

A Neuromorphic Approach to Path Integration: a Head-Direction Spiking Neural Network with Vision-driven Reset

Raphaela Kreiser^{*†}, Matteo Cartiglia^{*†}, Julien Martel^{*}, Jörg Conradt[†], and Yulia Sandamirskaya^{*}

^{*}Institute of Neuroinformatics, University of Zurich and ETH Zurich, Winterthurerstr. 190, 8057 Zurich, Switzerland

[†]Neuroscientific Systems Theory, Technical University of Munich, Karlstrasse 45-47, 80333 Munich, Germany

[‡] These authors have contributed equally to this work.

Email: {rakrei, ysandamirskaya}@ini.uzh.ch

Abstract—Simultaneous localization and mapping (SLAM) is one of the core tasks of mobile autonomous robots. Looking for power-efficient and embedded solutions for SLAM is an important challenge when building controllers for small and agile robots. Biological neural systems of even simple animals are until now unprecedented in their ability to localize themselves in an unknown environment. Neuromorphic engineering offers ultra low-power and compact computing hardware, in which biologically inspired neuronal architectures for SLAM can be realised. In this paper, we propose a low-power on chip approach for one of the components of SLAM – path integration. Our solution takes inspiration from biology and uses motor command information to estimate the orientation of an agent solely in a spiking neural network. We realise this network on a neuromorphic device that implements artificial neurons and synapses with analog electronics. A robotic agent receives visual input from an event-based camera and uses this information to correct the on-chip spiking neurons’ estimate of orientation. This system can be easily integrated with other localization and mapping components on chip and is a step towards a fully neuromorphic SLAM.

I. INTRODUCTION

Foraging in unknown, dynamically changing environments is a crucial skill for the survival of every animal. One of the tasks of foraging animals, as well as mobile robots, is to use self-motion cues in order to compute the direction and distance from the navigator’s current position to some reference point, e.g. the nest. This process, called path integration, is prone to accumulated error. In robots, this error is due to wheel slippage, uneven terrain, and fabrication mismatch. In humans and animals, drift in the estimated position occurs when navigating in the dark. In order to maintain an accurate position estimate, path integration can be corrected based on visual or other external cues.

Biological neural networks of even simple animals can perform the task of path integration effortlessly. Drawing inspiration from these networks could therefore lead to particularly efficient solutions in robotic navigation. A biological neural network that plays a key role in spatial navigation is composed of orientation selective head direction (HD) neurons. HD neurons have been found in many animals and have been referred to as a neural compass, as they spike in relation to the orientation of the head with respect to the outside environment. Contrary to a magnetic compass that can directly measure the orientation, neural circuits in the animal’s brain *compute* the orientation based on internal motion cues and external sensory input from the environment. Simulations of these neuronal circuits have been conducted [1], [2], however, emulating a neuronal architecture with its dynamics in software

is computationally expensive and is therefore not suitable for a real-time robotic application.

Neuromorphic hardware, to the contrary, offers a physical computational substrate for directly emulating the biophysics of neurons and synapses in real time [3]–[6], enabling low latency through massively parallel, event-based computation. In robotics, analog neuromorphic hardware has proven itself capable to achieve better time complexity and power savings in path planning [7].

In this work, we present a neural head direction architecture implemented on a neuromorphic device and tested with a robotic agent. We use the neuromorphic processor Reconfigurable On-Line Learning System (ROLLS) [8] that comprises hybrid analog/digital circuits to emulate the biological processes of spiking neurons and synapses with ultra low-power. Similarly to biological neural networks that face the problem of individual neurons and synapses being driven by stochastic processes, neuromorphic silicon neurons are noisy due to the variable nature of analog circuits and exhibit mismatch due to fabrication imperfections.

Inspired by existing models of HD cells [1], [2], we implement a HD network on a neuromorphic device. In our network, the recurrent connectivity in a population of HD neurons gives rise to a localized activity “bump”. Together with an angular velocity cue this activity induces a shifted localised “bump” in a second population of neurons, which feeds back into the HD population, making the activity bump in this latter population move in the correct direction. The implemented neural architecture demonstrates that the rotation of the robot shifts the HD activity according to the robot’s speed. In order to correct the accumulated error that is due to device mismatch, as well as to odometric drift, a visual cue can reset the activity of the HD cells to the correct location.

By functioning in a robotic sensory loop, the proposed neural architecture provides a sensory-motor embodiment for existing models of biological HD cell networks. A neural network implementation on neuromorphic hardware reduces dramatically the power consumption for embedded systems that require real-time processing.

We proceed with a brief introduction to the ROLLs neuromorphic processor and the robotic agent that was used for this work. We then introduce the proposed HD network and the visual reset mechanism. Finally, we validate the results of path integration using the IMU data as a ground truth.

II. MATERIALS AND METHODS

The setup used in this work consists of the ROLLS neuromorphic processor and the Pushbot robotic vehicle¹ with an embedded event based DVS camera (eDVS) [9]. Communication between the Pushbot and the ROLLS is coordinated by a Parallella board [10]: the Pushbot communicates through a dedicated wireless connection to the Parallella board receiving motor commands and sending address-events produced by the eDVS. The ROLLS chip is interfaced with the Parallella with an embedded FPGA which is used to configure the network connectivity and to stimulate neurons.

A. Neuromorphic processor




Fig. 1: Block-diagram of the chip architecture.

The ROLLS neuromorphic processor [8] is a full-custom mixed signal analog/digital VLSI device. It comprises analog neuromorphic circuits that emulate the biophysics of spiking neurons and dynamic synapses in real-time and asynchronous digital circuits that manage the transmission and routing of spikes, using the Address Event Representation (AER) [11]. The chip comprises a column of 256 neurons, an array of 256x256 non-plastic programmable synapses, an array of 256x256 plastic synapses, and 8x256 time-multiplexed “virtual” synapses that can be used to provide direct external input to neurons.

A block-diagram of the chip architecture is shown in Fig. 1. Peripheral input-output AER circuits for receiving and transmitting off-chip spikes in real time can be used to stimulate individual synapses or neurons on the chip. Silicon synapses process spikes as they arrive, and produce output currents with biologically plausible temporal dynamics. Silicon neurons integrate these currents to generate and transmit spikes in real time. The on-chip programmable bias generator [12] allows programming the properties of the synapses and neurons (time constants, pulse widths, etc.). The neuron circuit integrated in the ROLLS chip exhibits biologically realistic neuronal behaviors, such as spike-frequency adaptation, adjustable refractory period, and spiking threshold [13]. The neuron equations derived from the circuit closely resemble those of the adaptive exponential integrate-and-fire neuron model [14].

B. Neuromorphic robot and eDVS

The robot used in this work is a mobile autonomous platform called Pushbot, which consists of a 10x10 cm chassis with two motors driving two independent tracks for propulsion




Fig. 2: The pushbot with the eDVS (right), the ROLLS (left), and the visual beacon (center).

(left and right). It comprises an embedded dynamic vision sensor (DVS) – an event-based camera, inspired by the mammalian retina [15], [16]. Each pixel of the DVS is sensitive to the temporal change in luminance and sends out an event using the AER protocol. The robot has an integrated 9 DOF IMU which reports changes in velocity and orientation, which is used here to obtain the ground truth of robot’s orientation. The eDVS camera is configured to detect a high frequency LED, filtering out the noise that arises from stochastic spiking and movement of the robot. A high frequency LED (3kHz) is used as a visual landmark in order to reset the bump of activity in the HD neural network. In our implementation, when the visual landmark is observed within a set area of the visual array, the bump of activity is reset to the correct value.

III. HEAD DIRECTION NETWORK

In the past 20 years, HD cells have been characterized biologically [17] and modeled extensively. An individual HD cell has a maximum firing rate at one particular orientation and its firing rate decreases monotonically moving away from this ‘preferred direction’. In a population of HD cells with all directions equally represented, head direction is represented by a stable and localized activity bump [18]. Despite the strong visual component, HD cells also use inertia to update the representation of head direction in total darkness, presumably by integrating self motion information.

Many models have been proposed on how HD cells shift their activity according to angular movements. In the ring attractor model for HD cells, each cell features local excitation onto cells with similar orientation preference and global inhibition onto cells with different orientation preference [1]. Direct evidence for such an organization was recently found in insects [19]. The ‘hill of activity’ is moved around the ring to different directional headings following inputs from self-motion or external visual cues. An alternative model relies on cross-inhibition [2] and is based on data from angular head velocity cells that have been found in the thalamus. In this model, neural activity is mainly driven by external excitation making the network less prone to instability that might occur due to runaway feedback excitation [20]. Wilson and Cowan [21], and Amari [22] have proposed dynamic neural fields (DNFs) as a simplified mathematical model for neural processing based

¹<https://inilabs.com/products/pushbot/>




Fig. 3: The neural network architecture with the connections between five neural populations: head direction (HD), shift right (SR), integrated head direction (IHD), and angular velocity populations drive left (DL) and drive right (DR). The shift left (SL) population is not shown for clarity. The IMU and DVS populations reset the activity bump in the IHD population, driven by the compass- and DVS-signal, respectively. Arrows mark excitatory, circles inhibitory synapses (only a subset of recurrent connections is shown).




Fig. 4: The synaptic connectivity matrix that implements the architecture on the ROLLS chip. Colors encode synaptic weights of the recurrent connections on the chip (red for positive and blue for negative weights). Labels indicate the neural populations (see Fig. 3)

on soft winner take all (WTA) recurrent interactions, which have later been linked to representations in the HD system of rats [18]. Attractor models based on DNFs can account for various attributes of HD cells, e.g. neurons firing at a steady rate when facing the preferred direction even in darkness and shifting their activity according to self-motion.

A. Spiking Neural Network Architecture

The core of the presented system is a neuromorphic neural network architecture that integrates the robot’s angular velocity to compute its orientation. The architecture consists of six neuronal populations: An angular velocity population (AV) that is composed of a drive right (DR) and a drive left (DL) sub-populations, the head-direction neurons (HD), the integrated

head direction neurons (IHD), and two populations that shift the activity to the right (SR) or to the left (SL), respectively. A schematic representation of the networks is shown in Fig. 3. All populations are connected as ring networks. The figure does not depict the closure of the ring in order to obtain a better illustration of the synaptic connectivity.

HD neurons are connected in a WTA network, in which neurons amplify their local activity by being excitatorily connected to their nearest neighbors and globally inhibit each other. As a consequence, the network “selects” the local group that receives the strongest excitation and at the same time suppresses the activity of all other neurons via global inhibition, thus creating a bump of activity on the ring attractor network. Each active HD neuron inhibits every shift neuron except for the ones with the same index and the ones to the immediate right and left, allowing for local activity in the shift network when it is globally stimulated due to movement (by the AV populations). The two distinct shift populations (SL and SR) connect to the IHD population with asymmetric network connectivity [1], [18], [23], shifting the bump to the right (SR) and to the left (SL), respectively. With its one-to-one mapping to the HD neurons, the IHD population moves the HD activity bump to the estimated updated orientation.

When the robot turns, an angular velocity population (DR or DL) excites the whole shift population, but will only induce activity in the local region of the shift population due to the inhibition from the HD population. Shift neurons excite the IHD neurons to the immediate right or left and shift the bump of activity relative to the HD population. Finally, the IHD strongly feeds into the HD population and moves the activity bump smoothly with integration or abruptly by external reset. IMU is a neuronal population that, when activated, resets the IHD to the true position calculated with the IMU compass on the Pushbot. Moreover, when a visual landmark is detected, the DVS population is stimulated by the DVS events and moves the IHD activity bump to the correct position. The visual landmark corresponds to a blinking beacon and is seen once in every full rotation. Cross-inhibition between DL-SR and DR-SL populations reduces the chance of stochastic spiking. The excitatory WTA kernel of width two was chosen to make the localized activity more robust and stable. The synaptic connectivity matrix set on the neuromorphic chip is shown in Fig. 4.

IV. RESULTS

In our robotic experiments, we put the robot in an arena with a blinking LED fixed on one of the walls (see Fig. 2). In the beginning of each trial, the IMU state was read to obtain the true heading orientation. The robot turned around its axis and at the end of the rotation the IMU state was read again. The difference between the orientation represented by the activity of the HD neurons and the true orientation, given by the IMU data was calculated. This difference represents the error of the system and is referred to as Δ , measured in number of neurons. The spiking activity during two example trials can be seen in Fig. 5. It can be observed that the path integration with visual reset (right plot) is more precise due to the correction of the accumulated error.

The reliability of the system was tested by analyzing the difference between the HD neurons activated through




Fig. 5: Spiking activity of neurons on chip using the DVS-based reset (**right**) and without visual information (**left**). The figure shows how the visual correction aids the head direction network to have a more precise orientation estimation, limiting error accumulation (see “Accumulated error” in the **left** plot). IMU checks in the beginning and end of the trial are marked by red spikes of the IMU population. DVS population is activated when the visual landmark is detected (≈ 13 s and 26s).




Fig. 6: The mean of the accumulated error over time with and without visual reset. Error bars represent the one sigma spread of error measurements over dozens of trials. Data-points without visual feedback were shifted by one (to the right) to avoid graphic overlaps.

integration and the neurons that were activated by the IMU reset. Fig. 6 shows the mean error for the system with and without visual correction, measured as the robot turns for up to 45 seconds. We have conducted many runs of the system (between 16 and 52 per data point in the plot) with rotations both to the right (e.g., Fig. 6b) and to the left (e.g., Fig. 6a). This also allowed us to estimate the standard deviation of the error over time and revealed that it grows for trials without visual reset.

Fig. 6 shows that the mean deviation of the neurally estimated orientation does not diverge both with and without visual reset, but fluctuates around zero. This means that the neuronally estimated shift corresponded to the angular velocity of the robot and there was no constantly accumulating error due to bad calibration. However, using visual reset makes the system more reliable and less prone to neuronal or mismatch induced drift. This leads to a smaller mean error and a narrower standard deviation. By performing a χ^2 test, we assessed that

the system with visual reset is compatible with having an offset of zero as time progresses.

V. CONCLUSION

In this work, we realized a neural head direction architecture on a neuromorphic device with spiking silicon neurons that can path-integrate angular velocity information in order to maintain an internal representation of orientation. The activity of the silicon exponential integrate and fire neurons is driven in a closed sensory-motor loop. By detecting a visual landmark, the accumulated error can be corrected, as demonstrated in Fig. 5 and Fig. 6.

In future work, first, we will use multiple AV (drive) populations for different angular velocities, leading to a number of shift populations that produce shifts of different amplitude. Next, the visual information can be used to correct the path integration system by making the weights between shift and integrated head direction populations plastic and adjusting the activity shift to better correspond to the actual angular velocity. Finally, blinking LEDs with different frequencies can be placed at various orientations and during the rotation of the robot, on-chip plastic synapses of the silicon neurons can learn the corresponding place of each LED, effectively forming a map of the robot’s surroundings. The LEDs can be replaced by an object-recognition system.

Combining path integration with the visual correction is the first step towards robotic localization implemented in neuromorphic hardware. Future work will use a larger neuromorphic device to add path integration for translation and to establish a map of the environment using plastic on-chip connections, leading towards a complete neuromorphic simultaneous localization and mapping (SLAM) system.

ACKNOWLEDGMENT

The authors would like to thank Giacomo Indiveri and Ning Qiao for their help and support throughout the project. This project was funded by the UZH grant FK-16-106, ZNZ fellowship and SNF grant PZ00P2 168183.

REFERENCES

- [1] X. Xie, R. H. R. Hahnloser, and H. S. Seung, "Double-ring network model of the head-direction system," *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.*, vol. 66, no. 4, 2002.
- [2] P. Song, "Angular Path Integration by Moving "Hill of Activity": A Spiking Neuron Model without Recurrent Excitation of the Head-Direction System," *J. Neurosci.*, vol. 25, no. 4, pp. 1002–1014, 2005. [Online]. Available: <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.4172-04.2005>
- [3] G. Indiveri, E. Chicca, and R. J. Douglas, "Artificial cognitive systems: From VLSI networks of spiking neurons to neuromorphic cognition," *Cognit. Comput.*, vol. 1, no. 2, pp. 119–127, 2009.
- [4] S. B. Furber, D. R. Lester, L. A. Plana, J. D. Garside, E. Painkras, S. Temple, and A. D. Brown, "Overview of the SpiNNaker system architecture," pp. 2454–2467, 2013.
- [5] B. V. Benjamin, P. Gao, E. McQuinn, S. Choudhary, A. R. Chandrasekaran, J. M. Bussat, R. Alvarez-Icaza, J. V. Arthur, P. A. Merolla, and K. Boahen, "Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations," *Proc. IEEE*, vol. 102, no. 5, pp. 699–716, 2014.
- [6] E. Chicca, F. Stefanini, C. Bartolozzi, and G. Indiveri, "Neuromorphic electronic circuits for building autonomous cognitive systems," *Proc. IEEE*, vol. 102, no. 9, pp. 1367–1388, 2014.
- [7] S. Koziol, S. Brink, and J. Hasler, "A neuromorphic approach to path planning using a reconfigurable neuron array IC," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 22, no. 12, pp. 2724–2737, 2014.
- [8] N. Qiao, H. Mostafa, F. Corradi, M. Osswald, D. Sumislawska, G. Indiveri, and G. Indiveri, "A Re-configurable On-line Learning Spiking Neuromorphic Processor comprising 256 neurons and 128K synapses," *Frontiers in neuroscience*, vol. 9, no. February, 2015.
- [9] J. Conradt, R. Berner, M. Cook, and T. Delbruck, "An embedded AER dynamic vision sensor for low-latency pole balancing," 2009 *IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pp. 780–785, 2009. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.192.1328{\&}rep=rep1{\&}type=pdf>
- [10] A. Olofsson, T. Nordström, and Z. Ul-Abdin, "Kickstarting high-performance energy-efficient manycore architectures with Epiphany," *Conference Record - Asilomar Conference on Signals, Systems and Computers*, vol. 2015-April, no. May, pp. 1719–1726, 2015.
- [11] K. A. Boahen, "Point-to-Point Connectivity Between Neuromorphic Chips using Address-Events," *Ieee Transactions on Circuits & Systems*, vol. 47, no. 5, pp. 416–434, 1999. [Online]. Available: https://web.stanford.edu/group/brainsinsilicon/pdf/00{_}journ{_}IEEEtsc{_}Point.pdf
- [12] T. Delbruck, R. Berner, P. Lichtsteiner, and C. Dualibe, "32-bit configurable bias current generator with sub-off-current capability," in *ISCAS 2010 - 2010 IEEE Int. Symp. Circuits Syst. Nano-Bio Circuit Fabr. Syst.*, 2010, pp. 1647–1650.
- [13] G. Indiveri, E. Chicca, and R. Douglas, "A VLSI Array of Low-Power Spiking Neurons and Bistable Synapses With Spike-Timing Dependent Plasticity," vol. 17, no. 1, pp. 211–221, 2006.
- [14] R. Brette and W. Gerstner, "Adaptive exponential integrate-and-fire model as an effective description of neuronal activity," *J. Neurophysiol.*, vol. 94, no. 5, pp. 3637–3642, 2005. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/16014787>
- [15] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128 X 128 120db 30mw asynchronous vision sensor that responds to relative intensity change," 2006 *IEEE Int. Solid State Circuits Conf. - Dig. Tech. Pap.*, pp. 2004–2006, 2006.
- [16] S. C. Liu and T. Delbruck, "Neuromorphic sensory systems," pp. 288–295, 2010.
- [17] J. S. Taube, R. U. Muller, and J. B. Ranck, "Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis," *J. Neurosci.*, vol. 10, no. 2, pp. 420–35, 1990. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/2303851>
- [18] K. Zhang, "Representation of Spatial Orientation by the Intrinsic Dynamics of the Head-Direction Cell Ensemble: A Theory," *J. Neurosci.*, vol. 16, no. 6, pp. 2112–2126, 1996.
- [19] J. D. Seelig and V. Jayaraman, "Neural dynamics for landmark orientation and angular path integration," *Nature*, vol. 521, no. 7551, pp. 186–191, 2015. [Online]. Available: <http://www.nature.com/doi/10.1038/nature14446>
- [20] X. J. Wang, "Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory," *J. Neurosci.*, vol. 19, no. 21, pp. 9587–9603, 1999.
- [21] H. R. Wilson and J. D. Cowan, "A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue," *Kybernetik*, vol. 13, no. 2, pp. 55–80, 1973.
- [22] S. ichi Amari, "Dynamics of pattern formation in lateral-inhibition type neural fields," *Biol. Cybern.*, vol. 27, no. 2, pp. 77–87, 1977.
- [23] A. Redish, A. Elga, and D. Touretzky, "A coupled attractor model of the rodent head direction system," *Netw. Comput. Neural Syst.*, vol. 7, no. 4, pp. 671–685, 1996.