Spatiotemporal Pattern Recognition in Single Mixed-Signal VLSI Neurons with Heterogeneous Dynamic Synapses

Mattias Nilsson
EISLAB
Luleå University of Technology
Luleå, Sweden
mattias.1.nilsson@ltu.se

Foteini Liwicki
EISLAB
Luleå University of Technology
Luleå, Sweden

Fredrik Sandin
EISLAB
Luleå University of Technology
Luleå, Sweden

ABSTRACT

Mixed-signal neuromorphic processors with brain-like organization and device physics offer an ultra-low-power alternative to the unsustainable developments of conventional deep learning and computing. However, realizing the potential of such neuromorphic hardware requires efficient use of its heterogeneous, analog neurosynaptic circuitry with neurocomputational methods for sparse, spike-timing-based encoding and processing. Here, we investigate the use of balanced excitatory–inhibitory disynaptic lateral connections as a resource-efficient mechanism for implementing a thalamocortically inspired Spatiotemporal Correlator (STC) neural network without using dedicated delay mechanisms. We present hardware-in-the-loop experiments with a DYNAP-SE neuromorphic processor, in which receptive fields of heterogeneous coincidence-detection neurons in an STC network with four lateral afferent connections per column were mapped by random input-sampling. Furthermore, we demonstrate how such a neuron was tuned to detect a particular spatiotemporal feature by discrete address-reprogramming of the analog synaptic circuits. The energy dissipation of the disynaptic connections is one order of magnitude lower per lateral connection (0.65 nJ vs 9.6 nJ per spike) than in the former delay-based hardware implementation of the STC.

CCS CONCEPTS

• **Hardware → Neural systems**; • **Computing methodologies → Neural networks**.

KEYWORDS

Neuromorphic, Spatiotemporal pattern recognition, Ultra-low-power, Device mismatch, Neural heterogeneity, Temporal code, Excitatory–inhibitory balance

ACM Reference Format:
Mattias Nilsson, Foteini Liwicki, and Fredrik Sandin. 2022. Spatiotemporal Pattern Recognition in Single Mixed-Signal VLSI Neurons with Heterogeneous Dynamic Synapses . In *International Conference on Neuromorphic Systems (ICONS 2022), July 27–29, 2022, Knoxville, TN, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3546790.3546794

ICONS 2022, July 27–29, 2022, Knoxville, TN, USA

ACM ISBN 978-1-4503-9789-6/22/07.
https://doi.org/10.1145/3546790.3546794

1 INTRODUCTION

Neuromorphic engineering [22, 28] deals with the creation and use of physical substrates for information processing and sensing that imitate processes and structures observed in the brain. Following this approach, mixed-signal neuromorphic processors are very-large-scale integration (VLSI) systems that emulate the biophysical dynamics of neurons and synapses [8] in event-driven spiking neural networks (SNNs) [27] and have high potential for ultra-low-power pattern recognition and learning [16, 18, 38]. Such technology can offer a sustainable alternative to deep learning [29, 49] and enable always-on *edge intelligence* [53, 55] in sensing applications, such as wearable biomedical devices [12]. However, such analog-based systems are subject to transistor parameter-variance—so-called "device mismatch" [35, 51]—which gives rise to substantial static parameter variance within uniformly configured populations of neurons and synapses [37], and which poses a challenge to configurability and computational precision. Furthermore, the dynamic real-time operation of the neurosynaptic computational nodes of these processing systems needs to be employed in a way that realizes their potential for sparse, event-driven information processing [23] with spike-timing-based encoding [15, 19, 20, 47, 50]—as opposed to purely rate-based encoding, the use of which rather makes SNNs more or less analogous to the artificial neural networks often used within the deep learning paradigm.

Learning and recognition of spatiotemporal patterns—in contrast to static, purely spatial patterns, or even sequences of such—in artificial systems is an engineering problem for which spike-based neuromorphic computing holds particularly high potential for efficient solutions in terms of energy and latency [16, 21]. A central aspect of such pattern recognition with SNNs is the neural encoding of temporal relations, which may be performed in a number of different ways [44]. One bio-inspired approach to temporal encoding in SNNs is the use of neural spike-propagation delays [1, 50], which have been implemented in or for neuromorphic hardware in the form of axonal delays [33], delay interneurons [42], and emulation of dendrites [4, 24, 52]. However, such implementations come with a substantial cost, as they require allocation of additional hardware resources beyond what is already used to implement the neurons and synapses of a given SNN and sometimes beyond what is readily available in general-purpose neuromorphic hardware. Another relevant concept is that of the spiking time-difference encoder (TDE) [21, 30], which is a versatile neurocomputational primitive for temporal encoding, but which cannot be implemented in general-purpose analog-based neuromorphic systems without dedicated hardware.

The challenge to using mixed-signal neuromorphic hardware posed by the imprecision stemming from device mismatch has previously been addressed with usage of relatively costly compensation methods, such as circuit calibration procedures [26, 32] or neural resource redundancy [25]. Other mitigation methods include the use of neural–architectural robustness to noise and variability [7] and special training procedures that promote robustness during inference [6]. In contrast, however, some current approaches to spatiotemporal pattern recognition with mixed-signal neuromorphic hardware are—rather than attempting to mitigate the variance caused by device mismatch—based on neural processing frameworks that actually rely on variability in the neurosynaptic elements [3, 14, 17, 34, 39, 43], such as reservoir computing and ensemble learning. Similarly, Markov chain Monte Carlo sampling of heterogeneous memristors has been used to train a Bayesian machine learning model [13]. Furthermore, there is recent research suggesting that the existence of neurosynaptic variability, which is also observed in the brain, may support efficiency and robustness in neural processing and learning—especially for information that has a rich temporal structure [36, 54].

Here, we follow the approach of making use of device mismatch as a source of variance for resource-efficient employment of inherently inhomogeneous mixed-signal neuromorphic hardware. We investigate usage of a previously proposed form of excitatory–inhibitory (E–I) balanced disynaptic elements [40] as a heterogeneity-dependent low-resource alternative to dedicated delay-mechanisms for coincidence-detection-based spatiotemporal pattern recognition with a modified version of the Spatiotemporal Correlator (STC) network, which we propose here—the Disynaptic Spatiotemporal Correlator (dSTC). To investigate the feasibility and effectiveness of implementing the dSTC in mixed-signal neuromorphic hardware, we present hardware-in-the-loop experiments with a DYNAP-SE neuromorphic processor [31], in which we characterized spike-timing-based spatiotemporal receptive fields that form in coincidence-detection neurons of the dSTC. Furthermore, we performed tuning of one of these receptive fields by using digital synapse-address reprogramming to sample from the inherent parameter distributions of the analog synapses. Finally, we present an estimate of the difference in energy usage for hardware implementations of the dSTC and STC neural networks, which indicates a reduction by one order of magnitude per lateral connection (0.65 nJ vs 9.6 nJ per spike).

In summary, we present how heterogeneous neurosynaptic dynamics in neuromorphic processors like the DYNAP-SE can be used in a resource-efficient manner for spike-timing-based spatiotemporal pattern recognition in a way that enables synapse-address reprogramming as a discrete mechanism for feature tuning, which results in observable and reproducible state changes. This approach may serve as a complement to more accurate but resource-intensive delay-based coincidence detection or dendritic integration and use of volatile plastic synapses, which hamper observability and reproducibility. Our study also contributes a new perspective on the STC's mechanism and capability for pattern recognition on the level of single neurons—as opposed to the population-level view held in the original studies [11, 43].

2 MATERIALS AND METHODS

The experiments presented in this work were conducted with a DYNAP-SE (Dynamic Neuromorphic Asynchronous Processor – Scalable) [31] from SynSense interfaced in a closed loop with a PC using the software Legacy Samna [48] (formerly CTXCTL). All input stimuli were generated using the built-in spike-generator in the field-programmable gate array (FPGA) of the DYNAP-SE, which generates spike-events according to assigned temporal interspike intervals (ISIs) and virtual source-neuron addresses.

2.1 Neuromorphic Processor

The DYNAP-SE [31] is a reconfigurable, general-purpose, mixed-signal neuromorphic processor that uses subthreshold analog circuits to emulate the biophysical dynamics of neurons and synapses in real time, and asynchronous digital circuits for spike-event transmission according to an address-event representation (AER) protocol. One DYNAP-SE unit comprises four four-core neuromorphic chips—each of which comprises 256 silicon neuron-circuits implementing the Adaptive Exponential Integrate-and-Fire (AdEx) spiking neuron model [5]. Each neuron has a content-addressable memory (CAM) block that can contain up to 64 different addresses, each representing a presynaptic neuronal connection[1], see **Figure 1**. For each connection, dynamic synapses in the form of differential-pair integrator (DPI) circuits [2] are available in four different synaptic types: fast and slow excitatory, and subtractive and shunting inhibitory. The dynamic behaviors of the neuronal and synaptic circuits of the DYNAP-SE are governed by analog circuit parameters that are set by on-chip programmable bias generators, which provide 25 bias parameters independently for each core. In the present work, a certain kind of balanced E–I disynaptic elements [34, 40] are used to generate delayed excitations in the DYNAP-SE, in a way that is further described in Section 2.2.1.

2.2 Spiking Neural Network

The Spatiotemporal Correlator (STC) is a spiking neural network (SNN) for spatiotemporal pattern recognition and learning in neuromorphic hardware—based on coincidence detection of lateral, temporally delayed projections [42]—which was derived from a biologically plausible model of thalamocortical auditory processing [10] and has been implemented in real-time mixed-signal neuromorphic hardware [43, 45]. The originators of the STC argue that the functional principles of the network and of its more complex precursor can be used to produce spike-based feature extractors, which may form the basis of sensory–processing systems based on mixed-signal neuromorphic technology. Furthermore, the original conception and study of the STC was also motivated as a case-study of the neural mechanisms underlying feature selectivity in primary sensory processing—especially in the auditory domain—as well as a study of the intrinsic characteristics of such features. The robustness of the STC to variability in stimulus patterns—which is a prerequisite for most real-world sensing applications—has been further demonstrated in a subsequent study [11].

The STC network, see **Figure 2a**, consists of the following qualitatively distinct neuronal populations:

[1]Each CAM address can match the same local neuron-ID on different cores, which enables multiple presynaptic connections per input circuit.

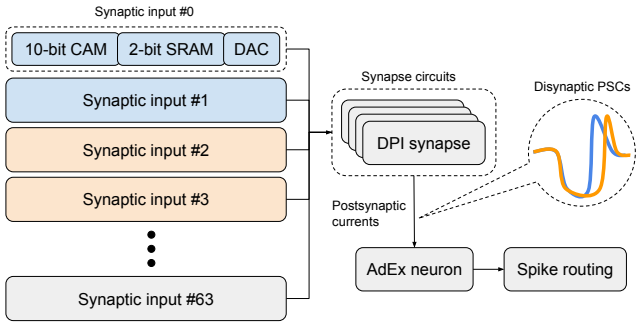

Figure 1: Implementation of balanced excitatory–inhibitory (E–I) disynaptic elements in one neuron of the DYNAP-SE neuromorphic processor. Based on a figure from [31]. Each neurosynaptic node contains 64 synaptic input-channels—each with a CAM cell for the presynaptic neuron address, an SRAM cell for the synapse-type information, and digital-to-analog signal-converting circuitry (DAC)—four analog synaptic DPI circuits of different types, and one analog AdEx-neuron circuit. The coloring indicates the use of two synaptic inputs for each E–I element [34, 40]. Two examples of the resulting heterogeneous postsynaptic currents (PSCs) for different E–I elements are illustrated, which, respectively, use synaptic inputs 0–1 and 2–3. The E–I elements are used for synapse-based generation of delayed excitations for coincidence detection, instead of using dedicated delay lines or emulation of dendrites.

A: Input neurons
B1: Secondary input neurons
B2: Coincidence-detection neurons
C: Delay neurons

The network is structured in columns—not to be confused with cortical columns—each consisting of one A, B1, and B2 neuron, respectively. Within each column, the input neuron, A, receives a signal from an input channel—such as a pixel of a visual sensor, or a frequency band of an auditory sensor—which it then projects to both the B1 and B2 neurons via excitatory synapses. B1 then generates a one-to-one spike-response to its input from A, which is projected by inhibition to the B2 neuron of the same column and by excitation to the B2 neurons of some number of adjacent columns—the number of which is subject to design-choice and pruning. Each lateral B1–B2 excitation is projected via a neuron from the C population, thereby inducing a temporal signal-propagation delay. The A–B2 excitation is slightly faster than the B1–B2 inhibition, thus creating a time-window during which B2 is primed to spike in response to coincident lateral projections from adjacent columns. Hence, each B2 neuron constitutes a coincidence detector that is sensitive to some particular set of spatiotemporal spike-patterns in a local sensory region, which forms a feature of the pattern that is recognized by the STC network as a whole.

2.2.1 The Disynaptic STC Network. Here, we propose a modified version of the STC—the Disynaptic Spatiotemporal Correlator (dSTC), see **Figure 2b**. In the dSTC, the delay interneurons of the STC are replaced with balanced E–I disynaptic elements—first proposed in [40] and further investigated in [34]—thereby theoretically reducing the amount of resources required to implement the network by a substantial amount. Each E–I element consists of one excitatory and one inhibitory synapse connected to the same presynaptic neuron. Thus, Dale's law is explicitly broken to save resources. The time-constants and weights of the two synapses are balanced in such a way that each presynaptic spike generates—by summation of the postsynaptic currents—an inhibition followed by a postinhibitory excitation, thus generating a delayed excitation that imitates postinhibitory rebound [40]. When implemented on different neuronal and synaptic hardware circuits in a mixed-signal system such as the DYNAP-SE, see **Figure 1**, the E–I elements generate varying temporal delays following a roughly Gaussian distribution, as demonstrated in [40] and [34]. This variation is the source of the differences in temporal delay between the lateral afferent connections of B2 neurons in the dSTC architecture. These differences make each B2 neuron sensitive to some range of spatiotemporal spike-patterns that are illustrated as receptive fields in the following, and which are dependent on the heterogeneity of the analog hardware circuits. In this manner, the dSTC performs coincidence-based spatiotemporal feature detection by relying on synaptic integration of the type investigated in [34].

2.3 Mapping of Receptive Fields

We investigated spike-based spatiotemporal receptive fields of the kind described in the previous section in B2 neurons, see **Figure 3b**, that have four lateral connections each—similar to the original STC after training [43]—by implementing a population of B2 neurons in one core of the DYNAP-SE neuromorphic processor. We focused on the forward and lateral connections and, for simplicity, omitted the B1–B2 inhibition, since its function is primarily to regulate spike-timing-dependent plasticity and to prevent sensitivity to excessive stimulation, and is therefore not required for an initial investigation of the receptive fields.

The receptive fields were mapped by stimulating the B2 neurons—independently from each other—with randomized spike-patterns (N = 10,000), each consisting of one spike per input channel, as illustrated by the sample pattern in **Figure 3a**. The spike-times of each of the lateral, delayed projections were independently drawn from a uniform random distribution ranging from 1–50 ms before the incidence of the direct forward excitation from A to B2, which we let define the reference time, $t = 0$, of the pattern. All spike-patterns for which a given B2 neuron generated one or more postsynaptic spikes in response were recorded and aggregated to form an approximation of the receptive field of that neuron.

2.4 Feature Tuning by Synapse Sampling

Based on the results in [34], we also investigated the possibility of tuning one of the receptive fields from the experiment described in Section 2.3 by changing the hardware configuration of the B2 neuron in question. More specifically, we investigated whether replacing the specific heterogeneous hardware circuits used for each synaptic input-connection of the neuron, see **Figure 1**, could affect its coincidence detection enough to make the neuron distinguish between two different prescribed patterns, see **Figure 4**, which

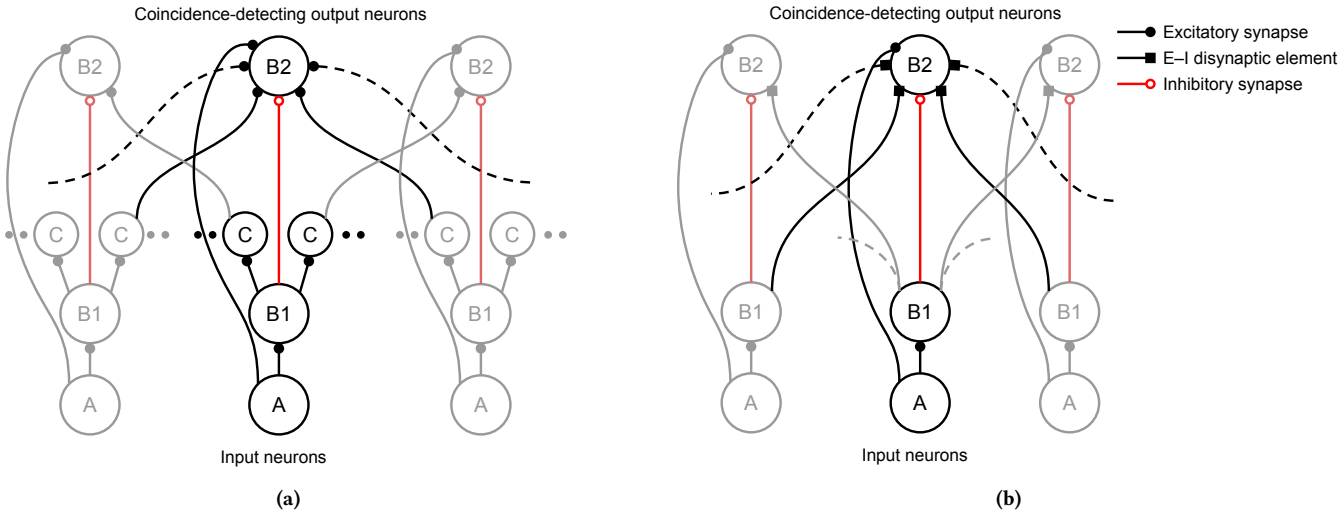

Figure 2: Spatiotemporal Correlator (STC) spiking neural network–architectures. (a) Original STC network, adapted from [45]. (b) Disynaptic STC (dSTC) network, which uses excitatory–inhibitory (E–I) disynaptic elements [34, 40]. For each architecture, one column is highlighted, and the two closest neighboring columns are included in shaded colors to illustrate lateral connections. Circles represent single neurons from populations A, B1, B2, and C.

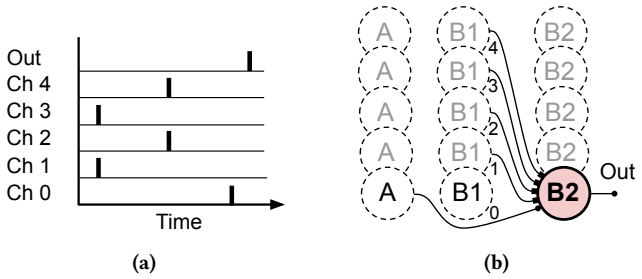

Figure 3: Receptive-field experiment. (a) Example input spike-pattern and response. (b) Feature-detection neuron of dSTC network. The B2 neurons receive spatiotemporal spike patterns on five input channels, of which Channels 1–4 are lateral connections with, in effect, delayed excitations via E–I disynaptic elements.

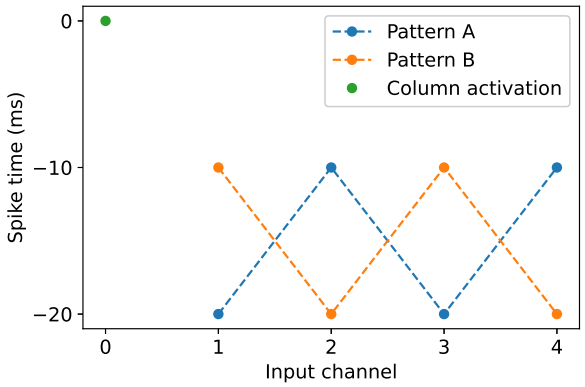

Figure 4: Spatiotemporal spike patterns for feature tuning. Each pattern consists of one spike per input channel, including the feed-forward input at reference time $t = 0$, resulting in five spikes per pattern.

the neuron initially could not. The synaptic reconfiguration was made by assigning to each of the synapses a random, unique input circuit drawn from the set of all the 64 input circuits of the neuron. Following each subsequent synaptic reconfiguration, the neuron was presented with both of the different patterns, separately, ten times—and its response in terms of the number of postsynaptic spikes was recorded.

3 RESULTS

3.1 Receptive Fields

In this section, results are presented from the mapping of receptive fields of B2 neurons implemented in a DYNAP-SE processor, as described in Section 2.3. **Figure 5** illustrates the receptive fields of four different B2 neurons, in the form of box-and-whisker plots, aggregated from presynaptic stimulation with different randomized input patterns (N = 10,000). The plots consist of all the stimulus patterns—that is, channel–spike-time combinations, see **Figure 3a**—that made the B2 neuron in question generate at least one postsynaptic spike in response. **Table 1** presents the hardware neuron IDs—local to the used DYNAP-SE core—for which the receptive fields in **Figure 5** were observed.

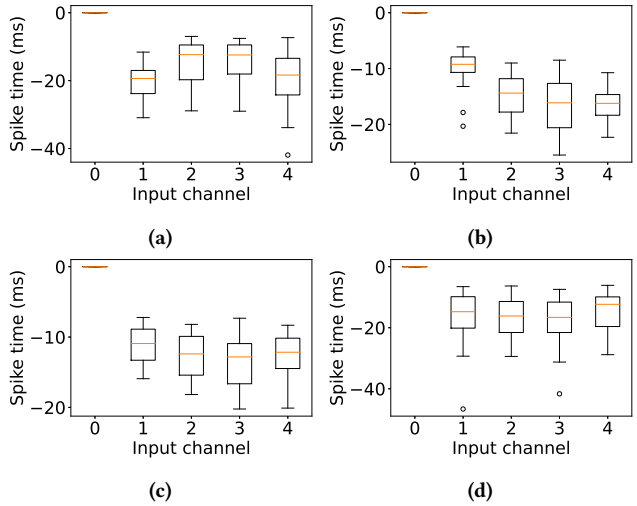

Figure 5: Spatiotemporal receptive fields of four different B2 neurons. Each B2 neuron has a feed-forward input (Channel 0) at reference-time $t = 0$ and four lateral inputs (Channels 1–4), each of which are stimulated with a single spike at a random time. Boxes extend from lower to upper quartiles, with lines at median presynaptic spike-times. Whiskers and flier points denote the range of data. Flier points are those past the end of the whiskers.

Table 1: DYNAP-SE circuit-IDs used to obtain the different receptive fields. The neuron IDs are local to the processor core, and the synapse IDs are local to each neuron.

Receptive field	Neuron ID ∈ [0,255]	Synapse IDs ∈ [0,63]
Figure 5a	74	0–8
Figure 5b	129	0–8
Figure 5c	222	0–8
Figure 5d	248	0–8

3.2 Feature Tuning

The B2 neuron with the receptive field of **Figure 5d** was successfully reconfigured by synapse-circuit sampling to distinguish between Pattern A and Pattern B in **Figure 4** with a spiking and non-spiking response, respectively. Prior to reconfiguration, the receptive field was fairly uniform, with the inputs on Channels 1–4 being almost interchangeable. The feature tuning was accomplished by randomly replacing the synapse circuits used for the connections, see **Table 2** and **Figure 7**. Out of 200 randomized configurations, four resulted in accurate discrimination between the two patterns. **Figure 6** shows the receptive fields of two of these successful configurations.

3.3 Energy Usage

Based on measurements of the energy dissipation for the main operations of the DYNAP-SE [31], we present an estimate of the energy

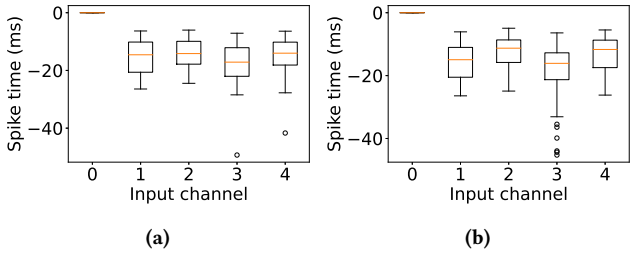

Figure 6: Receptive fields resulting from feature tuning. The receptive field of the neuron prior to tuning is presented in Figure 5d.

Table 2: DYNAP-SE circuit-IDs resulting from feature tuning. The neuron IDs are local to the processor core, and the synapse IDs are local to each neuron. The synapse IDs used for each E–I disynaptic connection are grouped inside parentheses. The original configuration of the receptive field of Figure 5d is included for reference.

Receptive field	Neuron ID ∈ [0,255]	Synapse IDs ∈ [0,63]
Figure 5d	248	0–8
Figure 6a	248	56, (2, 30), (54, 15), (46, 6), (4, 51)
Figure 6b	248	24, (25, 2), (45, 42), (28, 57), (16, 50)

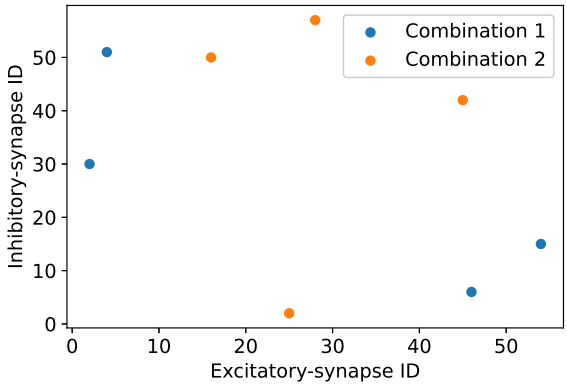

Figure 7: Synapse-circuit combinations satisfying the feature tuning of a B2 neuron. Each dot represents the hardware synapse IDs ∈ [0,63] used for the excitatory and inhibitory synapses of one of the four disynaptic elements, respectively. These synapse combinations correspond to the receptive fields presented in Figure 6.

usage per laterally projected spike-event in **Table 3** for the STC and dSTC, respectively. This estimate is made for the energy usage that is *added* by each lateral spike to the cost of the already present spike generation and routing in the networks—which is motivated by the fact that the lateral connections constitute more than half

Table 3: Estimated energy usage per lateral spike. The energy measurements for the DYNAP-SE were retrieved from [31].

SNN model	Hardware operation	Count	Energy dissipation
STC	Spike generation	1	883 pJ
	Spike-and-destination encoding	1	883 pJ
	Intracore event-broadcast	1	6.84 nJ
	Intercore event-routing	1	360 pJ
	CAM-match pulse-extension	2	324 pJ
	Sum		**9.6 nJ**
dSTC	CAM-match pulse-extension	2	324 pJ
	Sum		**0.65 nJ**

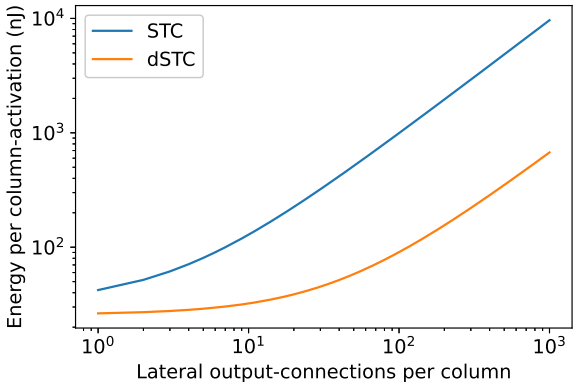

Figure 8: STC network energy-scaling with number of lateral connections. The energy is estimated in terms of the spiking activity that results from an STC column-activation—that is, for each spike of an A neuron and the subsequently elicited intercolumnar spike-events.

of connections in the network and have the potential of scaling to even higher numbers. The estimate suggests an improvement by one order of magnitude per lateral spike for the dSTC model—from 9.6 nJ to 0.65 nJ. In **Figure 8**, this energy comparison is put into the context of the energy usage of each whole network column and the energy-scaling with the number of lateral connections.

4 DISCUSSION

In the dSTC—in contrast to former implementations of STC networks—no interneuronal or axonal delays are required, since the dynamics of the postsynaptic E–I currents contribute, in effect, delayed excitations that are unique and tunable for each lateral input connection. Comparison with an STC network having dedicated delay interneurons, as in [43], shows a reduction of energy usage per lateral connection by about one order of magnitude, see **Table 3**, for implementation on the DYNAP-SE. **Figure 8** illustrates how this difference is substantial already for the number of lateral connections investigated so far, and increasingly so for potentially larger numbers of such connections. Furthermore, the original implementation of the STC used *three* synapses for each lateral, delayed B2-input—in order to mitigate device mismatch with hardware redundancy—but the energy estimate presented here is conservatively made for the ideal case of one synapse per STC connection.

The improved energy efficiency of the dSTC does, however, come at the cost of the relatively broadly tuned receptive fields of the output neurons, see **Figure 5**. This is a trade-off that could be motivated, for example, in a deep neural network of stacked STC layers with high fan-in on the B2 neurons. Moreover, the temporal widths of the receptive fields are comparable to that of a coincidence-detection based feature-detection circuit found in crickets [41], which originally inspired the disynaptic elements [40] used in the present work. During the experiments, we observed some irregularities in the results, such as widening or narrowing of the receptive fields. This variability is likely due to temperature effects [46] and may, for instance, be addressed with novel nanomaterials in future generations of neuromorphic hardware [9].

The feature-tuning experiment described in Section 2.4 provides proof-of-concept results, see Section 3.2, for hardware-circuit sampling as a possible approach for training or optimizing networks such as the dSTC. Judging by the illustration of the resulting circuit combinations in **Figure 7**, these are fairly dissimilar to each other and span a large range of the possible circuit IDs. This suggests that the solutions to the tuning problem investigated here are not unique, and, therefore, that the approach has potential for more complex tuning problems with a more constrained set of possible solutions.

5 CONCLUSION

We have investigated spike-timing-based spatiotemporal receptive fields of single mixed-signal spiking neurons using heterogeneous synaptic dynamics [2] in the DYNAP-SE neuromorphic processor and the possibility of tuning such receptive fields by discrete synapse-address reprogramming. The neurons were configured with balanced E–I synaptic dynamics [34, 40] and four lateral connections per neuron, such as the neurons in the output layer of an STC network [43, 45]. We find that the energy-cost per spike on the lateral connections is reduced by about one order of magnitude—from 9.6 nJ to 0.65 nJ—when disynaptic dynamics are used instead of dedicated delay-neurons. Furthermore, our conceptualization of the STC in terms of the receptive fields of single neurons contributes a more detailed view of its pattern-recognition mechanism than the population-level analysis of previous work [11, 43]. This may open up for further development of the network architecture, for instance by recognizing that the STC's receptive fields are similar to features in conventional deep neural networks.

ACKNOWLEDGMENTS

This work was funded by The Kempe Foundations under contract JCK-1809 and by ECSEL JU under grant no. 737 459. We thank Prof. Elisabetta Chicca for helpful feedback on an earlier version of this manuscript.

REFERENCES

[1] H. Agmon-Snir and I. Segev. 1993. Signal delay and input synchronization in passive dendritic structures. *Journal of Neurophysiology* 70, 5 (1993), 2066–2085. https://doi.org/10.1152/jn.1993.70.5.2066 PMID: 8294970.

[2] Chiara Bartolozzi and Giacomo Indiveri. 2007. Synaptic Dynamics in Analog VLSI. *Neural Computation* 19, 10 (2007), 2581–2603. https://doi.org/10.1162/neco.2007.19.10.2581 PMID: 17716003.

[3] F. C. Bauer, D. R. Muir, and G. Indiveri. 2019. Real-Time Ultra-Low Power ECG Anomaly Detection Using an Event-Driven Neuromorphic Processor. *IEEE Transactions on Biomedical Circuits and Systems* 13, 6 (2019), 1575–1582.

[4] Ben Varkey Benjamin, Nicholas A Steinmetz, Nick N Oza, Jose J Aguayo, and Kwabena Boahen. 2021. Neurogrid simulates cortical cell-types, active dendrites, and top-down attention. *Neuromorphic Computing and Engineering* 1, 1 (jul 2021), 013001. https://doi.org/10.1088/2634-4386/ac0a5a

[5] Romain Brette and Wulfram Gerstner. 2005. Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity. *Journal of Neurophysiology* 94, 5 (2005), 3637–3642. https://doi.org/10.1152/jn.00686.2005 PMID: 16014787.

[6] Julian Büchel, Fynn Firouz Faber, and Dylan Richard Muir. 2022. Network Insensitivity to Parameter Noise via Parameter Attack During Training. In *International Conference on Learning Representations*.

[7] Julian Büchel, Dmitrii Zendrikov, Sergio Solinas, Giacomo Indiveri, and Dylan R. Muir. 2021. Supervised training of spiking neural networks for robust deployment on mixed-signal neuromorphic processors. *Scientific Reports* 11, 1 (2021), 1–12. https://doi.org/10.1038/s41598-021-02779-x

[8] E. Chicca, F. Stefanini, C. Bartolozzi, and G. Indiveri. 2014. Neuromorphic Electronic Circuits for Building Autonomous Cognitive Systems. *Proc. IEEE* 102, 9 (2014), 1367–1388.

[9] Dennis V Christensen, Regina Dittmann, Bernabe Linares-Barranco, Abu Sebastian, Manuel Le Gallo, Andrea Redaelli, Stefan Slesazeck, Thomas Mikolajick, Sabina Spiga, Stephan Menzel, et al. 2022. 2022 roadmap on neuromorphic computing and engineering. *Neuromorphic Computing and Engineering* 2, 2 (may 2022), 022501. https://doi.org/10.1088/2634-4386/ac4a83

[10] Martin Coath, Robert Mill, Susan L. Denham, and Thomas Wennekers. 2011. Emergent Feature Sensitivity in a Model of the Auditory Thalamocortical System. In *From Brains to Systems*, Carlos Hernández, Ricardo Sanz, Jaime Gómez-Ramirez, Leslie S. Smith, Amir Hussain, Antonio Chella, and Igor Aleksander (Eds.). Springer New York, New York, NY, 7–17.

[11] Martin Coath, Sadique Sheik, Elisabetta Chicca, Giacomo Indiveri, Susan Denham, and Thomas Wennekers. 2014. A robust sound perception model suitable for neuromorphic implementation. *Frontiers in Neuroscience* 7 (2014), 278. https://doi.org/10.3389/fnins.2013.00278

[12] Erika Covi, Elisa Donati, Xiangpeng Liang, David Kappel, Hadi Heidari, Melika Payvand, and Wei Wang. 2021. Adaptive Extreme Edge Computing for Wearable Devices. *Frontiers in Neuroscience* 15 (2021), 429. https://doi.org/10.3389/fnins.2021.611300

[13] Thomas Dalgaty, Niccolo Castellani, Clément Turck, Kamel-Eddine Harabi, Damien Querlioz, and Elisa Vianello. 2021. In situ learning using intrinsic memristor variability via Markov chain Monte Carlo sampling. *Nature Electronics* (2021), 1–11.

[14] Anup Das, Paruthi Pradhapan, Willemijn Groenendaal, Prathyusha Adiraju, Raj Thilak Rajan, Francky Catthoor, Siebren Schaafsma, Jeffrey L. Krichmar, Nikil Dutt, and Chris Van Hoof. 2018. Unsupervised heart-rate estimation in wearables with Liquid states and a probabilistic readout. *Neural Networks* 99 (2018), 134 – 147. https://doi.org/10.1016/j.neunet.2017.12.015

[15] Simon Davidson and Steve B. Furber. 2021. Comparison of Artificial and Spiking Neural Networks on Digital Hardware. *Frontiers in Neuroscience* 15 (2021), 345. https://doi.org/10.3389/fnins.2021.651141

[16] Mike Davies, Andreas Wild, Garrick Orchard, Yulia Sandamirskaya, Gabriel A Fonseca Guerra, Prasad Joshi, Philipp Plank, and Sumedh R Risbud. 2021. Advancing neuromorphic computing with Loihi: A survey of results and outlook. *Proc. IEEE* 109, 5 (2021), 911–934.

[17] Elisa Donati, Melika Payvand, Nicoletta Risi, Renate Krause, Karla Burelo, Giacomo Indiveri, Thomas Dalgaty, and Elisa Vianello. 2018. Processing EMG signals using reservoir computing on an event-based neuromorphic system. In *2018 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. 1–4. https://doi.org/10.1109/BIOCAS.2018.8584674

[18] Charlotte Frenkel, David Bol, and Giacomo Indiveri. 2021. Bottom-Up and Top-Down Neural Processing Systems Design: Neuromorphic Intelligence as the Convergence of Natural and Artificial Intelligence. https://doi.org/10.48550/ARXIV.2106.01288

[19] J Göltz, L Kriener, A Baumbach, S Billaudelle, O Breitwieser, B Cramer, D Dold, AF Kungl, W Senn, J Schemmel, et al. 2021. Fast and energy-efficient neuromorphic deep learning with first-spike times. *Nature Machine Intelligence* 3, 9 (2021), 823–835.

[20] Wenzhe Guo, Mohammed E. Fouda, Ahmed M. Eltawil, and Khaled Nabil Salama. 2021. Neural Coding in Spiking Neural Networks: A Comparative Study for Robust Neuromorphic Systems. *Frontiers in Neuroscience* 15 (2021), 212. https://doi.org/10.3389/fnins.2021.638474

[21] Germain Haessig, Moritz B. Milde, Pau Vilimelis Aceituno, Omar Oubari, James C. Knight, André van Schaik, Ryad B. Benosman, and Giacomo Indiveri. 2020. Event-Based Computation for Touch Localization Based on Precise Spike Timing. *Frontiers in Neuroscience* 14 (2020), 420. https://doi.org/10.3389/fnins.2020.00420

[22] Giacomo Indiveri. 2021. Introducing 'Neuromorphic Computing and Engineering'. *Neuromorphic Computing and Engineering* 1, 1 (jul 2021), 010401. https://doi.org/10.1088/2634-4386/ac0a5b

[23] G. Indiveri and Y. Sandamirskaya. 2019. The Importance of Space and Time for Signal Processing in Neuromorphic Agents: The Challenge of Developing Low-Power, Autonomous Agents That Interact With the Environment. *IEEE Signal Processing Magazine* 36, 6 (2019), 16–28.

[24] Jakob Kaiser, Sebastian Billaudelle, Eric Müller, Christian Tetzlaff, Johannes Schemmel, and Sebastian Schmitt. 2021. Emulating Dendritic Computing Paradigms on Analog Neuromorphic Hardware. *Neuroscience* (2021). https://doi.org/10.1016/j.neuroscience.2021.08.013

[25] Eric Kauderer-Abrams, Andrew Gilbert, Aaron Voelker, Ben Benjamin, Terrence C Stewart, and Kwabena Boahen. 2017. A population-level approach to temperature robustness in neuromorphic systems. In *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 1–4.

[26] Juan A Leñero-Bardallo, Teresa Serrano-Gotarredona, and Bernabé Linares-Barranco. 2008. A calibration technique for very low current and compact tunable neuromorphic cells: Application to 5-bit 20-nA DACs. *IEEE Transactions on Circuits and Systems II: Express Briefs* 55, 6 (2008), 522–526.

[27] Wolfgang Maass. 1997. Networks of spiking neurons: the third generation of neural network models. *Neural networks* 10, 9 (1997), 1659–1671.

[28] Carver Mead. 2020. How we created Neuromorphic Engineering. *Nature Electronics* 3, 7 (2020), 434–435.

[29] A. Mehonic and A. J. Kenyon. 2022. Brain-inspired computing needs a master plan. *Nature* 604, 7905 (2022), 255–260. https://doi.org/10.1038/s41586-021-04362-w

[30] M. B. Milde, O. J. N. Bertrand, H. Ramachandran, M. Egelhaaf, and E. Chicca. 2018. Spiking Elementary Motion Detector in Neuromorphic Systems. *Neural Computation* 30, 9 (2018), 2384–2417. https://doi.org/10.1162/neco_a_01112 PMID: 30021082.

[31] S. Moradi, N. Qiao, F. Stefanini, and G. Indiveri. 2018. A Scalable Multicore Architecture With Heterogeneous Memory Structures for Dynamic Neuromorphic Asynchronous Processors (DYNAPs). *IEEE Transactions on Biomedical Circuits and Systems* 12, 1 (2018), 106–122.

[32] Emre Neftci and Giacomo Indiveri. 2010. A device mismatch compensation method for VLSI neural networks. In *2010 Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 262–265.

[33] C. Nielsen, N. Qiao, and G. Indiveri. 2017. A compact ultra low-power pulse delay and extension circuit for neuromorphic processors. In *2017 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. 1–4. https://doi.org/10.1109/BIOCAS.2017.8325234

[34] Mattias Nilsson, Foteini Liwicki, and Fredrik Sandin. 2020. Synaptic Integration of Spatiotemporal Features with a Dynamic Neuromorphic Processor. In *2020 International Joint Conference on Neural Networks (IJCNN)*. 1–7. https://doi.org/10.1109/IJCNN48605.2020.9207210

[35] Marcel JM Pelgrom, Aad CJ Duinmaijer, and Anton PG Welbers. 1989. Matching properties of MOS transistors. *IEEE Journal of solid-state circuits* 24, 5 (1989), 1433–1439.

[36] Nicolas Perez-Nieves, Vincent CH Leung, Pier Luigi Dragotti, and Dan FM Goodman. 2021. Neural heterogeneity promotes robust learning. *Nature communications* 12, 1 (2021), 1–9.

[37] Ning Qiao and Giacomo Indiveri. 2016. Scaling mixed-signal neuromorphic processors to 28 nm FD-SOI technologies. In *2016 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 552–555.

[38] B. Rajendran, A. Sebastian, M. Schmuker, N. Srinivasa, and E. Eleftheriou. 2019. Low-Power Neuromorphic Hardware for Signal Processing Applications: A Review of Architectural and System-Level Design Approaches. *IEEE Signal Processing Magazine* 36, 6 (2019), 97–110.

[39] Ole Richter, René Felix Reinhart, Stephen Nease, Jochen Steil, and Elisabetta Chicca. 2015. Device mismatch in a neuromorphic system implements random features for regression. In *2015 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. 1–4. https://doi.org/10.1109/BioCAS.2015.7348416

[40] Fredrik Sandin and Mattias Nilsson. 2020. Synaptic Delays for Insect-Inspired Temporal Feature Detection in Dynamic Neuromorphic Processors. *Frontiers in Neuroscience* 14 (2020), 150. https://doi.org/10.3389/fnins.2020.00150

[41] Stefan Schöneich, Konstantinos Kostarakos, and Berthold Hedwig. 2015. An auditory feature detection circuit for sound pattern recognition. *Science Advances* 1, 8 (2015), e1500325. https://doi.org/10.1126/sciadv.1500325

[42] Sadique Sheik, Elisabetta Chicca, and Giacomo Indiveri. 2012. Exploiting device mismatch in neuromorphic VLSI systems to implement axonal delays. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*. 1–6. https://doi.org/10.1109/IJCNN.2012.6252636

[43] Sadique Sheik, Martin Coath, Giacomo Indiveri, Susan Denham, Thomas Wennekers, and Elisabetta Chicca. 2012. Emergent Auditory Feature Tuning in a Real-Time Neuromorphic VLSI System. *Frontiers in Neuroscience* 6 (2012), 17. https://doi.org/10.3389/fnins.2012.00017

[44] Sadique Sheik, Michael Pfeiffer, Fabio Stefanini, and Giacomo Indiveri. 2013. Spatio-temporal Spike Pattern Classification in Neuromorphic Systems. In *Biomimetic and Biohybrid Systems*, Nathan F. Lepora, Anna Mura, Holger G. Krapp, Paul F. M. J. Verschure, and Tony J. Prescott (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 262–273.

[45] Sadique ul Ameen Sheik. 2013. *Autonomous learning in neuromorphic systems for recognition of spatio-temporal spike patterns.* Ph.D. Dissertation. ETH Zurich, Zürich. https://doi.org/10.3929/ethz-a-010118718 Diss., Eidgenössische Technische Hochschule ETH Zürich, Nr. 21425, 2013.

[46] Shihao Song, Jui Hanamshet, Adarsha Balaji, Anup Das, Jeffrey L. Krichmar, Nikil D. Dutt, Nagarajan Kandasamy, and Francky Catthoor. 2021. Dynamic Reliability Management in Neuromorphic Computing. arXiv:2105.02038 [cs.NE]

[47] Christoph Stöckl and Wolfgang Maass. 2021. Optimized spiking neurons can classify images with high accuracy through temporal coding with two spikes. *Nature Machine Intelligence* 3, 3 (2021), 230–238.

[48] SynSense. 2021. *Samna.* SynSense AG. https://synsense-sys-int.gitlab.io/samna/

[49] Neil C. Thompson, Kristjan Greenewald, Keeheon Lee, and Gabriel F. Manso. 2021. Deep Learning's Diminishing Returns: The Cost of Improvement is Becoming Unsustainable. *IEEE Spectrum* 58, 10 (2021), 50–55. https://doi.org/10.1109/MSPEC.2021.9563954

[50] Simon Thorpe, Arnaud Delorme, and Rufin Van Rullen. 2001. Spike-based strategies for rapid processing. *Neural networks* 14, 6-7 (2001), 715–725.

[51] Hans Tuinhout, Nicole Wils, and Pietro Andricciola. 2010. Parametric mismatch characterization for mixed-signal technologies. *IEEE journal of solid-state circuits* 45, 9 (2010), 1687–1696.

[52] Yingxue Wang and Shih-Chii Liu. 2012. Active processing of spatio-temporal input patterns in silicon dendrites. *IEEE transactions on biomedical circuits and systems* 7, 3 (2012), 307–318.

[53] Le Ye, Zhixuan Wang, Ying Liu, Peiyu Chen, Heyi Li, Hao Zhang, Meng Wu, Wei He, Linxiao Shen, Yihan Zhang, et al. 2021. The Challenges and Emerging Technologies for Low-Power Artificial Intelligence IoT Systems. *IEEE Transactions on Circuits and Systems I: Regular Papers* 68, 12 (2021), 4821–4834. https://doi.org/10.1109/TCSI.2021.3095622

[54] Fleur Zeldenrust, Boris Gutkin, and Sophie Denéve. 2021. Efficient and robust coding in heterogeneous recurrent networks. *PLoS computational biology* 17, 4 (2021), e1008673.

[55] Zhi Zhou, Xu Chen, En Li, Liekang Zeng, Ke Luo, and Junshan Zhang. 2019. Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proc. IEEE* 107, 8 (2019), 1738–1762.