

Even Faster SNN Simulation with Lazy+Event-driven Plasticity and Shared Atomics

1st Dennis Bautembach
FORTH-ICS & CSD-UOC
denniskb@ics.forth.gr

2nd Iason Oikonomidis
FORTH-ICS
oikonom@ics.forth.gr

3rd Antonis Argyros
FORTH-ICS & CSD-UOC
argyros@ics.forth.gr

Abstract—We present two novel optimizations that accelerate clock-based spiking neural network (SNN) simulators. The first one targets spike timing dependent plasticity (STDP). It combines lazy- with event-driven plasticity and efficiently facilitates the computation of pre- and post-synaptic spikes using bitfields and integer intrinsics. It offers higher bandwidth than event-driven plasticity alone and achieves a $1.5\times-2\times$ speedup over our closest competitor. The second optimization targets spike delivery. We partition our graph representation in a way that bounds the number of neurons that need be updated at any given time which allows us to perform said update in shared memory instead of global memory. This is $2\times-2.5\times$ faster than our closest competitor. Both optimizations represent the final evolutionary stages of years of iteration on STDP and spike delivery inside “Spice” (/spaik/), our state of the art SNN simulator. The proposed optimizations are not exclusive to our graph representation or pipeline but are applicable to a multitude of simulator designs. We evaluate our performance on three well-established models and compare ourselves against three other state of the art simulators.

Index Terms—AI, ML, spiking neural networks, SNN, simulation, HPC, GPGPU, CUDA

I. INTRODUCTION

SNNs are receiving a lot of attention since they more closely model biological neural networks than conventional ANNs do and therefore may prove more powerful. In contrast to ANNs, which can be inferred via matrix multiplications interleaved with activation functions, SNNs need to be *simulated* because their neurons can undergo *arbitrary* dynamics.

A SNN simulator works fairly simple:

- 1) Update neurons, note which ones fire
- 2) Update synapses
 - Compute pre- and post-synaptic spikes
- 3) Deliver spikes to neurons’ neighbors

Neurons are exclusively governed by their dynamics and can thus be updated independently of one another, making it trivial to parallelize step (1) on GPUs. The synapse update requires more attention: In order to advance synapses’ dynamics, a SNN simulator must first compute pre- and post-synaptic spikes which is a non-trivial task. Additionally, the sheer number of synapses (in the order of $O(\#neurons^2)$) means that most of the simulation time will be spent in step (2), requiring careful implementation that maps well to the hardware. Step (3) is another non-trivial and expensive part of the simulation loop since again the number of spikes

is in the order of $O(\#neurons^2)$ and neurons are randomly interconnected, leading to poor memory access patterns.

In this paper we present two optimizations targeting steps (2) and (3) of the SNN simulation loop. In order to accelerate step (2) we, for the first time, combine lazy plasticity [1] with event-driven plasticity [2]–[6] and efficiently facilitate the computation of pre- and post-synaptic spikes via per-neuron bit fields and integer intrinsics. Combined, these measures yield a cumulative speedup of $30\times$ over the naïve approach of blindly updating each synapse at every loop iteration and are $1.5\times-2\times$ faster than our closest competitor.

In order to accelerate step (3) we partition our adjacency list according to [7, Section III-D] into slices of equal width *in the neuron domain*. Knowing that each slice indexes a fixed, contiguous subset of the neuron pool we are able to load said subset into shared memory and perform all updates there. The resulting spike delivery algorithm is $2\times-2.5\times$ faster than relying on global memory and is *at least* $2\times$ faster than our closest competitor (depending on the benchmark).

These optimizations do not come at the cost of generality—users can still define custom models with arbitrary dynamics. Further, they are completely orthogonal to our multi-GPU acceleration scheme [7]. Finally, they are not exclusive to our graph representation or pipeline but are applicable to a multitude of simulator designs.

II. RELATED WORK

Tuckwell [8] notes that early attempts to mathematically model and simulate the function of biological neurons are found since the beginning of the 20th century [9], [10]. The research domain related with the efficient simulation of SNNs has a long history [2], [11], [12] and is actively research to this day [1], [13]–[17]. Within the domain, important points of research focus include the improvement of biological fidelity [18] and numerical stability of methods [17]. Furthermore, hardware acceleration is explored [19], [20], including platforms such as VLSI [11], FPGA [20], and even super-computers [21]. On the theoretical front, advances include methods to quantify the difference between two spike trains [22], exact solutions of differential equations that model membrane dynamics [23], as well as a method to implement back propagation on SNNs [24].

A. SNN Simulator Classification

The various SNN simulators that have been proposed feature different key characteristics, strengths and weaknesses. Based on these traits, one can define useful classifications. The biological fidelity level can be such a defining trait. According to it, simulators can be categorized into ones that aim to model the behavior of biological neurons as accurately as possible [18], [25]–[30], and into those that are merely based on the general principle of spiking neurons.

The versatility, and specifically the ability and ease of defining neuron models can serve as another way to categorize simulators. General purpose simulators allow the user to freely define the behavior of neurons and synapses [1], [5], [7], [31]–[35], whereas the rest only allow the use of a fixed number of predefined models [6], [18], [26], [30].

Another important classification can be made based on the approach adopted regarding the high-level software implementation. In one category, the time-drive approach simulates the behavior of neurons and synapses in lock-step, advancing all of them in every iteration of the simulation [1], [6], [7], [23], [33], [36], [37]. In contrast, event-driven simulators only update a neuron or synapse when a new event affects the (predictable, up until that time instant) evolution of their state [1], [6], [23], [33], [36], [37]. Ease of implementation and per-element efficiency are two strong points of the time-driven approaches, however if the time-step is set very low (desired to increase simulation accuracy), the simulation becomes computationally inefficient. Event-driven approaches on the other hand can achieve arbitrary time resolution with small or even negligible computational overhead, despite having low per-element computational efficiency. As a middle ground, approaches that adopt characteristics of both categories can be called hybrid [38], [39]. It is debatable whether approaches that adopt a time-driven neuron update with an event-driven synapse update (event-driven plasticity) [2]–[6] should be called hybrid, since this combination is less integral.

One last categorization can be made according to the target hardware platform of the simulator system. A wide variety of hardware has been employed, starting from regular CPUs [12], [15], [32], [36], [40]–[42], GPUs [1], [5]–[7], [13], [14], [16], [20], [27], [30], [33]–[35], [43], [44], combination of CPUs and GPUs [19], [38], [39], [45], and including even custom-built hardware [11], [20], [21], [25], [26], [46]. Recent simulators are increasingly adopting the GPU architecture.

B. Classification of Spice and Direct Competition

According to the categorizations presented above, Spice [1], [7] is general-purpose, time-driven, GPU-accelerated, and in this work it is fitted with event-driven plasticity.

For the quantitative evaluation of the proposed optimizations in our simulator, we choose as competitors three very recent works, closely related to ours: BSim [16], GeNN [5], and NeuronGPU [47]. All three simulators are time-driven, GPU-accelerated, and, to varying degrees, general. This allows the implementation of the same SNN models on all of them, in turn enabling a direct comparison between them and our

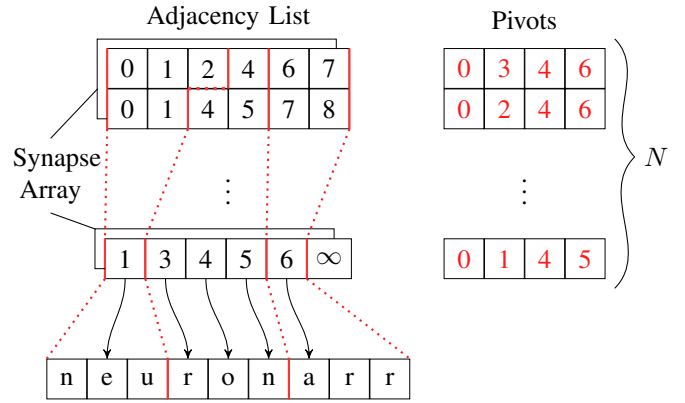


Fig. 1. Overview over our data structures. We represent the graph using a padded/rectangular adjacency list. It has $N = |\text{neurons}|$ many rows. Row i represents neuron i 's outgoing connections in the form of indices into the neuron array. They are sorted. Non-connections are represented by a special value (`INT_MAX` in practice). The synapse array has the same number of entries as the adjacency list with an implicit 1:1 mapping between the two (`adj[i][j]` corresponds to `synapses[i][j]`). We divide the neuron array into chunks of even size, in this case 3. We then binary-search pivots so that entries `adj[i][pivots[i][j] . pivots[i][j+1]-1]` only index neurons $[3j, 3j+3)$. We explicitly store the implicit “0” so as to avoid special cases in the implementation.

approach. We do not include event-driven simulators in the comparison since they pursue different goals with fundamentally different designs and imposed trade-offs, and so are not directly comparable. Even within the time-driven domain, the target of each simulator, and the associated performance can vary quite a lot. GeNN [5], [33], [48] has improved quite a bit since its introduction, and is usually among the fastest direct competitors. BSim [16] among other features supports multi-GPU setups. Finally, NeuronGPU [47] is the least comparable of the three simulators. It strives to maximize biological fidelity and as such features double precision arithmetic and “exact integration” [49]. We still included it in this comparison, mostly for completeness, to gauge what performance trade-off these design choices entail.

III. METHOD

SNNs are essentially directed graphs (neurons = vertices, synapses = edges). They are typically represented as a (binary) adjacency matrix, adjacency list, or compressed sparse row (CSR). All representations are grouped by source neuron (that is, a neuron’s outgoing connections are stored in contiguous memory) because it is suitable for computations. The connections are also sorted to improve cache coherency. Since a SNN’s topology is static, the graph needs to be constructed only once at the beginning of a simulation.

We will illustrate the optimization techniques described in this paper using a (padded/rectangular) adjacency list because it is a very simple representation and also happens to be what Spice uses (Fig. 1). The presented optimizations are applicable to a variety of representations though, including all of the commonly used ones mentioned above.

```
extern int delay; // synaptic delay (in time steps) of the network
struct neuron { int64 hist; ... }; // firing history, each bit represents one time step, LSB = most recent
struct synapse { neuron& src; neuron& dst; ... };
void update(synapse& syn, bool preSynapticSpike, bool postSynapticSpike, int numSteps = 1); // user-defined callback
```

<pre>1 parfor n in neurons: 2 parfor syn in outgoingSynapses(n): 3 4 5 6 update(7 syn, 8 n.hist[delay], 9 syn.dst.hist[0]); 10 11 12</pre> <p style="text-align: center;">(a)</p>	<pre>parfor n in firing neurons: parfor syn in outgoingSynapses(n): age = now - n.timeOfLastUpdate; for i in age..0: update(syn, i == 0, syn.dst.hist[i]);</pre> <p style="text-align: center;">(b)</p>	<pre>parfor n in firing neurons: parfor syn in outgoingSynapses(n): age = now - n.timeOfLastUpdate; j = -1; for i in setBits(syn.dst.hist[age..0]): update(syn, i == age, true, i - j); j = i; update(syn, j < age, false, age - j);</pre> <p style="text-align: center;">(c)</p>
--	---	---

Fig. 2. Pseudocode showing the evolution of our implementation from naïve, over lazy, to event-driven plasticity. **(a) Naïve plasticity.** We loop over all synapses (lines 1–2) and invoke the user-provided callback `update()` (line 6) which is part of the model definition. We pass it a reference to each synapse (line 7) as well as information about the occurrence of pre- and post-synaptic spikes by reading the corresponding bits from the source and destination neuron’s firing histories (lines 8–9, $x[i] := (x \gg i) \& 1$). The `parfor` loops can be trivially parallelized because each loop iteration is independent. In the case of two nested `parfors` the outer loop iterations could be assigned to different CUDA blocks and the inner loop iterations could be assigned to different threads within a block. **(b) Lazy plasticity.** Instead of updating all synapses, we only update those originating from firing neurons (line 1) since they are about to transmit a spike and thus need be up to date. We compute the synapses’ “age”, i.e. the number of iterations for which we have intentionally neglected them (line 3) and “replay” all the missed updates (line 5). **(c) Event-driven plasticity** by contrast does not replay the entire firing history since the last spike, but only iterates over set bits, i.e. post-synaptic spikes, skipping ahead multiple update steps in-between (line 5). This is communicated to the user via `update()`’s fourth parameter which represents the number of time steps that have elapsed since the callback was last invoked (line 10). The `update()` call in line 12 handles the “tail”.

A. Lazy+Event-driven Plasticity

“Plasticity” refers to models with changing (i.e. plastic) synapse state—neurons are always plastic. The sheer number of synapses makes this a very costly operation worth optimizing. The naïve approach (Fig. 2(a)) would be to simply update every synapse at each simulation step. This is incredibly slow with the limiting factor being memory bandwidth: Brunel+ uses 12 B per plastic synapse which make up 40% of all synapses. In order to run a network with 1B synapses total in real time, it would require 96 TB/s of bandwidth (every plastic synapse would have to be read and written 10K times per second of biological time). In comparison, a NVIDIA Tesla V100 offers ~ 0.7 TB/s of bandwidth.

We already presented a simple yet effective optimization to this naïve approach called “Lazy Plasticity” [1, Section III-A1] (Fig. 2(b)). It takes advantage of the facts that

- neurons fire infrequently and
- a synapse need only be up to date to transmit a spike

Rather than updating all synapses at every step we intentionally keep them in a stale state. When a neuron fires, we compute its $age = now - timeOfLastSpike$ and perform that many update steps on its outgoing synapses. While the total number of computations stays the same, they can now be performed inside registers with only the initial and final state

having to be read from/written to global memory, increasing our effective bandwidth (Section IV-C).

Plasticity need not only update synapses’ states but also facilitate the efficient computation of pre- and post-synaptic spikes. A pre-synaptic spike occurs if a synapse’s source neuron fires, a post-synaptic spike occurs if a synapse’s destination neuron fires. We must inform the synapse of both types of events. Pre-synaptic spikes are trivial to compute, especially if synapses are grouped by source neuron: The fact that a neuron fires, implies that all its outgoing synapses experience a pre-synaptic spike (potentially after some delay). Luckily for us these same synapses are already laid out in a known, contiguous region of memory, making it very efficient to iterate over and update them.

Post-synaptic spikes are more tricky to compute since we need to consolidate a neuron’s *incoming* synapses. An infeasibly slow solution would be to iterate over the *entire* adjacency list and search for all occurrences of said neuron, and repeat this for all firing neurons. Some simulators [5], [6] tackle this issue by storing a “reverse” adjacency list. That is, in addition to storing outgoing synapses they also store incoming synapses per neuron. This has several disadvantages:

- It doubles the memory footprint of the graph representation.

- While the adjacency data of incoming synapses are now laid out in contiguous memory and can thus be efficiently iterated, the synapses they index are spread all over memory resulting in very poor bandwidth when updating them.

We instead store the network’s recent firing history in the form of per-neuron bit fields (typically 64 bit large). When it is time to update a row of synapses (because their source neuron fired), we read each destination neuron’s firing history and loop over the most recent *age* bits which represent the post-synaptic spikes. This is also sub-optimal in terms of bandwidth since the read accesses are scattered over memory, however,

- they are confined to a smaller region of memory (kilobytes vs. gigabytes) resulting in better cache utilization,
- their addresses are steadily increasing due to the sorted nature of the adjacency list, resulting in better cache coherency, and
- they are only performed once for (up to) 64 update steps, increasing our effective bandwidth.

From here, implementing event-driven plasticity is trivial: Instead of looping over all *age* bits, we only loop over the set ones which can be done very efficiently using integer intrinsic `__clz()` (Fig. 2(c)). In order to benefit from event-driven plasticity, synapse dynamics must have a closed-form solution so they can “skip ahead” multiple steps at a time. If this is not possible, users of our simulator can simply loop internally in which case we gracefully degrade back to lazy plasticity.

B. Spike Delivery with Shared Atomics

Whenever a neuron fires, its neighbors must be informed (potentially after a model-specific delay). The naïve approach of (in parallel, using a CUDA block) looping over the corresponding row in the adjacency list and invoking the user-defined `deliver()` callback for each neighbor (Fig. 3(a)), works fairly well: Reads from the adjacency list get fully coalesced. Writes to the neurons *are* scattered, the negative effects of which get mostly mitigated by the cache.¹

In [7, Section III-C] we introduced an effective optimization which we call “cache-aware” spike delivery: Instead of traversing the adjacency list row-wise, we traverse it per column or, more accurately, per 32-column slice so as to maintain full coalescing of the adjacency data. While writes are still scattered, they are now statistically expected to be closer together, resulting in a better cache hit rate. Cache-aware spike delivery is strictly faster than the naïve one but only really shines once neurons stop fitting into cache. At that point naïve delivery starts to grow quadratically with network size while cache-aware delivery continues to grow linearly. Cache-aware spike delivery is slightly sub-optimal in the sense that metadata (such as information about the source neuron) need to be re-read for every warp whereas previously they only had to be read once per block. This is negligible in practice though.

¹Writes are scattered so far apart that we would expect our effective bandwidth to be $1/32$ nd of the theoretical bandwidth as each 4 B word read/write would result in its own memory transaction (128 B large on CUDA devices). Thanks to the cache, this is far from the case in practice.

```

// adjacency list
extern int adj[/*rows*/][/*columns*/];
// offsets into 'adj' marking the start index of each slice
extern int pivots[/*rows*/][/*columns*/];
struct neuron { int id; ... };
// user-defined callback
void deliver(const neuron& source, neuron& destination);

1  parfor src in firing neurons:
2      parfor dst in neighbors(src):
3          deliver(src, dst);

1  parfor slice in 0..|neurons| / 1024 - 1:
2      __shared__ neuron shared[1024];
3
4      parfor i in 0..1023:
5          shared[i] = neurons[slice * 1024 + i];
6
7      parfor src in firing neurons:
8          parfor dstID in adj[src.id][
9              pivots[src.id][slice]..
10             pivots[src.id][slice+1]-1]:
11             deliver(src, shared[dstID % 1024]);
12
13     parfor i in 0..1023:
14         neurons[slice * 1024 + i] = shared[i];

```

Fig. 3. Pseudocode contrasting naïve with shared memory-based spike delivery. **(a) Naïve delivery.** We loop over all firing neurons (line 1) and inform their neighbors (line 2) about the spike by invoking the user-provided callback `deliver()` (line 3) which is part of the model definition. **(b) Shared memory-based delivery.** We cooperatively load a 1024-neuron slice from global into shared memory (lines 4–5). Once again we deliver all spikes, but only to neurons that fall into our slice. The delivery is carried out in shared memory (lines 7–11). When done we write the new neuron state back to global memory (lines 13–14). In practice we assign a 1024-thread block to each slice, establishing a 1:1 mapping between threads and neurons which greatly simplifies code (the `parfors` in lines 4 & 13 turn into simple assignments). We then assign each warp within the block to one spike and finally each thread within the warp to one destination neuron. This code assumes that the total number of neurons is a multiple of 1024, for brevity. Additional logic is needed to guard against out-of-bounds accesses.

In this paper we drive cache-aware spike delivery to its conclusion: We logically divide the neuron domain into chunks of equal size, say 1024 neurons each. We then partition the adjacency list into slices that index a single and only a single chunk, which is simply achieved by pre-computing $\lceil \frac{n}{1024} \rceil$ pivots per row via binary search (Fig. 1 (red dotted lines)). This turns our statistical expectation into a mathematical certainty: Knowing that slice i only indexes neurons $[i * 1024, (i + 1) * 1024)$ allows us to load them into shared memory and deliver all spikes there (Fig. 3(b)). Doing so:

- minimizes global memory traffic (only the initial/final neuron states have to be read from/written to global memory),




Fig. 4. Simulation time as a function of network size for each model: We measure the time it takes to simulate 10 s of biological time for various synapse counts and report *wall-clock time* \div *biological time*. Note that in the Brunel(+) case both axes are logarithmic.

- uses faster shared memory,
- uses faster shared atomics, and
- circumvents the cache entirely (avoiding cache misses/e-
victions/etc.)

Shared memory-based spike delivery is strictly faster than cache-aware spike delivery (regardless of the number of neurons) for practical network densities and shared memory sizes.

IV. RESULTS

We compare the performance of our simulator to that of BSim [16], GeNN [5], and NeuronGPU [47] using three well-established models: Vogels-Abbott (“Vogels”) and Brunel (with and without plasticity, “Brunel(+)”), detailed in [6, Appendixes A & B]. We apply a scaling factor to synaptic weights allowing us to vary the network size while maintaining the overall firing pattern [1]. The models are visualized in Fig. 6.

When comparing ourselves to other simulators we enable all optimizations. The impact of individual optimizations is analyzed in Section IV-C.

All the code used for the experiments can be found at:

- **BSim** github.com/denniskb/bsim, forked from master as of Feb 19, 2020.
- **GeNN** github.com/denniskb/genn, forked from tag “GeNN 4.4.0” as of Jan 5, 2021.
- **NeuronGPU** github.com/denniskb/neurongpu, forked from master as of Oct 20, 2020.
- **Spice** github.com/denniskb/spice/tree/gather, as of May 7, 2021.

All benchmarks were performed on a Google Cloud VM with an Intel Xeon E5-2699 v3, a Nvidia Tesla V100 16 GB, and 256 GB RAM, running a headless Ubuntu 20 with CUDA 11 and GCC 9.

A. Simulation Time as a Function of Network Size

We measure the time it takes to simulate 10 s of biological time for various network sizes (synapse counts). We report *wall-clock time* \div *biological time* (Fig. 4).

NeuronGPU uses double precision arithmetic and “exact integration” [49] as opposed to single precision arithmetic and

Euler integration used by the other simulators—we do not expect it to perform on par but include it for completeness, to gauge how much performance one has to sacrifice for biological fidelity.

GeNN offers the choice between three different connectivity types: `SPARSE_GLOBALG` is similar to an adjacency list, `BITMASK_GLOBALG` is similar to a binary adjacency matrix, and `PROCEDURAL_GLOBALG` does not store the graph at all but generates it on the fly. We found that `SPARSE_GLOBALG` is faster for Vogels while `BITMASK_GLOBALG` is faster for Brunel(+). The latter is also more memory efficient for dense networks, allowing GeNN to simulate much larger instances of Brunel(+).

While both BSim and NeuronGPU support STDP, their synapse types do not quite match the behavior of Brunel+. According to the authors, modifying them “currently is a task for developers, not for users”.

B. Setup Time as a Function of Network Size

Fast setup is important as it allows more experiments to run in quick succession and thus speeds up network design and parameter tuning. We measure the time in seconds it takes to initialize Vogels for various network sizes (Fig. 5(left)). GeNN and Spice perform this on the GPU which is why they are orders of magnitude faster. In GeNN, *any* parameter change requires recompilation which takes ~ 15 s on the test machine (green, dashed line). Our network construction is so fast (~ 200 M synapses/ms) that setup is dominated by memory allocations and thus virtually constant w.r.t. network size.

BSim and NeuronGPU also use a lot of RAM during setup, peaking at 200 GB for BSim and 40 GB for NeuronGPU [7, Table 1], which might be prohibitive for some users.

C. Impact of Optimizations on Simulation Time

In the preceding plots all optimizations were enabled for Spice. In this section we analyze the relative speedup that each optimization contributes to the final simulation time, compared to previous implementations (Fig. 5(middle+right)).

To the best of our knowledge, nobody uses naïve plasticity due to how slow it is (Fig. 5(middle)). However, it makes for




Fig. 5. (left) Setup time as a function of network size for Vogels. The y-axis is logarithmic. The green, dashed line represents setup time + compilation time. (middle) Speedup due to various plasticity optimizations. The baseline is naïve plasticity which updates every synapse at each simulation step. Shown are the speedups achieved due to lazy plasticity alone, as well as lazy plasticity combined with event-driven plasticity. (right) Speedup due to shared memory-based spike delivery. The baseline is naïve spike delivery relying on global atomics and the cache.

a good baseline because it is very close to a “first implementation” when porting/writing an algorithm to/in CUDA. Lazy plasticity (with a firing history of 64 steps) is 5 times faster. Note that both algorithms perform the exact same number of computations. The speedup is only due to decreased global memory traffic. Beyond 64 steps, diminishing returns kick in: While the effective bandwidth continues to increase, dynamics’ computational costs stay constant and start to dominate the simulation time. Going from 64 to 128 steps is only $\sim 15\%$ faster while it increases code verbosity (emulating 128-bit integers) and might not even generalize to other, higher-activity models. Event-driven plasticity is another 6 times faster, which might seem little when considering that most neurons reach their maximum age and their firing histories only have 0–1 bits set. However, internal tests revealed that we are operating close to the theoretical limit of our approach: When bypassing the event-driven plasticity logic entirely and “blindly” updating synapses in a single step, we only observed a 25% increase in performance. That is not to say that there is not a *fundamentally different*, potentially much faster approach.

Once again, the straightforward, row-wise spike-delivery algorithm serves as the baseline (Fig. 5(right)). Shared memory-based spike delivery achieves its $2\times$ – $2.5\times$ speedup with a slice width of 1024 neurons. The slice width is a delicate, tunable parameter. Increasing it increases the average gap between pivots which leads to better memory bandwidth when reading adjacency data and higher parallelism when delivering spikes. At the same time though, it leads to higher shared memory consumption which can inhibit parallelism again. 1024 neurons strike a good balance in our benchmarks and have the additional benefit of allowing us to map them 1:1 to CUDA threads, simplifying code.

V. SUMMARY AND FUTURE WORK

We presented two algorithms for the efficient facilitation of STDP and spike delivery inside SNNs, which significantly outperform the state of the art. They can be retrofitted to existing simulators with minimal code additions and without

the need for fundamental changes to the code architecture. STDP requires the allocation and update of per-neuron firing histories and in turn makes reverse adjacency lists obsolete. Spike delivery requires the computation of pivots which index the existing graph representation. In the case of adjacency matrices, the pivots do not even have to be stored but can be inferred.

A great feature that would interoperate seamlessly with our presented optimizations is graph compression. The way we partition the adjacency list lends itself to a very simple block compression scheme: Since every entry indexes only 1024 neurons, 10 bits are sufficient to represent it. If we packed 3 consecutive indices into one 32-bit integer, we could reduce memory consumption by almost $2/3$ rd. This would allow us to triple the size of static models regardless of their density.




Fig. 6. Visualization of Vogels and Brunel(+). (top left) Graph of Vogels. (top right) Graph of Brunel(+). (middle) Firing pattern of Vogels. Each row represents a neuron, each column a time step, each dot a spike. (bottom) Firing pattern of Brunel (Brunel+ fires very similarly).

REFERENCES

- [1] D. Bautembach, I. Oikonomidis, N. Kyriazis, and A. Argyros, "Faster and simpler snn simulation with work queues," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–8.
- [2] M. Mattia and P. Del Giudice, "Efficient event-driven simulation of large networks of spiking neurons and dynamical synapses," *Neural Computation*, vol. 12, no. 10, pp. 2305–2329, 2000.
- [3] S. Song, K. D. Miller, and L. F. Abbott, "Competitive hebbian learning through spike-timing-dependent synaptic plasticity," *Nature neuroscience*, vol. 3, no. 9, pp. 919–926, 2000.
- [4] A. Morrison, A. Aertsen, and M. Diesmann, "Spike-timing-dependent plasticity in balanced random networks," *Neural computation*, vol. 19, no. 6, pp. 1437–1467, 2007.
- [5] E. Yavuz, J. Turner, and T. Nowotny, "GeNN: A code generation framework for accelerated brain simulations," *Scientific Reports*, vol. 6, no. June 2015, pp. 1–14, 2016. [Online]. Available: <http://dx.doi.org/10.1038/srep18854>
- [6] N. Ahmad, J. B. Isbister, T. S. C. Smithe, and S. M. Stringer, "Spike: A GPU Optimised Spiking Neural Network Simulator," *bioRxiv*, p. 461160, 2018. [Online]. Available: <https://www.biorxiv.org/content/early/2018/11/06/461160>
- [7] D. Bautembach, I. Oikonomidis, and A. Argyros, "Multi-gpu snn simulation with static load balancing," 2021.
- [8] H. C. Tuckwell, *Introduction to theoretical neurobiology. Vols. 1 and 2*. Cambridge University Press, 1988.
- [9] W. S. McCulloch and W. H. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics*, vol. 5, pp. 115–133, 1943.
- [10] L. Lapique, "Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation." *Journal of Physiology and Pathology*, vol. 9, pp. 620–635, 1907.
- [11] F. J. Pelayo, E. Ros, X. Arreguit, and A. Prieto, "VLSI Implementation of a Neural Model Using Spikes," *Neuron*, vol. 121, pp. 111–121, 1997.
- [12] J. Reutimann, "Event-driven simulation of spiking neurons with stochastic dynamics," no. 2593, 2002.
- [13] M. A. Van Der Vlag, "Multi-GPU Brain: A multi-node implementation for an extended Hodgkin-Huxley simulator," 2019.
- [14] M. Mozafari, M. Ganjtabesh, A. Nowzari-Dalini, and T. Masquelier, "SpykeTorch: Efficient Simulation of Convolutional Spiking Neural Networks with at most one Spike per Neuron," no. Mm, pp. 1–16, 2019. [Online]. Available: <http://arxiv.org/abs/1903.02440>
- [15] S. Panagiotou, R. Miedema, H. Sidiropoulos, G. Smaragdus, C. Strydis, and D. Soudris, "A novel simulator for extended Hodgkin-Huxley neural networks," pp. 395–402, 2020.
- [16] P. Qu, Y. Zhang, X. Fei, and W. Zheng, "High Performance Simulation of Spiking Neural Network on GPGPUs," *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 11, pp. 2510–2523, 2020.
- [17] Z.-Q. K. Tian and D. Zhou, "Library-based Fast Algorithm for Simulating the Hodgkin-Huxley Neuronal Networks," pp. 1–19, 2021. [Online]. Available: <http://arxiv.org/abs/2101.07257>
- [18] T. S. Chou, H. J. Kashyap, J. Xing, S. Listopad, E. L. Rounds, M. Beyeler, N. Dutt, and J. L. Krichmar, "CARLsim 4: An Open Source Library for Large Scale, Biologically Detailed Spiking Neural Network Simulation using Heterogeneous Clusters," *Proceedings of the International Joint Conference on Neural Networks*, vol. 2018-July, pp. 1158–1165, 2018.
- [19] G. Smaragdus, G. Chatzikonstantis, R. Kukreja, H. Sidiropoulos, D. Rodopoulos, I. Sourdis, Z. Al-Ars, C. Kachris, D. Soudris, C. I. De Zeeuw, and C. Strydis, "BrainFrame: A node-level heterogeneous accelerator platform for neuron simulations," *Journal of Neural Engineering*, vol. 14, no. 6, 2017.
- [20] A. Sripad, G. Sanchez, M. Zapata, V. Pirrone, T. Dorta, S. Cambria, A. Marti, K. Krishnamourthy, and J. Madrenas, "SNAVA—A real-time multi-FPGA multi-model spiking neural network simulation architecture," *Neural Networks*, vol. 97, pp. 28–45, 2018. [Online]. Available: <https://doi.org/10.1016/j.neunet.2017.09.011>
- [21] S. Kunkel, M. Schmidt, J. M. Eppler, H. E. Plesser, G. Masumoto, J. Igarashi, S. Ishii, T. Fukai, A. Morrison, M. Diesmann, and M. Helias, "Spiking network simulation code for petascale computers," *Frontiers in Neuroinformatics*, vol. 8, no. October, pp. 1–23, 2014.
- [22] M. C. W. Van Rossum, "A novel spike distance," *Neural Computation*, vol. 13, no. 4, pp. 751–763, 2001.
- [23] M. Rudolph and A. Destexhe, "Analytical integrate-and-fire neuron models with conductance-based dynamics for event-driven simulation strategies," *Neural Computation*, vol. 18, no. 9, pp. 2146–2210, 2006.
- [24] M. S. Tomlinson, "Spike Transmission for Neural Networks," 1990.
- [25] E. Ros, E. M. Ortigosa, R. Agís, R. Carrillo, and M. Arnold, "Real-time computing platform for spiking neurons (RT-spike)," *IEEE Transactions on Neural Networks*, vol. 17, no. 4, pp. 1050–1063, 2006.
- [26] J. Schemmel, D. Brüderle, A. Grünbl, M. Hock, K. Meier, and S. Millner, "A wafer-scale neuromorphic hardware system for large-scale neural modeling," *ISCAS 2010 - 2010 IEEE International Symposium on Circuits and Systems: Nano-Bio Circuit Fabrics and Systems*, pp. 1947–1950, 2010.
- [27] C. M. Thibault, R. Hoang, and F. C. Harris, "A novel multi-GPU neural simulator," *3rd International Conference on Bioinformatics and Computational Biology 2011, BICoB 2011*, no. 1, pp. 146–151, 2011.
- [28] A. Antonietti, C. Casellato, J. A. Garrido, N. R. Luque, F. Naveros, E. Ros, E. D'Angelo, and A. Pedrocchi, "Spiking neural network with distributed plasticity reproduces cerebellar learning in eye blink conditioning paradigms," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 1, pp. 210–219, 2016.
- [29] D. Lee, G. Lee, D. Kwon, S. Lee, Y. Kim, and J. Kim, "Flexon: A flexible digital neuron for efficient spiking neural network simulations," *Proceedings - International Symposium on Computer Architecture*, pp. 275–288, 2018.
- [30] M. A. Van Der Vlag, G. Smaragdus, Z. Al-Ars, and C. Strydis, "Exploring complex brain-simulation workloads on multi-GPU deployments," *ACM Trans. on Architecture and Code Optimization*, no. 4, 2019.
- [31] D. F. M. Goodman, "The Brian simulator," *Frontiers in Neuroscience*, vol. 3, no. 2, pp. 192–197, 2010.
- [32] D. Pecevski, D. Kappel, and Z. Jonke, "NEVESIM: event-driven neural simulation framework with a Python interface," *Frontiers in Neuroinformatics*, vol. 8, no. August, pp. 1–20, 2014.
- [33] J. C. Knight and T. Nowotny, "GPUs Outperform Current HPC and Neuromorphic Solutions in Terms of Speed and Energy When Simulating a Highly-Connected Cortical Model," *Frontiers in Neuroscience*, vol. 12, no. December, pp. 1–19, 2018.
- [34] M. Stimberg, D. F. M. Goodman, and T. Nowotny, "Brian2GeNN: a system for accelerating a large variety of spiking neural networks with graphics hardware," *bioRxiv*, p. 448050, 2018. [Online]. Available: <https://www.biorxiv.org/content/early/2018/10/20/448050>
- [35] H. Hazan, D. J. Saunders, H. Khan, D. T. Sanghavi, H. T. Siegelmann, and R. Kozma, "BindsNET: A machine learning-oriented spiking neural networks library in Python," vol. 12, no. December, pp. 1–18, 2018. [Online]. Available: <http://arxiv.org/abs/1806.01423> {%}0Ahttp://dx.doi.org/10.3389/fninf.2018.00089
- [36] E. Ros, R. Carrillo, E. M. Ortigosa, B. Barbour, and R. Agís, "Event-driven simulation scheme for spiking neural networks using lookup tables to characterize neuronal dynamics," *Neural Computation*, vol. 18, no. 12, pp. 2959–2993, 2006.
- [37] A. Hanuschkin, S. Kunkel, M. Helias, A. Morrison, and M. Diesmann, "A General and Efficient Method for Incorporating Precise Spike Times in Globally Time-Driven Simulations," *Frontiers in Neuroinformatics*, vol. 4, no. October, pp. 1–19, 2010.
- [38] F. Naveros, N. R. Luque, J. A. Garrido, R. R. Carrillo, M. Anguita, and E. Ros, "A Spiking Neural Simulator Integrating Event-Driven and Time-Driven Computation Schemes Using Parallel CPU-GPU Co-Processing: A Case Study," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 7, pp. 1567–1574, 2015.
- [39] F. Naveros, J. A. Garrido, R. R. Carrillo, E. Ros, and N. R. Luque, "Event- and Time-Driven Techniques Using Parallel CPU-GPU Co-processing for Spiking Neural Networks," *Frontiers in Neuroinformatics*, vol. 12, no. February, pp. 1–22, 2018.
- [40] A. Delorme and S. J. Thorpe, "SpikeNET: An event-driven simulation package for modelling large networks of spiking neurons," *Network: Computation in Neural Systems*, vol. 14, no. 4, pp. 613–627, 2003.
- [41] H. E. Plesser, J. M. Eppler, A. Morrison, M. Diesmann, and M. O. Gewaltig, "Efficient parallel simulation of large-scale neuronal networks on clusters of multiprocessor computers," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 4641 LNCS, pp. 672–681, 2007.
- [42] J. A. Garrido, R. R. Carrillo, N. R. Luque, and E. Ros, "Event and Time Driven Hybrid Simulation of," no. June 2011, 2014.
- [43] B. Kasap and A. J. van Opstal, "Dynamic parallelism for synaptic updating in GPU-accelerated spiking neural network simulations,"

- Neurocomputing*, vol. 302, pp. 55–65, 2018. [Online]. Available: <https://doi.org/10.1016/j.neucom.2018.04.007>
- [44] P. Szykiewicz, “A novel GPU-enabled simulator for large scale spiking neural networks,” *Journal of Telecommunications and Information Technology*, vol. 2016, no. 2, pp. 34–42, 2016.
- [45] P. Krishnamani and V. Venkittaraman, “Acceleration of spiking neural networks on single-GPU and multi-GPU systems,” no. May, p. 81, 2010. [Online]. Available: <http://gradworks.umi.com/14/75/1475559.html>
- [46] 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,” *IEEE Transactions on Computers*, vol. 62, no. 12, pp. 2454–2467, 2013.
- [47] B. Golosio, G. Tiddia, C. De Luca, E. Pastorelli, F. Simula, and P. S. Paolucci, “A new gpu library for fast simulation of large-scale networks of spiking neurons,” *arXiv preprint arXiv:2007.14236*, 2020.
- [48] J. C. Knight and T. Nowotny, “Larger gpu-accelerated brain simulations with procedural connectivity,” *Nature Computational Science*, vol. 1, no. 2, pp. 136–142, Feb 2021. [Online]. Available: <https://doi.org/10.1038/s43588-020-00022-7>
- [49] S. Rotter and M. Diesmann, “Exact digital simulation of time-invariant linear systems with applications to neuronal modeling,” *Biological Cybernetics*, vol. 81, no. 5, pp. 381–402, Nov 1999. [Online]. Available: <https://doi.org/10.1007/s004220050570>