or in recent years memristors [7], [8]. However, these systems show severe device variability and it is challenging to scale them to larger systems. In contrast, digital neuromorphic hardware systems such as TrueNorth [9], Loihi [10] or SpiNNaker [11] emulate neural processing by means of digital circuits or embedded software. Due to their purely digital realization they can be implemented in the latest CMOS technologies, operate very reliably/reproducibly and can be scaled to large system sizes.

To make digital neuromorphic systems as energy-efficient as analog ones, we can take inspiration from biology: The brain seems to maximize the ratio of information transmitted/computed to energy consumed [12]. That is, it consumes energy proportional to task complexity or activity levels. The brain also seems to use this in a spatial dimension, i.e. energy is allocated to different regions according to task demand [13]. The cost of a single spike limits the number of neurons that are concurrently active to one percent of the brain.

Transferring this concept to digital neuromorphic hardware, where large neural networks are mapped to many processing cores, leads to the requirement for local, fine-grained scalability for the trade-off between processing power and energy efficiency for highly dynamic workloads, i.e. spikes to be processed in biological real time. The temporal demand for high computational performance to be able to process large neural networks in biological real time requires high processor clock frequencies, and thereby the operation at nominal or even overdrive supply voltages, leading to high-power consumption. In contrast, low power operation can only be achieved by operating the processor core at lower supply voltages at the cost of higher logic delays and thereby slower maximum clock frequencies. Dynamic voltage and frequency scaling (DVFS) during system operation can break this trade-off, by lowering supply voltage and frequency during periods of low processor performance requirements and increasing

supply voltage and clock frequency only if peak performance is required temporarily [14].

DVFS is widely used in modern MPSoCs, like for mobile communication [15]–[17] or database acceleration [18]. In the mentioned scenarios DVFS is typically controlled by a task scheduling unit which assigns supply voltage and frequency settings to the worker processing elements dynamically. In contrast, neuromorphic many core systems do not contain a central task management unit, since the neuromorphic application runs in parallel on processing elements [11]. Their actual workload for neuron state updates and synaptic event processing mainly depends on both the neural network topology and the input (e.g. spikes) of the experiment or application. It is therefore not known in advance or by any central control unit. However, the DVFS technique is beneficial for those systems, since neural networks show significant temporal variations of activity, making them very energy efficient for the neuromorphic task to be solved. This technique is to be employed by future digital neuromorphic hardware, such as the second generation of the SpiNNaker [11] system, which is currently in the prototyping phase.

This work presents a DVFS architecture for neuromorphic many core systems, where each processing element can autonomously change its supply voltage and frequency, only determined by its local workload. It presents a complete validation of the proposed concepts from software level down to a silicon prototype. It extends a conference publication [19] by a technique for workload estimation and performance level selection for neuromorphic DVFS and a numerical energy consumption model for architecture exploration. Both aspects are demonstrated by means of chip measurement results for several benchmarks. In Sec. II the architecture of the neuromorphic SoC is presented including the hardware architecture and software flow for DVFS and Sec. III describes the corresponding power consumption model. Sec. IV presents a test chip in 28nm CMOS, with measurement results including 3 neuromorphic benchmarks summarized in Sec. V. Sec. VI shows an exploration of DVFS architecture parameters for future neuromorphic many core systems, such as SpiNNaker2.

II. NEUROMORPHIC SoC ARCHITECTURE

A. Overview

Fig.1 shows the block diagram of the many core system, based on the SpiNNaker architecture [11]. The neuromorphic computation problem is mapped to processing elements (PEs), which are responsible for neuron state computation and synaptic processing. In this work, the PEs are based on ARM M4F processors with local SRAM and a globally asynchronous locally synchronous (GALS) clocking architecture for dynamic power management. Off chip DRAM serves as synaptic weight storage. A SpiNNaker router [20] is used for on-chip and off-chip neuromorphic spike communication. The on-chip components are connected by a packet based network-on-chip (NoC), carrying spike packets, DMA packets for DRAM access and various types of control packets. In the periphery, shared modules are included, such as true random

Fig. 1. System architecture.

Fig. 2. PE DVFS architecture [19].

number generators (TRNGs) [21] and timers for the neuromorphic real time simulation time step (e.g. 1ms) generation.

B. Power Management Hardware Architecture

The power management architecture of the PE is shown in Fig. 2. It is based on the concepts from [22] and [16]. The PE is equipped with a local all-digital phase-locked-loop ADPLL [23] with open-loop output clock generation [24], enabling ultra-fast defined frequency switching. The core domain, including the processor and local SRAM is connected to one out of several global supply rails by PMOS header power switches. If all switches are opened, the core is in power-shut-off. This allows for individual, fine-grained power management and GALS operation.

A PE performance level (PL) consists of a supply voltage and frequency (V_{DD}, f) pair. Switching between PLs is scheduled by the power management controller (PMC) [22]. The supported scenarios include the change between two PLs when the core is active as well as power-shut-off (PSO) and power-on after PSO. Therefore, the PMC schedules a sequence of the events clock disable, supply selection, net pre-charge, frequency selection and clock enable as shown in Fig. 3. The timings are derived from the reference clock signal (period typically 10 ns) and are fully configurable.

When connecting a PE to another supply rail at a new target V_{DD} level, rush currents must be reduced to prevent unwanted supply voltage drops of other PE, being currently

Fig. 3. PE DVFS Timing of performance level change [19].

Fig. 4. Illustration of PE supply switching.

operated at the target rail, as illustrated in Fig. 4. Therefore, a pre-charge scheme is used, where a small (configurable) number of power switches is connected to the target supply net to reduce the slew rate of the switched net [22]. The PE contains 642 header PMOS switches per supply rail. Two of the switches per rail have reduced width by factor of 2 and 4. Therefore their on-resistance R_{on} is guaranteed to be $< 100\Omega$ at 1.0V drain-source voltage difference, limiting the power-on current to $< 10\text{mA}$, ensured that it does not exceed the current consumption of the operational PE. After power-on all switches are enabled to ensure $< 100\text{m}\Omega$ R_{on} for PE operation. Measurement results of this scheme for the 28nm SoC implementation in this work are shown in Sec. IV-B. As result PL changes of active PEs can be achieved within approximately 100ns, i.e., *instantaneously* from the perspective of the software running on the PE.

PL changes are triggered by sending commands via the NoC interface to the PMC. For power-up or remotely controlled DVFS, these commands can be sent by another core, orchestrating system boot-up. During PE operation (e.g. distributed neuromorphic application) the PE can trigger PL changes on its own by sending a NoC packet to its own PMC (*self DVFS*). Thus, the PE software can actively change its PL without significant latency or software overhead. The application specific power management algorithms can be completely implemented in software at the local PEs.

C. Spiking Neural Network Simulation

We follow the approach of SpiNNaker [25] to implement real-time spiking neural networks (SNN). Each core simulates the dynamics of a number of neurons and their inbound synapses (Fig. 5(a)). A real-time tick from a peripheral timer regularly triggers the neuron state updates and synapse processing which must complete within each simulation cycle t_{sys} (e.g. 1 ms) otherwise the neurons may fall out of sync. The spike communication between neurons is established by the

Fig. 5. Spiking neural network simulation architecture. (a) Logical SNN realization. (b) Physical implementation.

SpiNNaker router that forwards multicast packets containing the identifiers of the sending neurons to target PEs according to a configurable routing table. At the target PE the spike events are inserted into a hardware FIFO attached to the local SRAM and processed in the subsequent cycle. Note that this is different to the typical SpiNNaker operation where incoming spikes are processed immediately.

Details about the memory layout for neural processing are shown in Fig. 5(b). Being accessed in every simulation cycle, the neuron state variables and parameters are stored in local SRAM. The synapse parameters, which require more memory and are only needed at an incoming spike, are stored in external DRAM. They are organized in so-called *synapse rows* which are contiguous memory blocks containing the synapses between one source neuron and all neurons of a core [26]. In the synapse row each existing synaptic connection is represented by a 32-bit word with a 16-bit weight, an 8-bit target neuron identifier, one synapse type bit (excitatory/inhibitory), and a 4-bit delay. The size of a synapse row depends on the fan-out of each source neuron, cf. Fig. 5(a).

When processing a received spike event, the core extracts start address and size of the synapse row belonging to the source neuron from a lookup table in the SRAM. Then, a direct memory access (DMA) for reading the synapse row from the external DRAM is scheduled with a dedicated DMA controller. During the DMA transfer the processor is not idle but can execute other tasks like neuron updates from a job queue [27]. Upon completion of the DMA, the synapses are processed and the weights are added to input ring buffers of the

Fig. 6. Software flow for neuromorphic simulation with DVFS [19], PE tasks run sequentially, spikes are received in parallel by the hardware spike FIFO.

target neurons. These ring buffers accumulate the synaptic inputs for the next 15 clock cycles and enable configurable synaptic delays. When calculating the neuron state update, the synaptic input from the buffer corresponding to the current cycle is injected into the neuron model. If neurons have fired, spike events are generated and sent to the SpiNNaker router.

D. Power Management Software Architecture

The computational load in neuromorphic simulations is determined by the neuron state updates and synaptic events. While the neuron processing cost is constant in each simulation cycle, the number of synaptic events to be processed per time and core strongly varies with network activity. Our approach for neuromorphic power management exploits this by periodically adapting the performance level to the current workload. Fig. 6 visualizes the flow of a neuromorphic simulation with DVFS. Within a simulation cycle of length t_{sys} spikes are received by the PE and registered in the hardware spike FIFO. While spikes of cycle k are received those from cycle $k-1$ are processed without interrupting the processor at incoming spikes. At the beginning of each cycle k the workload is estimated based on the spikes in the queue and the performance level is set to the lowest possible level that guarantees the completion of the neural processing within the cycle t_{sys} .

E. Workload Estimation

The workload for neuron processing is constant in each cycle and depends on the number of neurons per core n_{neur} and the cost for each neuron state update c_{neur} (in clock cycles):

$$c_{neur,tot} = n_{neur} \cdot c_{neur} \quad (1)$$

Instead, the cost for synapse processing varies with the number of spike events l received and the fan-out g_i of respective source neurons to the neurons on this core (Fig. 5(a)). The cost for decoding the synapse words and applying the weights to the input buffers of the neurons is given by

$$c_{syn,tot} = \sum_l g_i(l) \cdot c_{syn}, \quad (2)$$

where c_{syn} is the cost for processing one synapse.

In addition, there is a non-negligible overhead $c_{pre-spike}$ for each received presynaptic spike for looking up the synapse row address and the DMA transfer, which we summarize as

$$c_{pre-spike,tot} = l \cdot c_{pre-spike} \quad (3)$$

Hence, the total workload c in clock cycles is given by

$$c = c_{neur,tot} + c_{syn,tot} + c_{pre-spike,tot} + c_{other} \quad (4)$$

where c_{other} subsumes all remaining tasks such as the main experiment control or sending of spike events.

F. Performance Level Selection

In each cycle k the lowest performance level is chosen that achieves the completion of neuron and synapse processing within time step t_{sys} . The PL is determined by comparing the workload c to thresholds $c_{th,1}$ and $c_{th,2}$ representing the compute capacities of PL1 and PL2 (in clock cycles per time step):

$$PL(k) = \begin{cases} PL1, & \text{if } c < c_{th,1} \\ PL2, & \text{if } c_{th,1} \leq c < c_{th,2} \\ PL3, & \text{if } c_{th,2} \leq c \end{cases} \quad (5)$$

Then synaptic event processing and neuron state computation is performed at $PL(k)$. When these tasks are completed after the spike processing time $t_{sp}(k)$ the processor is set back to PL1 and sleep mode (clock gating) is activated. It reads

$$t_{sp}(k) = c(k) \cdot f_{PL,k}. \quad (6)$$

The optimization target for PL selection is to maximize t_{sp} within a single t_{sys} period, since this relates to the usage of the minimum required PL to complete the neuron and synapse processing tasks while maintaining biological real-time operation.

To obtain a good estimate of the workload $c(k)$ at the beginning of the simulation cycle, we must compute $c_{syn,tot}$ and iterate over all spike events in the FIFO and add up their fan-outs which are implicitly contained in the synapse row lookup table as the synapse row sizes. This extra loop over the spike events creates an additional workload that consumes part of the compute performance per simulation cycle and increases the energy demands at the first glance. Yet, the possibility to precisely adapt the performance level to the workload offers great power saving capabilities, and it must be evaluated for each application whether the detailed strategy (Eq. 5) pays off.

For this paper, however, we consider a simple performance level selection model based on the number of received spikes l in the spike queue:

For this paper, however, we consider a simple performance level selection model that compares the number of received spikes l in the spike queue with thresholds $l_{th,1}$ and $l_{th,2}$:

$$PL(k) = \begin{cases} PL1, & \text{if } l < l_{th,1} \\ PL2, & \text{if } l_{th,1} \leq l < l_{th,2} \\ PL3, & \text{if } l_{th,2} \leq l \end{cases} \quad (7)$$

The relation between selected performance level and the thresholds can be seen at the bottom of Fig. 7. This strategy

Fig. 7. Secure strategy for selection of performance level thresholds: Depending on the fan-out of the input neurons per core, the worst-case workload for l received spikes can be calculated assuming that the spikes belong to the l inputs with the largest fan-out. The PL thresholds $l_{th,1}$, $l_{th,2}$ are obtained by the intersections with the compute capacities $c_{th,1}$, $c_{th,2}$ of the PLs. The PL is selected according to Eq. 7. The strategy is illustrated for one core of the synfire chain model investigated in Sec. V-A1.

was used in the preceding work [19] and allows a fast selection of the performance level as l is immediately available at the beginning of the cycle k . In turn, the thresholds $l_{th,1}$ and $l_{th,2}$ must be tuned for each application for the best exploitation of energy savings. Care has to be taken that all spike events can be processed at the chosen PL before the end of the time step. To ensure this, we employ a worst-case strategy for setting the performance level thresholds based on the fan-out of source neurons per core, as illustrated in Fig. 7: For l spikes received, the worst case for the workload is when these spikes belong to the l input neurons with the highest fan-out. As we know the fan-out of each neuron before the simulation, we can set the performance level thresholds according to this worst case: For this, we sort the fan-out of all inputs per core in descending order, and compute the worst-case workload for l spikes received according to Eq. 4. Then, the thresholds $l_{th,1}$ and $l_{th,2}$ are given by the intersection points of the worst-case workload and the compute capacities $c_{th,1}$ and $c_{th,2}$, as sketched in the bottom of Fig. 7. The advantage of this approach is two-fold: On the one hand, it guarantees the completion of synapse processing within the time step, on the other hand it is universal and can be employed to any network architecture. In real applications, however, higher thresholds might suffice when the worst-case that the l input neurons with the highest fan-out fire at the same time never occurs.

The PL selection requires less than 500 clock cycles, which is $< 0.4\%$ of $t_{sys} = 1\text{ms}$ timestep for the lowest PL at 125MHz.

III. POWER AND ENERGY MODEL

To derive a model for power consumption and energy efficiency [28] of the PE within the many core system,

a breakdown into the individual contributors (from application perspective) is required. Based on the definitions from [29], the total PE power consumption is split into baseline power, and the energies for neuron and synapse processing in the simulation cycles, as explained in the following. The parameter extraction is shown in Sec. IV-C.

A. Baseline Power

The baseline power is consumed by the processor running the neuromorphic simulation kernel without processing any neurons or synapses. The t_{sys} timer events are received but trigger no neuron processing or synapse processing of the PEs. The baseline power also includes the PE leakage power, when connected to the particular V_{DD} rails of the PL. The baseline power of PL i is $P_{BL,i}$.

B. Neuron Processing Energy

The neuron processing energy per simulation time step at PL i is modeled by,

$$E_{neur,i} = E_{neur,0,i} + e_{neur,i} \cdot n_{neur} \quad (8)$$

assuming a linear relation with an offset energy $E_{neur,0,i}$ (as extracted in Sec. IV-C) and an incremental neuron processing energy $e_{neur,i}$ per neuron. The total energy depends on the number of neurons n_{neur} mapped to the particular PE, independent from the network activity.

C. Synapse Processing Energy

The synapse energy is the PE contribution caused by processing synaptic events within each simulation time step. The PE synapse energy at PL i is modeled by,

$$E_{syn,i} = E_{syn,0,i} + e_{syn,i} \cdot n_{syn} \quad (9)$$

assuming a linear relation with an offset energy $E_{syn,0,i}$ and an incremental energy $e_{syn,i}$ per synaptic event. The total energy depends on the number of synaptic events n_{syn} within a simulation cycle on the particular PE, and thereby from the network activity.

From these definitions, the total energy consumed within a simulation cycle k of length t_{sys} at PL i reads

$$\begin{aligned} E_{cycle}(k) = & P_{BL,i} \cdot t_{sp} + P_{BL,i} \cdot (t_{sys} - t_{sp}) \\ & + E_{neur,0,i} + e_{neur,i} \cdot n_{neur} \\ & + E_{syn,0,i} + e_{syn,i} \cdot n_{syn} \end{aligned} \quad (10)$$

assuming that after the completion of neuron and synapse processing within that cycle after t_{sp} the PE is set back to PL1 (as illustrated in Fig. 6). This reduces baseline power in that idle time of length $t_{sys} - t_{sp}$. The average power consumption over the total experiment of k_{max} cycles reads

$$P_{avg} = \frac{1}{k_{max} \cdot t_{sys}} \cdot \sum_{k=1}^{k_{max}} E_{cycle}(k) \quad (11)$$

Fig. 8. Test chip block diagram.

Fig. 9. Chip Photo, with major building blocks marked [19].

Fig. 10. Measurement setup, PCB with power supply and up to 4 test chip modules, FPGA board and host PC GUI.

IV. TEST CHIP

A. Overview

A test chip of the SpiNNaker2 neuromorphic many core system has been implemented in GLOBALFOUNDRIES 28 nm SLP CMOS technology. Its block diagram is shown in Fig. 8. It contains 4 PEs with ARM M4F processors and hardware accelerators for exponentials [30] and true random number generators [21]. Each PE includes 128kB local SRAM and the proposed power management architecture with three PLs. 128MByte off-chip DRAM is interfaced by LPDDR2. All on-chip components are connected by a NoC, where the longer range point to point connections are realized using the serial on-chip link [31]. The chip photo is shown in Fig. 9. For lab evaluation a power supply PCB is used which hosts up to 4 chip modules. This is connected to an FPGA evaluation board via SerDes links which then connects to the host PC via standard Ethernet. The setup is shown in Fig. 10, consisting of a power supply PCB hosting up to 4 chip modules and an FPGA board for host PC communication over Ethernet. A graphical user interface running on the host PC allows exploration of the DVFS measurements.

Fig. 11. Processing element (ARM M4F) DVFS measurements. (a) Maximum PE clock frequency. (b) Power and energy per operation.

B. PE Measurement Results

Fig. 11(a) shows the measured maximum PE frequency versus the supply voltage V_{DD} including a polynomial fitting curve. The PE is operational with high yield (including SRAM) down to 0.7V. In this testchip the three PLs for the neuromorphic application are defined as PL1 (0.70 V, 125 MHz), PL2 (0.85 V, 333 MHz) and PL3 (1.00 V, 500 MHz) spanning a range from the nominal supply voltage of this process node of 1.00 V down to the minimum SRAM supply of 0.70 V. A multi-mode-multi-corner (MMC) has been used for physical implementation and sign-off. The design has been implemented for target performance at PL3. Hold timing has been fixed in all corners. The achieved performance at PL1 and PL2 has been analyzed from sign-off timing analyses. Fig. 11(b) shows the scaling of the PE energy-per-task metric and power consumption when scaling V_{DD} and f . The dynamic energy consumption of the operating processor, consisting of internal and switching power of logic gates and interconnect structures dominates in this design implementation. Therefore a fitting of $E_{\text{task}} = e_{\text{task,norm}} \cdot V_{DD}^2$, where $e_{\text{task,norm}}$ is the normalized energy per specific task, is applicable here. By voltage and frequency scaling in the mentioned ranges, the energy consumption per task can be reduced by 50% and the power consumption by 85% relative to the nominal operation point of PL3 (1.00 V, 500 MHz).

The robustness of the PL switching has been analyzed by measurements in which a particular PE (aggressor) performs PL switches to a target supply net while another PE (victim) is actively operating at this target net, performing a task which can be checked for successful execution. The PL switching

Fig. 12. Robustness measurements of PL changes, color maps show the pass-rate of a software test case running of the victim PE, when the aggressor PE is switching its PL, parameters are the number of pre-charge switches and the pre-charge time. (a) PU to 0.7 V. (b) PU to 0.85 V. (c) PU to 1.0 V (always-on-rail). (d) SC from 0.7 V to 0.85 V.

is repeated multiple times. The experiment is repeated for various settings of supply net pre-charge time and number of activated pre-charge switches (see Sec. II-B). Fig. 12 shows the measurement results for various scenarios for power-up (PU) and supply (PL) change (SC).

In Fig. 12(a) PU to the 0.7 V rail (PL1) is measured, which is very robust versus the switching time and the number of pre-charge switches. This is caused by the fact that the 125 MHz frequency setting of PL1 has significant headroom to the maximum possible frequency of 166 MHz (see Fig. 11(a)), thereby tolerating temporary supply drops during switching.

In case of PU to PL2 (Fig. 12(b)), a strong dependency of the robustness on the pre-charge time and number of switches is visible. For a large number of pre-charge switches the victim PE fails due to the large rush current induced voltage drop at PL2. A smaller number of pre-charge switches anyway requires a minimum pre-charge time, which decreases with increasing number of switches. This is caused by the fact that if the pre-charge time is too short, all power switches are activated although the PE supply net has not yet settled to PL2, resulting in a significant rush current and PL2 supply drop directly after the pre-charge phase. The PU behavior to PL3 is similar (Fig. 12(c)) but much more robust, since the 1.0 V supply net PL3 is the always on-domain of the chip toplevel, therefore having high on-chip decoupling capacitance which tolerates larger rush currents during power switching. Fig. 12(d) shows the SC scenario from PL1 to PL2 which is most critical (PL2 has smaller on-die decoupling capacitance than PL3). Also high switching robustness is achieved here. In summary, for safe operation well within the PASS regions

Fig. 13. Neural and synapse energy parameter extraction. The data points at 16,000 synaptic events per ms in the bottom plot also show up in Table III as power figures for the locally-connected network.

of Fig. 12, a setting of 31 pre-charge switches enables save PU in $\approx 1\mu\text{s}$ and SC in $< 100\text{ ns}$.

C. DVFS Power Model Parameter Extraction

The power management model parameters from Sec. III have been extracted. For baseline power P_{BL} extraction the simulation kernel without neuron processing code and synapse processing code has been executed at different PLs, including the t_{sys} trigger to the PE. The neuron processing power has been determined by running the neuron state calculation for different numbers of leaky integrate-and-fire (LIF) neurons with conductance-based synapses as shown in Fig. 13.

TABLE I
MEASURED PARAMETERS OF POWER MANAGEMENT MODEL

	PL1 (0.70 V)	PL2 (0.85 V)	PL3 (1.00 V)
$P_{BL,leak}$ [mW]	8.94	20.03	28.53
P_{BL} [mW]	14.92	37.44	71.17
$E_{neur,0}$ [nJ]	1000	1410	1540
e_{neur} [nJ]	2.19	2.88	3.96
$E_{syn,0}$ [nJ]	730	990	1490
e_{syn} [nJ]	0.45	0.65	0.90

The power for synapse processing has been measured by running the synapse calculation code with a varying number of synaptic events. Therefore, we used a locally-connected network with 80 neurons per core that are connected to all neurons on the same core, but not to other cores. Then we modified the number of spiking neurons per core and time step to vary the number of synaptic events. This network topology is identical to the local network used in [29] and allows the direct comparison of energy efficiency to the first generation SpiNNaker system in Sec. V-C. The results of the power model parameter extraction are shown in Fig. 13. Linear interpolation has been applied to extract the model parameters for neuron energy $E_{neur,0} = P_{neur,0} \cdot t_{sys}$ and e_{neur} and synapse energy $E_{syn,0} = P_{syn,0} \cdot t_{sys}$ and e_{syn} , respectively. All parameters are summarized in Tab. I (for 4 PEs). As expected, the energy per task (e_{neur} or e_{syn}) scales with V_{DD}^2 , while the baseline power at PL1 is only 20% of the power at PL3, which is consistent with the results in Sec. IV-B.

V. RESULTS

A. Benchmark Networks

To show the capabilities of neuromorphic power management we implement three diverse spiking neural networks for execution on the chip. To show the capabilities of neuromorphic DVFS we implement three diverse spiking neural networks for execution on the chip. The benchmarks were selected to cover the wide range of neural activity found in the brain [32], ranging from asynchronous irregular activity over synchronous spike packets to network bursts. In all networks we use leaky integrate-and-fire neurons with conductance-based synapses. Table II lists characteristics of the networks and the thresholds l_{th} of received spikes used for the performance level selection. The thresholds were determined using the worst-case approach described in Section II-F, except for the synfire network, which uses the same numbers as in the preceding work [19] for consistency.

1) *Synfire Chain*: The first benchmark is a synfire chain network [33] which was already used in the preceding work [19]: Synfire chains are feedforward networks that propagate synchronous firing activity through a chain of neuron groups [34]. We implement a synfire chain with feedforward inhibition [33] consisting of 4 groups (Fig. 14(a)), each with 200 excitatory (E) and 50 inhibitory (I) neurons. As in [33] the neurons receive a normally distributed noise current. Excitatory neurons are connected to both excitatory and inhibitory neurons of the next group, while inhibitory neurons only

Fig. 14. Benchmark networks. (a) Synfire chain network [19]. (b) Bursting network. (c) Asynchronous irregular firing network.

TABLE II
BENCHMARK NETWORKS

	synfire	bursting	async
neuron model	LIF + noise	LIF	LIF
neurons per core	250	250	250
synapses per core	20000	25000	10000
avg. fan-out	80	21	8
$l_{th,1}$	20	47	47
$l_{th,2}$	100	214	229

connect to the excitatory population of the same group. There are 25 presynaptic connections per neuron from I to E and 60 presynaptic connections per neuron from E of the previous group, the delays are 8 ms within a group and 10 ms between groups. There are no recurrent connection within a population. We simulate one group per core and connect the last group to the first one. At start, the first group receives a Gaussian stimulus pulse packet generated on core 3 (400 spikes, $\sigma = 2.4$ ms). As shown in Fig. 15, the pulse packet propagates stably from one group to another, where the feedforward inhibition ensures that the network activity does not explode.

As shown in Fig. 15, cores adapt their PLs to the number of incoming spikes within the current 1 ms simulation cycle. Fig. 18(a) shows histograms of the cycles being processed at a particular PL versus t_{sp} . Within some cycles being processed at PL3 many spikes occur simultaneously such that their processing t_{sp} requires up to 0.8 ms, where 1 ms is the real-time constraint. Thus, the system is close to its performance limit. A conventional system without DVFS would have to be operated at PL3. In the DVFS approach only a little percentage of cycles are processed at higher PLs, thereby achieving nearly the energy efficiency of the low voltage operation at PL1.

2) *Bursting Network*: The second benchmark is a sparse random network generating network bursts, where the neurons collectively switch between UP states with high firing rate and DOWN states with low baseline activity. Such network events have been found both *in-vitro* and *in-vivo* at different spatial and temporal scales [35]. The implemented model is based on [36] and consists of $N = 1000$ excitatory neurons recurrently connected with probability $p_{rec} = 0.08$ (Fig. 14(b)).

Fig. 15. Synfire chain benchmark spike train, sent spikes (blue), number of received spikes per core (green), PL selection thresholds (orange) and core PL (violet), plot shows zoom in to 100 ms simulation time for illustration of fast PL changes.

Fig. 16. Bursting network benchmark spike train, sent spikes (blue), number of received spikes per core (green), PL selection thresholds (orange) and core PL (violet). Only core 0 is shown as activity on other cores is nearly identical.

The neurons are equally distributed over the 4 cores. Spike frequency adaptation (SFA) is implemented by creating an inhibitory synapse from each neuron to itself with a long synaptic decay time constant. A background population (BG) of 200 Poisson neurons is connected with $p_{ext} = 0.1$ to the excitatory population to generate a baseline firing activity in the network. Spikes from the BG population are stored before simulation in the DRAM and thus do not add significant workload. Typical network dynamics are shown in Fig. 16: The network quickly enters an UP state with average firing rate higher than 100 Hz until the SFA silences the neurons into a DOWN state until the next network burst is initiated. The histogram of simulation cycles versus t_{sp} is shown in Fig. 18(b). Here the peak performance PL3 is not utilized.

3) *Asynchronous Irregular Firing Network*: A sparse random network with asynchronous irregular firing activity serves as the third benchmark. The limited size of the prototype system does not allow to implement the standard benchmark for sparse random networks [37], which is commonly used to benchmark SNN simulators [38], [39] or neuromorphic hardware [40]. Instead, we use the same network architecture as for the bursting network (Fig. 14(c)) and disable the spike frequency adaption. Additionally, the recurrent connection

Fig. 17. Asynchronous network benchmark spike train, sent spikes (blue), number of received spikes per core (green), PL selection thresholds (orange) and core PL (violet). Only core 0 is shown as activity on other cores is nearly identical.

probability is reduced to $p_{rec} = 0.02$ such that the network stays in a low-rate asynchronous irregular firing regime, as can be seen in Fig. 17. Fig. 18(c) shows the corresponding histogram of simulations cycles versus t_{sp} . In this case only the lowest PL is required. The system automatically remains in its most energy efficient operation mode.

B. Power Management Results

To assess the benefit of the dynamic power management, we compare the power consumption for the benchmarks at the highest performance level (PL3) and using DVFS. The power measurement is done differentially. First the total power P_0 is measured with the neuromorphic experiment running on the chip. Then spike sending is deactivated and the power P_1 is measured. Since no spike is sent, no spike is received and processed. The difference $P_{syn} = P_0 - P_1$ is the synapse processing power. Then all neural processing is deactivated and upon t_{sys} timer interrupt only empty interrupt handlers are called. The power at this stage is measured as P_2 . $P_{neur} = P_1 - P_2$ is the neuron processing power. Then the ARM cores are deactivated and the power is measured as P_3 . $P_{BL} = P_2 - P_3$ is the baseline power. When DVFS is enabled, the PL is determined by the network activity. Thus, after measuring P_0 with the complete software autonomously switching PLs during simulation, the percentage of simulation time at the 3 PLs are recorded and then applied to the simulations when measuring P_1 and P_2 .

Tab. III summarizes the power measurement results of the system for the benchmarks. For comparison, we also include the locally-connected network used for the power model parameter extraction in Sec. IV-C. The testchip as shown in Fig. 8 contains only 4 PEs for prototyping purposes, for which the relative impact of infrastructure power, including the LPDDR2 memory interface is relatively high. It is expected that for future neuromorphic SoCs the infrastructure overhead is somehow balanced with respect to the number and performance of the PEs on the chip. Thus, the further comparison of energy efficiency is focused on the PEs only, which benefit from the proposed DVFS technique. Using DVFS, baseline power can be reduced by $\approx 80\%$, neuron power by up to $\approx 50\%$ and synapse power by between $\approx 35\%$ and $\approx 50\%$, depending on the experiment. The total PE power reduction by means of DVFS is $\approx 73\%$. For comparison with other systems, we also calculate the energy per synaptic event (total energy vs. total synaptic events), which reaches its minimum value of 1.5 nJ PE energy at highest utilization within the locally connected network, similar to the benchmark used in [29].

Fig. 18. Histograms of simulation cycles (1 ms) processed at different PLs versus t_{sp} . (a) Synfire chain network. (b) Bursting network. (c) Asynchronous irregular firing network.

TABLE III
BENCHMARK POWER RESULTS

		local		synfire		bursting		async	
		only PL3	only PL1	only PL3	DVFS	only PL3	DVFS	only PL3	DVFS
power [mW]	total ¹	138.2	72.7	135.6	71.2	136.5	71.6	133.8	67.3
	infrastructure ²	48.2	48.2	48.2	48.2	48.2	48.2	48.2	48.2
	baseline	71.2	14.9	76.2	17.0	76.0	17.0	76.4	14.7
	neuron	3.0	1.7	7.7	3.6	7.8	3.5	7.5	3.6
	synapse	15.9	7.9	3.5	2.4	4.5	2.9	1.7	0.8
	PE	90.0	24.5	87.4	23.0	88.3	23.4	85.6	19.1
SynEvents/s		16,000,000		3,030,000		2,240,000		489,000	
E/SynEvent ³ [nJ]		8.6	4.5	44.7	23.5	61.0	32.0	274	138
E/SynEvent ⁴ (PE only) [nJ]		5.6	1.5	28.8	7.59	39.4	10.5	175	39.1

¹excluding unused test chip components

²timer, router, LPDDR2

³total power divided by SynEvents/s

⁴PE power divided by SynEvents/s

C. Results Comparison

Although this work is focused on a dynamic power management technique for event-based digital neuromorphic systems, thereby not being directly comparable to other neuromorphic approaches, Tab. IV compares the achieved energy consumptions also of different neuromorphic systems. For realistic comparison two metrics for synaptic energy are considered. First, the energy per synaptic event as in total energy divided by synaptic events processed across the system and second the incremental energy for one more synaptic event. These metrics are not scaled to the semiconductor technology, since completely different circuit approaches provide their optimal results in different technology nodes. For a fair comparison these metrics would have to be extracted from the same benchmark running on the different neuromorphic hardware systems [41]. We did not put memristor systems in the table, as the approach seems too different and no large-scale memristor systems have been reported. However, for comparison we would still name a few figures for memristor synapse arrays: Du et al. [7] report measurements of 4.7 pJ for potentiation/depression at a single synapse, but without the circuit overhead. Chakma et al. [42] simulate memristors and CMOS neurons, reporting between 0.1 and 10 pJ/synapticoperation for both learning and passive memristors, with additional 10 pJ/spike for the CMOS neurons. Analog subthreshold systems [3], [4], [6], [43] and mixed-signal systems [44]–[46]

use analog circuits to mimic neural and synaptic behavior. Custom-digital neuromorphic chips [10], [47], [48] use event based processing in custom non-processor units. They can be implemented in nanometer CMOS technologies and show low energy per synaptic event regarding total power in the same order of magnitude as the analog and mixed-signal approaches. Multi-processor based neuromorphic systems, such as this work, trade off much higher system flexibility due to software defined neuromorphic processing by two to three orders of magnitude higher energy consumption. However, the scaling of semiconductor technology together with dynamic power management techniques such as the proposed DVFS in this work reduces this gap. Compared to SpiNNaker approximately 10x reduction of neuron processing energy and PE energy per synaptic event is achieved. For our benchmarks on this work, as summarized in Tab. III, the application of DVFS results in 73% total PE power reduction. The proposed technique is also applicable to custom-digital neuromorphic systems which operate in an event-driven fashion.

VI. DVFS ARCHITECTURE EXPLORATION

Based on a numerical power management model from Sec. III and the extracted model parameters from the 28 nm test chip, an exploration of the DVFS architecture is performed. The synfire chain from Sec. V-A1 is chosen as benchmark, since it shows highly dynamic workload with the temporary demand for peak performance at the highest PL.

TABLE IV
EFFICIENCY COMPARISON OF NEUROMORPHIC HARDWARE SYSTEMS

System	Reference	Type	Techn. node [nm]	Neuron power [nJ/ms]	E/SynEvent [pJ]	E/additional SynEvent [pJ]
Neurogrid	[6]	analog sub-Vt	180		119	n.a.
ROLLS	[4], [5]	analog sub-Vt	180		n.a.	0.077
HiAER-IFAT	[3]	analog sub-Vt	130		50	n.a.
cxQuad	[5], [43] ⁵	mixed-signal	180		46	0.134
BrainScales	[44]	mixed-signal accelerated	180		100	n.a.
Titan	[45], [46] ⁶	mixed-signal	28		18	n.a.
TrueNorth	[47]	custom digital	28		26	n.a.
Odin	[48]	custom digital	28		12.7	8.4
Loihi	[10]	custom digital	14	0.052	n.a.	23.6
SpiNNaker	[29] ⁷	MPSoC	130	26	13300 ⁸ (19300) ⁹	8000
SpiNNaker2 prototype	this work ¹⁰	MPSoC	28	2.19	1500 (4500)	450

For all systems we report the best case with highest energy efficiency found in literature. Note that results are not fully commensurable, especially for the metric considering total power, as different benchmarks are applied.

⁵Synaptic input event of 0.134 fJ is broadcast to all 256 neurons per core. Total power assumes 1024 neurons firing at 30Hz connected to 256 targets each consuming 360 μ W at 1.3 V.

⁶Assuming accelerated operation with speed-up factor of 100, power draw of 15 mW, 128 inputs firing at 1 kHz connected to 64 targets each.

⁷Results extracted from 1st column of Table III in [29], infrastructure power assumed as 11 W per board, cf. Figure 6 of [29]

⁸regarding PE power only

⁹regarding total power

¹⁰at PL1 0.7 V

Fig. 19. Measured power consumption versus DVFS model (synfire chain).

Fig. 19 shows the measured PE power consumption compared to the DVFS model parameters, extracted with a fanout of 80 similar to the synfire chain configuration as shown in Sec. IV-C. The power values match with acceptable accuracy.

For DVFS architecture exploration the number of PLs with dedicated supply rails is considered as parameter, since the hardware overhead of separated supply rails, power switches and external voltage regulators is the main system overhead of the DVFS approach. The chip in this work contains 264 pins of which are 8 PL1 supply, 13 PL2 supply and 26 PL3 supply (where the PL3 supply also connects to other components) For scaled chips using this approach with larger numbers of PEs the required pin count for PE DVFS power supply increases as well. Therefore, the power saving potential versus the number of PLs is analyzed. It is based on the workload threshold method described in this paper, based on 3 PLs. If the number of PLs is lower, e.g. 2, it is directly switched to PL3 when $l_{th,1}$ is achieved. PL2 is omitted in this case. Fig. 20 shows the simulated and measured PE power for different numbers of PLs. They always include the 500 MHz at 1.0 V PL, since this one provides the required peak performance. It can be seen that already with 2 PLs the PE power can be reduced by 70%. Addition of a third level results in 73% power reduction. A hypothetical 4th PL has been added to a scheme of (0.70V, 0.80V, 0.90V, 1.0V) with (125MHz, 300MHz,

Fig. 20. PE power prediction (baseline power, neuron power, synapse power) of synfire chain benchmark for different number of PLs and one scenario with an additional DFS of 10 MHz level at the lowest available PL.

400MHz, 500MHz) and analyzed using the model. Adding this results in only 1% additional power reduction compared to three PLs. From this it is concluded that more than 2 PLs with distinct supply rails do not gain much additional efficiency, justifying their additional overhead.

Without the insertion of an additional supply rail, a dynamic frequency scaling (DFS) performance level can be added. This is based on the supply rail selection of the lowest PL but has a much lower ADPLL clock generator frequency setting. The PE can switch to this DFS level, after the neuron and synapse computation within a simulation time step is done after t_{sp} . This allows the reduction of baseline power after the computation is done, while still allowing the core to react on interrupts. Fig. 20 shows an example DFS analysis result where an additional DFS level with 10 MHz clock frequency is assumed. This reduces the total PE baseline power after t_{sp} close to the leakage power value of $P_{BL,leak,1} = 8.94\text{mW}$, $P_{BL,leak,2} = 20.03\text{mW}$ and $P_{BL,leak,3} = 28.53\text{mW}$, respectively. Due to the high baseline power portion in this particular implementation, this results in 62% power reduction at the highest PL. Adding one or two more PLs, power can be reduced by additionally 45% and 49%, respectively.

VII. CONCLUSION

A DVFS power management approach for event-based neuromorphic real-time simulations on MPSoCs has been presented. Its effectiveness has been demonstrated with a 28 nm CMOS prototype. For a neuromorphic benchmark application, the PE power including baseline power and energy consumption for neuromorphic processing can be significantly reduced by up to $\approx 73\%$ compared to non-DVFS operation while maintaining biological real-time operation. This approach is usable in event-based systems, where the workload of processing elements can be predicted. A supply change can be realized robustly in $< 100\text{ns}$, which is just 0.01% overhead for a 1ms timestep. Using the presented neuromorphic power management model, energy consumption of the next generation large scale neuromorphic many core systems can be estimated. It helps to design both the power management hardware architecture, the software flow and the strategy for mapping a neuromorphic problem to the system with energy awareness. The results flow into the SpiNNaker2 neuromorphic many core system, which is currently under development.

ACKNOWLEDGMENT

The authors thank Arm and Synopsys for IP and the Vodafone Chair at Technische Universität Dresden for contributions to RTL design.

REFERENCES

- [1] S. Furber, "Large-scale neuromorphic computing systems," *J. Neural Eng.*, vol. 13, no. 5, 2016, Art. no. 051001.
- [2] S. J. van Albada *et al.*, "Performance comparison of the digital neuromorphic hardware spinnaker and the neural network simulation software nest for a full-scale cortical microcircuit model," *Frontiers Neurosci.*, vol. 12, p. 291, May 2018.
- [3] T. Yu, J. Park, S. Joshi, C. Maier, and G. Cauwenberghs, "65 k-neuron integrate-and-fire array transceiver with address-event reconfigurable synaptic routing," in *Proc. IEEE Biomed. Circuits Syst. Conf. (BioCAS)*, Nov. 2012, pp. 21–24.
- [4] N. Qiao *et al.*, "A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128 k synapses," *Frontiers Neurosci.*, vol. 9, p. 141, Apr. 2015. [Online]. Available: http://www.frontiersin.org/neuromorphic_engineering/10.3389/fnins.2015.00141/abstract
- [5] G. Indiveri, F. Corradi, and N. Qiao, "Neuromorphic architectures for spiking deep neural networks," in *Proc. IEEE Int. Electron Devices Meeting*, Dec. 2015, pp. 4.2.1–4.2.4.
- [6] B. V. Benjamin *et al.*, "Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations," *Proc. IEEE*, vol. 102, no. 5, pp. 699–716, May 2014.
- [7] N. Du *et al.*, "Single pairing spike-timing dependent plasticity in BiFeO₃ memristors with a time window of 25 ms to 125 μs ," *Frontiers Neurosci.*, vol. 9, p. 227, Jun. 2015.
- [8] H. Mostafa, A. Khia, A. Serb, C. G. Mayr, G. Indiveri, and T. Prodromakis, "Implementation of a spike-based perceptron learning rule using TiO_{2-x} memristors," *Frontiers Neurosci.*, vol. 9, p. 357, Oct. 2015.
- [9] F. Akopyan *et al.*, "TrueNorth: Design and tool flow of a 65 mW 1 million neuron programmable neurosynaptic chip," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 34, no. 10, pp. 1537–1557, Oct. 2015.
- [10] M. Davies *et al.*, "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, vol. 38, no. 1, pp. 82–99, Jan. 2018.
- [11] E. Painkras *et al.*, "SpiNNaker: A 1-W 18-core system-on-chip for massively-parallel neural network simulation," *IEEE J. Solid-State Circuits*, vol. 48, no. 8, pp. 1943–1953, Aug. 2013.
- [12] J. J. Harris, R. Jolivet, and D. Attwell, "Synaptic energy use and supply," *Neuron*, vol. 75, no. 5, pp. 762–777, 2012. doi: 10.1016/j.neuron.2012.08.019.
- [13] P. Lennie, "The cost of cortical computation," *Current Biol.*, vol. 13, no. 6, pp. 493–497, Mar. 2003.
- [14] D. Ma and R. Bondade, "Enabling power-efficient DVFS operations on silicon," *IEEE Circuits Syst. Mag.*, vol. 10, no. 1, pp. 14–30, 1st Quart., 2010.
- [15] M. Winter *et al.*, "A 335 Mb/s 3.9 mm² 65 nm CMOS flexible MIMO detection-decoding engine achieving 4G wireless data rates," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2012, pp. 216–218.
- [16] B. Noethen *et al.*, "A 105 GOPS 36 mm² heterogeneous SDR MPSoC with energy-aware dynamic scheduling and iterative detection-decoding for 4G in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2014, pp. 188–189.
- [17] S. Haas *et al.*, "A heterogeneous SDR MPSoC in 28 nm CMOS for low-latency wireless applications," in *Proc. 54th Annu. Design Automat. Conf.*, Jun. 2017, Art. no. 47.
- [18] S. Haas *et al.*, "An MPSoC for energy-efficient database query processing," in *Proc. 53rd ACM/EDAC/IEEE Design Automat. Conf.*, Jun. 2016, pp. 1–6.
- [19] S. Höppner *et al.*, "Dynamic voltage and frequency scaling for neuromorphic many-core systems," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2017, pp. 1–4.
- [20] J. Navaridas, M. L. A. Luján, L. S. Plana, S. Temple, and S. B. Furber, "Spinnaker: Enhanced multicast routing," *Parallel Comput.*, vol. 45, pp. 49–66, Jun. 2015.
- [21] F. Neumarker, S. Höppner, A. Dixius, and C. Mayr, "True random number generation from bang-bang ADPLL jitter," in *Proc. IEEE Nordic Circuits Syst. Conf.*, Nov. 2016, pp. 1–5.
- [22] S. Höppner, C. Shao, H. Eisenreich, G. Ellguth, M. Ander, and R. Schüffny, "A power management architecture for fast per-core DVFS in heterogeneous MPSoCs," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2012, pp. 261–264.
- [23] S. Höppner, S. Haenzsche, G. Ellguth, D. Walter, H. Eisenreich, and R. Schüffny, "A fast-locking ADPLL with instantaneous restart capability in 28-nm CMOS technology," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 60, no. 11, pp. 741–745, Nov. 2013.
- [24] S. Höppner, S. Henker, H. Eisenreich, and R. Schüffny, "An open-loop clock generator for fast frequency scaling in 65 nm CMOS technology," in *Proc. 18th Int. Conf. Mixed Design Integr. Circuits Syst.*, Jun. 2011, pp. 264–269.
- [25] S. B. Furber, F. Galluppi, S. Temple, and L. A. Plana, "The SpiNNaker project," *Proc. IEEE*, vol. 102, no. 5, pp. 652–665, May 2014.
- [26] M. Noack, J. Partzsch, C. Mayr, S. Henker, and R. Schüffny, "Biology-derived synaptic dynamics and optimized system architecture for neuromorphic hardware," in *Proc. 17th Int. Conf. Mixed Design Integr. Circuits Syst.*, Jun. 2010, pp. 219–224.
- [27] T. Sharp, L. A. Plana, F. Galluppi, and S. Furber, *Neural Information Processing: 18th International Conference, ICONIP 2011, Shanghai, China, November 13–17, 2011, Proceedings, Part III*. Berlin, Germany: Springer, 2011, pp. 424–430.
- [28] J. Haid, C. Bachmann, A. Genser, C. Steger, and R. Weiss, "Power emulation: Methodology and applications for HW/SW power optimization," in *Proc. 8th ACM/IEEE Int. Conf. Formal Methods Models Codesign*, Jul. 2010, pp. 133–138.
- [29] E. Stomatias, F. Galluppi, C. Patterson, and S. Furber, "Power analysis of large-scale, real-time neural networks on SpiNNaker," in *Proc. 25th Int. Joint Conf. Neural Netw.*, Dallas, TX, USA, Aug. 2013, pp. 1–8.
- [30] J. Partzsch *et al.*, "A fixed point exponential function accelerator for a neuromorphic many-core system," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2017, pp. 1–4.
- [31] S. Höppner *et al.*, "An energy efficient multi-Gbit/s NoC transceiver architecture with combined AC/DC drivers and stoppable clocking in 65 nm and 28 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 3, pp. 749–762, Mar. 2015.
- [32] G. Buzsáki and K. Mizuseki, "The log-dynamic brain: How skewed distributions affect network operations," *Nature Rev. Neurosci.*, vol. 15, no. 4, pp. 264–278, Apr. 2014.
- [33] J. Kremkow, L. Perrinet, G. Masson, and A. Aertsen, "Functional consequences of correlated excitatory and inhibitory conductances in cortical networks," *J. Comput. Neurosci.*, vol. 28, no. 3, pp. 579–594, 2010.
- [34] M. Abeles, "Synfire chains," *Scholarpedia*, vol. 4, no. 7, p. 1441, 2009. [Online]. Available: http://www.scholarpedia.org/article/Synfire_chains

- [35] G. Gigante, G. Deco, S. Marom, and P. D. Giudice, "Network events on multiple space and time scales in cultured neural networks and in a stochastic rate model," *PLoS Comput. Biol.*, vol. 11, no. 11, Nov. 2015, Art. no. e1004547.
- [36] M. Mattia and M. V. Sanchez-Vives, "Exploring the spectrum of dynamical regimes and timescales in spontaneous cortical activity," *Cogn. Neurodyn.*, vol. 6, no. 3, pp. 239–250, Jun. 2012.
- [37] T. P. Vogels and L. F. Abbott, "Signal propagation and logic gating in networks of integrate-and-fire neurons," *J. Neurosci.*, vol. 25, no. 46, pp. 10786–10795, Nov. 2005.
- [38] R. Brette *et al.*, "Simulation of networks of spiking neurons: A review of tools and strategies," *J. Comput. Neurosci.*, vol. 23, no. 3, pp. 349–398, Jul. 2007.
- [39] F. Zenke and W. Gerstner, "Limits to high-speed simulations of spiking neural networks using general-purpose computers," *Frontiers Neuroinform.*, vol. 8, p. 76, Sep. 2014.
- [40] J. C. Knight and S. B. Furber, "Synapse-centric mapping of cortical models to the spinnaker neuromorphic architecture," *Frontiers Neurosci.*, vol. 10, p. 79, Sep. 2016.
- [41] A. Diamond, T. Nowotny, and M. Schmuker, "Comparing neuromorphic solutions in action: Implementing a bio-inspired solution to a benchmark classification task on three parallel-computing platforms," *Frontiers Neurosci.*, vol. 9, p. 491, Jan. 2016.
- [42] G. Chakma, M. M. Adnan, A. R. Wyer, R. Weiss, C. D. Schuman, and G. S. Rose, "Memristive mixed-signal neuromorphic systems: Energy-efficient learning at the circuit-level," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 8, no. 1, pp. 125–136, Mar. 2018.
- [43] S. Moradi, N. Qiao, F. Stefanini, and G. Indiveri, "A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs)," *IEEE Trans. Biomed. Circuits Syst.*, vol. 12, no. 1, pp. 106–122, Feb. 2018.
- [44] S. Schmitt *et al.*, "Neuromorphic hardware in the loop: Training a deep spiking network on the Brainscales wafer-scale system," in *Proc. Int. Joint Conf. Neural Netw.*, May 2017, pp. 2227–2234.
- [45] C. Mayr *et al.*, "A biological-realtime neuromorphic system in 28 nm CMOS using low-leakage switched capacitor circuits," *IEEE Trans. Biomed. Circuits Syst.*, vol. 10, no. 1, pp. 243–254, Feb. 2016.
- [46] M. Noack *et al.*, "Switched-capacitor realization of presynaptic short-term-plasticity and stop-learning synapses in 28 nm CMOS," *Frontiers Neurosci.*, vol. 9, p. 10, Feb. 2015.
- [47] P. A. Merolla *et al.*, "A million spiking-neuron integrated circuit with a scalable communication network and interface," *Science*, vol. 345, no. 6197, pp. 668–673, Aug. 2014.
- [48] C. Frenkel, M. Lefebvre, J.-D. Legat, and D. Bol, "A 0.086-mm² 12.7-pJ/SOP 64 k-Synapse 256-neuron online-learning digital spiking neuromorphic processor in 28-nm CMOS," *IEEE Trans. Biomed. Circuits Syst.*, vol. 13, no. 1, pp. 145–158, Feb. 2019.

Sebastian Höppner received the Dipl.-Ing. (M.Sc.) degree in electrical engineering and the Ph.D. degree from Technische Universität Dresden, Germany, in 2008 and 2013, respectively. He is currently a Research Group Leader and Lecturer with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. He has experience in designing full-custom circuits for multi-processor system-on-chip (MPSoC), such as ADPLLs, register files, and high-speed on-chip and off-chip links, in academic and industrial research projects. He has been managing the full-custom circuit design and SoC integration for more than 12 MPSoC chips in 65-, 28-, and 22-nm CMOS technology. Currently, he leads the chip design of the SpiNNaker2 neuromorphic computing system within the Human Brain Project (HBP). His research interest includes circuits for low-power system-on-chip in advanced technology nodes, with a special focus on clocking, data transmission, and power management. He has authored or coauthored more than 56 publications, and he holds ten patents (five issued and five pending) in the above fields. He received the Barkhausen Award for his Ph.D. degree.

Bernhard Vogginger received the Diploma degree in physics from the University of Heidelberg, Heidelberg, Germany, in 2010. He is currently pursuing the Ph.D. with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden, Germany, under the supervision of Prof. C. Mayr. He is also a Research Associate with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. His research interests include neuromorphic engineering, neural computation, and deep learning.

Yexin Yan received the Dipl.-Ing. (M.Sc.) degree in electrical engineering from Technische Universität Dresden, Germany, in 2016, where he is currently pursuing the Ph.D. degree with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits. His research interests include hardware–software co-design for low-power neuromorphic applications.

Andreas Dixius received the Dipl.-Ing. (M.Sc.) degree in information systems engineering from Technische Universität Dresden, Germany, in 2014. Since 2014, he has been a Research Assistant with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. His research interests include on-chip timing detection.

Stefan Scholze received the Dipl.-Ing. (M.Sc.) degree in information systems engineering from Technische Universität Dresden, Germany, in 2007. Since 2007, he has been a Research Assistant with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. His research interests include design and implementation of low-latency communication channels and systems.

Johannes Partzsch received the M.Sc. degree in electrical engineering and the Ph.D. degree from Technische Universität Dresden, Germany, in 2007 and 2014, respectively. He is currently a Research Group Leader with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits. He has authored or coauthored more than 45 publications. His research interests include neuromorphic system design, topological analysis of neural networks, and technical application of bio-inspired systems.

Felix Neumärker received the Dipl.-Ing. (M.Sc.) degree in electrical engineering from Technische Universität Dresden, Germany, in 2015. He is currently a Research Associate with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. His research interests include software and circuit design and MPSoCs with a special focus on neuromorphic computing.

Luis A. Plana (M'97–SM'07) received the Ingeniero Electrónico degree (*Cum Laude*) from Universidad Simón Bolívar, Venezuela, and the Ph.D. degree in computer science from Columbia University, USA. He was with Universidad Politécnica, Venezuela, for over 20 years, where he was a Professor of electronic engineering. He is currently a Research Fellow with the School of Computer Science, The University of Manchester, U.K.

Stephan Hartmann received the Dipl.-Ing. (M.Sc.) degree in electrical engineering from Technische Universität Dresden, Germany, in 2007. He is currently a Research Associate with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. His research interests include circuit design with a special focus on FPGA.

Jim Garside received the Ph.D. degree in computer science from The University of Manchester, Manchester, U.K., in 1987, for work in signal processing architecture. His post-doctoral work on parallel processing systems based on Inmos Transputers was followed by a spell in industry writing air traffic control software. Returning to academia gave an opportunity for integrated circuit design work, dominated by design and construction work on asynchronous microprocessors in the 1990s. He has been involved in dynamic hardware compilation, GALS interconnection, and the development of the hardware and software of the SpiNNaker neural network simulators.

Stefan Schiefer received the Dipl.-Ing. (M.Sc.) degree in electrical engineering from Technische Universität Dresden, Germany, in 2008. He is currently a Research Associate with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. His research interests include full system signal and power integrity as well as regulator behavioral modeling, including non-linearities.

Steve Furber received the B.A. degree in mathematics and the Ph.D. degree in aerodynamics from the University of Cambridge, U.K. He spent the 1980s at Acorn Computers, where he was a Principal Designer of the BBC Microcomputer and the ARM 32-bit RISC Microprocessor. He is currently an ICL Professor of computer engineering with the School of Computer Science, The University of Manchester, U.K. Over 120 billion variants of the ARM processor have since been manufactured, powering much of the world's mobile and embedded computing. He moved to the ICL Chair at Manchester in 1990 where he leads research into asynchronous and low-power systems and, more recently, neural systems engineering, where the SpiNNaker project is delivering a computer incorporating a million ARM processors optimized for brain modeling applications. He is a CBE, a fellow of the Royal Society, and a fellow of the Royal Academy of Engineering.

Georg Ellguth received the Dipl.-Ing. (M.Sc.) degree in electrical engineering from Technische Universität Dresden, Germany, in 2004. Since 2004, he has been a Research Assistant with the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden. His research interests include low-power implementation techniques in multi-processor system-on-chip.

Christian Mayr received the Dipl.-Ing. (M.Sc.) degree in electrical engineering and the Ph.D. and Habilitation degrees from Technische Universität Dresden (TU Dresden), Germany, in 2003, 2008, and 2012, respectively. From 2003 to 2013, he has been with TU Dresden, with a secondment to Infineon from 2004 to 2006. From 2013 to 2015, he held a Post-Doctoral Position with the Institute of Neuroinformatics, University of Zurich, and ETH Zürich, Switzerland. Since 2015, he has been the Head of the Chair of Highly-Parallel VLSI-Systems

Love Cederstroem (M'14) received the Civ.ing. (M.Sc.) degree in applied physics and electrical engineering from Linköping University in 2009. In 2013, he joined the Chair of Highly-Parallel VLSI-Systems and Neuromorphic Circuits, Technische Universität Dresden, where he is currently a Research Assistant. His work focuses on system design and methodologies to bridge the chip-package-system boundaries, currently with an emphasis on power delivery.

and Neuromorphic Circuits, TU Dresden, where he is currently a Professor of electrical engineering. He has authored or coauthored over 80 publications and he holds four patents. His research interests include bio-inspired circuits, brain-machine interfaces, AD converters, and general mixed-signal VLSI-design. He has acted as an editor/reviewer for various IEEE and Elsevier journals. His work has received several awards.