



Transitioning Spiking Neural Network Simulators to Heterogeneous Hardware

Quang Anh Pham Nguyen
TUM Create Ltd. and
Nanyang Technological University, Singapore
nqapham@ntu.edu.sg

Philipp Andelfinger
TUM Create Ltd. and
Nanyang Technological University, Singapore
pandelfinger@ntu.edu.sg

Wentong Cai
Nanyang Technological University, Singapore
aswtcai@tum-create.edu.sg

Alois Knoll
Technische Universität München, Germany and
Nanyang Technical University, Singapore
knoll@in.tum.de

ABSTRACT

Spiking neural networks (SNN) are among the most computationally intensive types of simulation models, with node counts on the order of up to 10^{11} . Currently, there is intensive research into hardware platforms suitable to support large-scale SNN simulations, whereas several of the most widely used simulators still rely purely on the execution on CPUs. Enabling the execution of these established simulators on heterogeneous hardware allows new studies to exploit the many-core hardware prevalent in modern supercomputing environments, while still being able to reproduce and compare with results from a vast body of existing literature. In this paper, we propose a transition approach for CPU-based SNN simulators to enable the execution on heterogeneous hardware (e.g., CPUs, GPUs, and FPGAs) with only limited modifications to an existing simulator code base, and without changes to model code. Our approach relies on manual porting of a small number of core simulator functionalities as found in common SNN simulators, whereas unmodified model code is analyzed and transformed automatically. We apply our approach to the well-known simulator NEST and make a version executable on heterogeneous hardware available to the community. Our measurements show that at full utilization, a single GPU achieves the performance of about 9 CPU cores.

ACM Reference Format:

Quang Anh Pham Nguyen, Philipp Andelfinger, Wentong Cai, and Alois Knoll. 2019. Transitioning Spiking Neural Network Simulators to Heterogeneous Hardware. In *SIGSIM Principles of Advanced Discrete Simulation (SIGSIM-PADS '19)*, June 3–5, 2019, Chicago, IL, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3316480.3322893>

1 INTRODUCTION

Spiking neural networks (SNNs) are artificial neural networks that are used to model and understand the mammalian brain. Due to the scale and complexity of relevant networks, simulations of SNNs

often require enormous amounts of computational resources and substantial running times. For example, on a single CPU core, simulating 250ms of activity for a network of 11 250 neurons and about 127 million synapses using the NEST simulator [17] takes more than 30 seconds. Meanwhile, the brain of a mammal contains on the order of 10^7 to 10^{11} neurons, with thousands of synapses per neuron on average [1, 22, 23]. To tackle the computational demands of large-scale SNN simulations, many of the existing simulators are designed with the capability for parallel execution using multi-core CPUs in a single-node or multi-node environment. In past few years, the focus of works on high performance SNN simulation has shifted towards accelerators such as graphics processing units (GPUs), which are now present in most computers ranging from individual workstations to the largest supercomputers. Several GPU-enabled SNN simulators have been developed with substantial performance improvements over a CPU-based execution [6, 10, 15, 25, 40, 48]. Typically, separate CPU and GPU variants of the relevant segments of the simulator code base have been developed manually. While this approach allows for a large degree of optimization to the target hardware, the simulator and model code is duplicated for each hardware platform, posing significant challenges in terms of maintainability and extensibility. Further, developing code to be executed on accelerators is widely considered to be more cumbersome and error-prone than CPU programming (e.g., [40]). Thus, it is desirable to minimize the need to develop and modify accelerator code directly. Since in contrast to the core simulator functionalities, neuron and synapse models may frequently be added, modified, or extended, the model development process should abstract from the target hardware as much as possible.

Newly developed simulators can achieve a certain degree of hardware independence by a template-based model specification [48]. However, a transition path is needed for existing and widely used SNN simulators such as NEST [17] and Brian 2 [18], which have been developed in general-purpose languages such as C++, targeting the execution on CPUs. Adding support for the execution on heterogeneous hardware to these simulators allows researchers to conduct new simulation studies in a timelier manner, while still relying on widely studied and tested simulator and model implementations. In particular, new studies can easily reproduce previous results from the literature and benefit from existing validation results, using existing configuration files and tool flows.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGSIM-PADS '19, June 3–5, 2019, Chicago, IL, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6723-3/19/06...\$15.00

<https://doi.org/10.1145/3316480.3322893>

This paper presents a semi-automated approach for transformation of SNN simulators to enable execution on heterogeneous hardware. While our experiments are executed on CPUs and GPUs, the transformed simulator can make use of CPUs, GPUs, FPGAs, and DSPs. Our main contributions are as follows:

- (1) We analyze the common architecture of CPU-based SNN simulators, differentiating static components implementing basic simulator functionality, and the portions specifying neuron and synapse models.
- (2) We present an approach based on an automated model code transformation to support the transition from a purely CPU-based simulator code base to an implementation executable on heterogeneous hardware. We propose optimizations for synapse models to reduce their memory consumption.
- (3) We demonstrate our approach on the NEST simulator and present performance measurement results when executing NEST on GPUs, which was not possible previously. Substantial performance gains are achieved over a purely CPU-based execution. Our transformation code is publicly available¹.

2 BACKGROUND AND RELATED WORK

In this section, we briefly introduce spiking neural networks as well as the simulators for this class of network. Further, we outline fundamentals and existing work on the execution of spiking neural network simulations on heterogeneous hardware.

2.1 Spiking neural network models

Spiking neural networks (SNNs) [37] are artificial neural networks that are considered to be more biologically accurate representations of mammal brains than the more abstract neural network models used by common machine learning applications. An SNN is a directed graph with so-called integrate-and-fire neurons as nodes and synapses between two neurons as edges. Integrate-and-fire neurons generate rapid increases or decreases of their membrane potential, so-called spikes, according to the potential changes on the incoming synapses. The spikes are transmitted to the neighboring neurons across the outgoing synapses, which may increase or decrease the spike's potential. The main application areas of SNN models are in computational neuroscience, which aims at understanding the nervous systems, and in machine learning, e.g., in the robotics field [8, 31].

In addition to the network topology, SNNs are characterized by the neuron and synapse models. Neuron models define neuron state variables and actions on incoming spikes, commonly using differential equations. Typically, these equations are solved numerically by time-stepped integration. Each time step updates a neuron's state according to the incoming spikes and its current state, potentially generating outgoing spikes. Synapse models define the way spikes are affected by the transmission through a synapse.

With *static* synapse types, each spike's potential is affected in a constant way as it travels to the target neuron. In contrast, *Spike-timing dependent plasticity* (STDP) models dynamically vary the current spike's potential depending on the current synapse state, which may change with each transmitted spike. An SNN may use several different types of neuron and synapse models.

¹<https://github.com/opencl-nest/opencl-nest>

2.2 SNN simulators

SNN simulators are programs that execute the network's activities in the form of neuron and synapse behaviors over a period of time. Although the local state of each neuron leads to spike patterns that are usually not synchronous across the network, the synchronous execution of SNN models permits a simple barrier-base parallelization. The simulation state is advanced by updating each neuron's state according to the chosen neuron model, which is typically reflected computationally by one time step of numerical integration. The number of integration steps before synchronization is required (*super step*), depends on the *lookahead*, which is a delta in time during which state changes of a neuron are guaranteed not to affect other neurons, as defined by the minimum time required for a spike to travel through a synapse. At the end of each super step, synchronization is achieved by exchanging newly generated spikes among the neurons. In parallel and distributed simulation terminology, this execution scheme is similar to synchronous conservative synchronization using the YAWNS algorithm [42].

A number of SNN simulators have been proposed in the literature, frequently focusing on scalability to large networks. The number of neurons in the brain of a mammal is on the order of 10^7 to 10^{11} , with about 1 000 to 10 000 synaptic connections per neuron on average [1, 22, 23]. Thus, the number of synapses is substantial even in small networks. To support large scale SNN simulation, simulators have been developed targeting high-performance computing environments employing multi-core CPUs and GPUs.

Established simulators such as CARLsim 4 [10], Brian 2 [18], NEURON 7.5 [24], NCS 6 [25], Nemo 0.7 [15], Nengo 2.6 [6], NEST 2.14, HRLSim [40], and PCSIM 0.5 [43] support multi-threaded execution on CPUs, some of them with support for execution on multiple nodes. GENN 3 [48] can be executed on a single GPU. Some other simulators additionally support the execution in GPU clusters [6, 10, 10, 15, 25, 40, 40]. Among these simulators, CARLsim 4 and NCS 6 support a CPU-GPU co-execution.

Several authors have explored the execution of SNN simulations on GPUs independently of the established simulators (e.g., [7, 44]). These works share the approach of manual development and optimization of GPU code. While manually tuned GPU code can provide high performance, maintenance and extensibility is impacted by the presence of hardware-specific code or separate CPU and GPU implementations. This issue is exacerbated when considering the execution on further accelerator types such as FPGAs (e.g., [38]). Among the simulators supporting GPU-based execution, to our knowledge GENN [48] is the only one to support automatic GPU code generation. Users of GENN define neuron models using a C++ interface provided by GENN. The main inputs provided by the user are the model parameter names and the neuron update rules in the form of C++ statements. A drawback of both hand-tuned and generated GPU code is the lack of comparability of results to those generated using the well-tested and well-studied code and models of established simulators. The transformation approach proposed in our present paper enables the use of unmodified models of an existing CPU-based simulator, while reducing the simulation running time using hardware accelerators. While our experiments are performed on GPUs, the generated OpenCL code enables the execution on further hardware types such as FPGAs or DSPs.

2.3 Heterogeneous computing and graphics processing units

In the past two decades, high-performance computing has moved from mostly CPU-based platforms to heterogeneous architectures. While a range of accelerator types have emerged that provides performance benefits for certain classes of computational problems, graphics processing units (GPUs) are currently the dominant hardware type used to supplement CPU-based host systems with resources for massively data-parallel processing. Of the list of the top 500 supercomputers published in November 2018², 126 rely on GPUs. The benefits of CPU-GPU co-processing or purely GPU-based execution have been shown for many scientific applications (e.g., [14, 19, 35]), including several types of network simulations [2, 3, 46].

In NVIDIA's terminology, a GPU contains several *streaming multiprocessors (SMs)* with up to thousands of arithmetic units in total. This enables a GPU to execute thousands of lightweight threads in parallel to provide teraflops of computing power. However, the design of efficient GPU algorithms requires exploiting the GPU's unique architectural characteristics. The programs to be executed by the GPU's threads are referred to as *kernels*. On current NVIDIA GPUs, 32 consecutive GPU threads are grouped into a *warp*, in which threads run the same sequence of instructions in lockstep. If the program executed by a warp includes branches taken by only some of the threads, the branches are serialized, decreasing performance. Another significant factor for GPU performance is the memory access pattern in the GPU's DRAM. If threads in a warp access 32 consecutive memory locations, the accesses are served by a single memory transaction. Conversely, if the memory accesses by a warp do not follow this pattern, multiple transactions are issued, which can severely reduce the overall performance. GPU programs are commonly implemented using frameworks such as CUDA or OpenCL. While the former only targets NVIDIA GPUs, OpenCL supports cross-platform development and deployment on a variety of hardware devices ranging from CPUs to accelerators such as GPUs, APUs, FPGAs, and Intel Xeon Phi. In recent years, these accelerators have shown promising performance results for various simulation problems [47]. To permit execution on a wide range of hardware platforms, the translation approach presented in our present paper generates OpenCL code. In the remainder of the paper, we refer to the portion of the hardware controlled directly by the CPU as the *host*, and portion executing OpenCL kernels as the *device* or *accelerator*.

2.4 Code transformation and generation for heterogeneous computing

An automated transformation of a CPU program for efficient execution on heterogeneous platforms often requires the parallelization of suitable program portions. When transforming sequential into parallel code, loops are important sources of parallelism. In automated loop transformations, data dependencies between loop iterations are analyzed. Subsequently, the computations are reordered to maximize the amount of parallelizable operations, while satisfying

all data dependencies. The code analysis frequently relies on knowledge of the problem class at hand. For instance, existing works propose tiling approaches for the successive over-relaxation method in linear algebra [13] and for stencil-based loop computations [9]. Da Li et al. [34] present a template-based approach to parallelization of irregular nested loops and recursive computations for GPUs. Aside from maximizing the exploited parallelism, a second challenge in loop transformation targeting GPUs lies in achieving memory access locality [4]. Hou et al. [27] address the issue of memory access locality by reordering computational operations in wavefront loops without affecting the correctness of the results. Fully automated parallelization is possible for a certain class of nested loops using the so-called polyhedral (or polytope) model [33], which permits a compile-time analysis of data dependencies for affine programs. Several works propose approaches to transform affine programs for parallelized execution on GPUs [5, 45]. While the above approaches can be carried out without domain knowledge, their applicability is limited to programs of a certain structure, e.g., nested loops with predictable control flow.

Another development approach for heterogeneous computing environments is to formulate programs in a *domain-specific language* (DSL) and to generate code for the target platforms. By providing abstractions tailored to the given problem domain, DSLs are often more concise than general-purpose languages. Further, DSLs can be used to hide low-level or hardware-dependent implementation details. Many DSLs have been proposed to solve computational tasks such as graph problems [26], image processing [30, 39], and social network analysis [16]. A number of DSLs have been proposed to achieve heterogeneous execution for simulations. The DSL by Hawick et al. [20] permits the generation of simulation programs for many problems based on partial differential equations. Similarly, Devito et al. [12] propose a language to build portable mesh-based PDE solvers that can run on multiple platforms. In 2018, Cosenza et al. [11] presented the OpenABL language, which allows users to formulate agent-based simulation models in a hardware-independent manner. The generated code can be executed using a number of existing frameworks running either on CPUs or GPUs.

Since our goal is to provide support for heterogeneous execution of existing simulators, we do not define our own DSL. Instead, we rely on knowledge of the structure of SNN simulators to enable source-to-source translation from C++ to OpenCL.

3 ANALYSIS OF SPIKING NEURAL NETWORK SIMULATORS

Conceptually, SNN simulations follow a common structure that is reflected in the software architecture of simulators and in the computations involved in the process of an SNN simulation run. In the following, we describe the general structure of SNN simulations as a basis for the transition approach presented in Section 4.

3.1 Architecture

An SNN simulation can be seen as a collection of simulations for individual neurons that interact among each other by the exchange of spikes. Often, each neuron is simulated in a time-stepped fashion by updating the neuron's internal state at each step. The changes in neuron state at a time step may trigger zero or more spikes

²<https://www.top500.org>




Figure 1: A sequence of SNN simulation steps. The synaptic delay guarantees that neurons cannot interact within a super step. Thus, inter-neuron synchronization is required only after a super step has been processed.

that are then sent to neighboring neurons. Since the spikes must be considered in the target neurons' future state updates, a neuron's state can only be updated once it has received all spikes with smaller timestamps. From this outline of an SNN simulation, we can see that the two main operations are the **neuron update** and the **spike delivery**, the facilities for which commonly form the main components of an SNN simulator: The **neuron update** component receives the current states and incoming spikes as input and performs a simulation step for each neuron. The output of this computation step is composed of the neuron's new state and any newly emitted spikes. The **spikes delivery** component is responsible for transferring the spikes among neurons in the network. The delivery may either occur locally, if the connected neurons are simulated in the same process, or remotely via inter-process communication or a physical network.

An important aspect of SNN simulations is the delay in spike transmission reflecting the synaptic delay, i.e., the number of time steps taken by the signal to arrive at and have impact on the target neuron. Simulators such as NEST [17] exploit this characteristic in a synchronization scheme similar to the YAWNS algorithm [42] in order to avoid transferring all new spikes at every time step. Instead, the simulation time is divided into super steps with the size of the minimal delay in the neural network, with all spikes generated in a super step only being sent at the end of the super step (cf. Figure 1). With this approach, it is guaranteed that all neurons receive the incoming spikes in time, while the frequency of synchronization is reduced. The super steps also enable larger amounts of computation between synchronization points, which is beneficial for parallel computation on a GPU.

Since there are many different neuron and synapse models that can be used in SNN simulations, simulators typically provide interfaces to allow other simulator components to interact with the *neuron update* and *spike delivery* components regardless of the specifics of the underlying models. Since neuron and synapse models may be added, extended, or modified frequently, we consider the models the **dynamic** parts of a simulator.

The remaining simulator components are independent of the models and typically remain unchanged during modeling and when executing simulations using different combinations of models and parameters. We refer to these as **static** components. For instance, the static components in NEST are the *SimulationManager*, which

orchestrates the overall simulation (e.g., by advancing time and triggering synchronization) and the *ConnectionManager*, which provides access to the topology defined by the synapses.

3.2 Neuron and synapse models

Neuron models are differentiated by two aspects:

- (1) **Neuron state:** Each type of neuron uses a set of variables (e.g., the current membrane potentials) to represent the neuron's state at a given time. Further, some configurable constants (e.g., the membrane capacitance and resting potential) may exist that affect the neuron's behavior.
- (2) **Update function:** The update function defines how the state in the next time step is calculated from the current state and the incoming spikes. If certain conditions are met, output spikes are generated to be transmitted to other neurons. Since neuron models are often specified in terms of differential equations, the neuron update function frequently involves performing one iteration of a numerical integration per super step.

As described in Section 2.1, synapse models can be either of *STDP* or *static* type. Similarly to neuron models, they can be differentiated according to their state and their update functions:

- (1) **Synapse state:** Each synapse in the network may have state variables and parameters. The most important state variables are the weight, which affects a spike's weight on arrival at the target node, and the synaptic delay, which is the time required for a new spike to be received by the target neuron. *STDP* models reflect a synapse's plasticity, i.e., the synapse state variables may be modified each time a spike is transmitted. In contrast, *static* synapse models affect the spike weight in a fixed manner, independently of the state variables. Since *STDP* models rely on the current value of the state variables to determine each spike's weight, *STDP* synapses may require substantially larger amounts of memory for their internal data than static models.
- (2) **Synapse update function:** In *STDP* models, each transmitted spike may modify the values of the synapse's state variables. To achieve correct updates of the variables, spikes must be processed according to timestamp order, which is an important constraint when considering parallelized execution. Static synapse models do not require updates of the state variables. Instead, a constant weight value is assigned to the spike before it is forwarded to the target neuron. Thus, while the transmit function for *STDP* models may require significant amounts of computation, the execution of the transmit function of static synapse models is reflected mainly by the memory accesses required to store new spikes at the target neurons.

An implication of the above descriptions is that while the neuron updates within each super step can be trivially parallelized across neurons, it must be ensured that the spike transmission and update steps at a *STDP* synapse follow the temporal order of the spikes.

4 PROPOSED TRANSITION APPROACH

In this section, we propose an approach to generate a GPU implementation from a CPU-based SNN simulator in a semi-automated fashion. The approach follows the differentiation of static and dynamic simulator components from Section 3.1, which makes it applicable to several CPU-based SNN simulators: while some basic simulator functionalities are ported manually, neuron and synapse models are analyzed and translated to code executable on heterogeneous hardware in an automated fashion. In contrast, the computationally inexpensive portions of the simulator code remain on the CPU.

4.1 Static and dynamic components

As discussed in the previous section, SNN simulators can be regarded as being comprised of **static** and **dynamic** components. The *static* components constitute the simulator core of the simulator, carrying out the basic operations in an SNN simulation such as the creation of neurons and synapses, the triggering of neuron updates and the exchange of spikes, synchronization among processes, and the time advancement. These essential procedures are carried out by every SNN simulation, regardless of the neuron and synapse models used. Thus, the corresponding components are rarely modified between simulation experiments or in the development cycle of the simulator. The remaining part of an SNN simulator is comprised of the neuron and synapse models, which are used to construct networks. A simulation experiment may use any combination of models from a model library. Further, experiments may require changes or extensions to existing neuron or synapse models, or the development of new models. Thus, the models constitute the *dynamic* part of the simulator code base. A common interface for neuron and synapse models is often provided so that the static components can communicate with the dynamic components (cf. Figure 2). The resulting loose coupling among components simplifies the maintenance of the simulator code base, as *static* components are not affected by changes to the *dynamic* components such as the addition, modification, or extension of models. The loose coupling also enables the development of multiple implementations, i.e., targeting CPUs and GPUs, for the same neuron or synapse model. The GPU-accelerated simulators NCS [25] and Nemo [15] as well as the purely CPU-based NEST simulator rely on such an architecture. In contrast, CARLsim [10], which provides support for GPU execution, does not follow this structure. Instead, it defines separate interfaces for CPU and GPU implementation and the simulator explicitly selects an implementation at runtime.

Since we assume that the *static* components of an SNN simulator are modified only rarely, the effort of a manual translation of code for execution on heterogeneous hardware can be justified. In addition, the static components are loosely coupled to the model implementations, allowing for the required code changes to be reasonably small and self-contained. For example, in our implementation described in Section 5, the key code modifications to add accelerator support were limited to only a few source files.

The same strategy cannot be applied to the *dynamic* components: firstly, the variety of neuron and synapse models makes a manual translation process cumbersome and error-prone. For instance, version 2.14 of the NEST simulator distribution includes more than 50

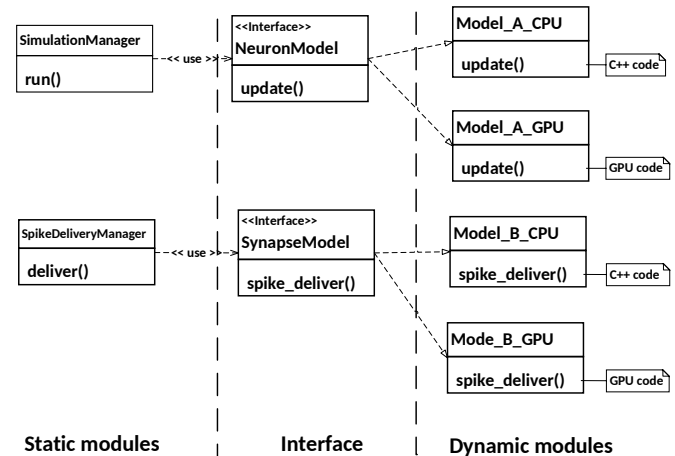


Figure 2: Static and dynamic components in an SNN simulator.

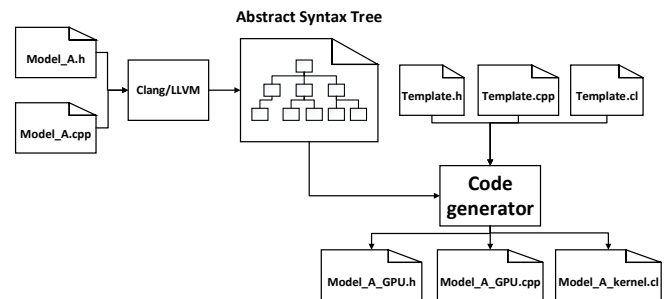


Figure 3: Compilation workflow from C++ SNN simulator code to OpenCL.

neuron and 10 synapse models. Further, it may frequently be necessary to add, modify, or extend models to carry out SNN studies. To handle the diverse and dynamic nature of the models, we propose a tool to automate the transformation of neuron and synapse models for execution on heterogeneous hardware. The transformation relies on the fact that most models share a similar basic structure, with major differences only in the specific computations performed. The transformation tool exploits the similarities among models by applying a template common to a class of models, which is then populated with the model-specific computations.

4.2 Transformation of neuron models

In this section, we present the steps required for porting neuron models to code executable on heterogeneous hardware. The overall workflow, which applies both to neuron and synapse models, is illustrated in Figure 3. In the first step, we parse and analyze the model’s source code. We assume that each neuron model is defined by a class with the neuron variables being class members. The result of the code analysis is an abstract syntax tree (AST) describing the computations performed in the neuron model code, as well as their order. Relying on the fact that the models follow the same high-level structure, we can identify the code snippet

Algorithm 1 A super step in a CPU-based SNN simulation.

```

1: function SIMULATIONUPDATE(from, to)
2:   for i ← number of neurons do
3:     neuron[i].Update(from, to)
1: function UPDATE(from, to)
2:   for lag ← from; lag < to; lag ← lag + 1 do
3:     ModelBehavior()

```

Algorithm 2 A super step in a heterogeneous SNN simulator.

```

1: function SIMULATIONUPDATE(from, to)
2:   MassUpdate(neurons, from, to)
1: function MASSUPDATE(neurons, from, to)
2:   if first update then
3:     CopyDataHostToDevice(neurons)
4:   for lag ← from; lag < to; lag ← lag + 1 do
5:     ModelBehaviorGPU()
6:     CopyDataDeviceToHost(neurons)

```

Figure 4: A super step in a CPU-based SNN simulator and in a simulator supporting heterogeneous hardware.

carrying out the main computations that define the model behavior (*ModelBehavior* in Algorithm 1). The *ModelBehavior* routine will be translated to OpenCL. To do so, the neuron state variables used in the *ModelBehavior* routine must be identified. The AST allows us to determine several properties of the variables used in the routine: from the scope of a variable, we can determine whether it is local to the current function local data or a member variable of the neuron. If a variable represents a state variable of a neuron, information on its type (e.g., integer or double-precision floating point) and dimensionality (scalar or array) allows us to allocate memory for the variable in the OpenCL code. To conserve memory, state variables that are not used in the *ModelBehavior* routine are not allocated in the OpenCL code. Finally, we identify function calls to later replace them with calls to equivalent OpenCL functions (cf. Section 5.2).

After the analysis, OpenCL code is generated based on predefined code templates (cf. Figure 3) for accelerator and host code. A critical part in the generation of the accelerator code is producing data access patterns suitable for the chosen hardware. Since our performance experiments of Section 6 are performed using GPUs, we automatically apply a common memory access optimization applicable to all common GPU accelerators: due to the large number of neurons in a typical SNN, the neuron state variables reside in the GPU’s high-capacity DRAM. Unfortunately, CPU code typically relies on a “array-of-struct” (AoS) representation, where neurons are stored as an array of objects, with the state variables as object members. On a GPU, the AoS representation is known to result in large numbers of memory transactions when operating on many array elements in a data-parallel fashion. Thus, we convert the code to follow a “struct-of-array” (SoA) access pattern by flattening out the neuron data into a one-dimensional array for each variable and accessing the entries using each neuron’s numerical index. Since within a super step the state updates are independent across neurons, we assign one GPU thread to each neuron. In contrast to the neuron state variables, function-local variables rely on GPU registers by default and thus can be accessed efficiently without further code transformations.

Depending on the data used by the model, host-side OpenCL API calls are generated for GPU memory allocation and data transfer from and to the accelerator. Figure 4 provides pseudo code of the original and the transformed host code. The result of the code generation step is a set of C++ files containing host and accelerator code. Further details on the transformation of the neuron models as well as example code are provided in Section 5.2.

4.3 Transformation of synapse models and spike delivery

Since there are two synapse model types (*static* and *STDP*) that differ in their characteristics, we employ separate strategies for their transformation. One of the factors affecting this design decision is the high degree of nodes in neural networks. Motivated by the node degrees in mammals’ brains, a neuron is commonly connected to about 10^4 other neurons. When relying on STDP models, the resulting large numbers of synapses in the network incur substantial requirements both in terms of memory capacity and in terms of computation to transfer spikes across the network. As we will see in Section 6, the spike delivery in fact constitutes the largest portion of the simulation workload. Therefore, it is critical to achieve high performance and low memory usage for STDP models.

In contrast, in the simple case of *static* synapse models, the computational and memory demands are relatively low. Since the network topology is generated at the start of the simulation, the states of all static synapses can be copied to device memory and remain there over the course of the entire simulation, reducing the overall amount of data transfer between the host and the device. Note that this also applies to the neuron data, as illustrated in Figure 5. The graph is stored in compressed sparse row (CSR) format, connections being grouped by their source node. Each time we perform the spike delivery, the list of firing nodes is copied to the device. The CSR format allows us to quickly determine the outgoing connections from the source node indexes and perform the delivery. Because of the high degree of connectivity, on an NVIDIA GPU, we assign one entire warp (32 threads) to process the outgoing connections of each source node. In *static* synapse models, the order of transferred spikes can be ignored, since their weights are constant. Therefore, we can deliver spikes with different timestamps in parallel. At the target node, atomic operations are utilized to avoid race conditions with respect to multiple spikes received by a target node at the same time.

Since *STDP* synapses consume significantly more memory than *static* models, storing all the *STDP* synapses in device memory would severely limit the feasible network size. To achieve scalability, we instead carry out the spike delivery through *STDP* synapses one batch of spikes at a time. On a GPU, one thread is assigned to one spike. The spikes traveling through the same synapse must be processed sequentially in timestamp order. To avoid additional processing, we group only spikes with the same timestamp in each batch. When a batch of spike is processed, the states of *STDP* synapses

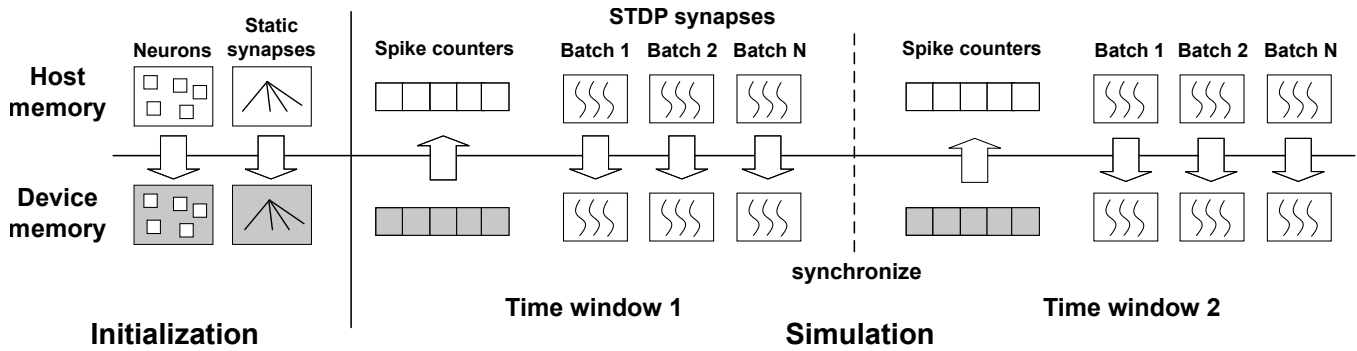


Figure 5: Data transfer between host and device memory in the initialization step and during simulation. The shaded shapes illustrate memory allocated permanently on the device.

transmitting these spikes are transferred to device memory. After a batch has been processed, the memory allocated for the STDP synapse data is freed (cf. Figure 5). The efficiency of this approach depends on the batch size, which is limited by the available graphics memory. The corresponding device kernels are generated automatically according to the same procedure as applied to the neuron models (cf. Figure 3).

5 HETEROGENEOUS NEST

We demonstrate our transition approach on the example of the well-known SNN simulator NEST [17]. NEST is an SNN simulator supporting various neuron and synapse models. In the following, we briefly sketch NEST’s architecture and describe our modifications as well as the automatic transformation of neuron models.

The process to enable automatic transformation of synapse models follows the same principles as the transformation of neuron models. Thus, we leave the implementation of the automatic transformation for synapses to future work. Our publicly available implementation currently includes manually ported versions of the synapse models required for the experiments in Section 6.

5.1 NEST overview

NEST is implemented using C++ and supports CPUs in shared and distributed-memory environments using OpenMP and MPI. The neural network is divided into a number of OpenMP threads referred to as *virtual processes* (VPs). Nodes are assigned to VPs in a round-robin fashion. A synapse is stored in the same VP as the target neuron, which guarantees thread-safety for the spike delivery, since only one thread can write data to a given neuron. The VPs can be assigned to multiple processes communicating via MPI in a distributed-memory environment.

A super step in NEST involves the following operations:

- (1) Each VP executes a super step on its assigned neurons, which involves executing the state update function of each neuron according to the selected model. Since the super step is divided into multiple time steps, the internal state of each neuron is updated multiple times, potentially creating new spikes at each iteration. If an emitted spike has a target node in a remote MPI process, it is stored in a local buffer.

- (2) The stored spikes are inserted into an MPI buffer and issued using an MPI_Allgather call to provide all processes with the spike data.
- (3) Each VP scans through its MPI receive buffer to obtain the list of source neurons of incoming spikes. If the VP holds one or more target neurons, it delivers the spikes through the corresponding synapses according to the synapse model.

The division of work across VPs makes it convenient to extend NEST with support for heterogeneous hardware: since each VP controls a device, multiple VPs can easily exploit the resources of a multi-device system. Further, the support for MPI communication enables a multi-device execution across multiple execution nodes without further development efforts.

The core component of NEST is the *SimulationManager*, which executes the operations of a super step as described above. For spike delivery, the *SimulationManager* relies on the *EventDeliveryManager*, which delivers spikes within and across VPs. The *ConnectionManager* stores the network topology of the current VP.

The above components constitute the core of the framework and do not require modification when varying the neuron or synapse models across simulation runs. Thus, since the code changes required to enable support for heterogeneous hardware are permanent and independent of the chosen models being simulated, we carry out these changes manually.

The main remaining components of the NEST framework are the neuron and synapse models. NEST 2.14 provides more than 50 neuron models and 10 synapse models. To perform a simulation run, the user prepares a script that describes the scenario, which is defined by the neuron and synapse models with their parameters, the number of neurons and connections, the simulated time, etc. Based on the script, NEST constructs the network with the selected models and executes the simulation. These models are defined as C++ classes and share a common interface so that the core components, particularly the *SimulationManager*, can invoke the model behavior (e.g., as defined by the state update function) in a generic fashion. Using the techniques described in Sections 4.2 and 4.3, we developed a tool that is supplied with a model implementation in the form of C++ code as input and generates an OpenCL implementation executable on OpenCL-supported accelerators. The compilation workflow follows the description in Section 4.2

Listing 1: Original C++ code

```

1  if ( S_.r_ == 0 ) {
2      S_.y3_ = V_.P30_ * ( S_.y0_ +
3          P_.I_e_ ) + ... ;
4
5      S_.y3_ = ( S_.y3_ < P_.LowerBound_ ?
6          P_.LowerBound_ : S_.y3_ );
7  } else {
8      --S_.r_;
9  }
10 S_.I_ex_ = V_.P21_ex_ * S_.dI_ex_ +
11     ...;
12 ...
13 // Ring buffer
14 V_.weighted_spikes_ex_ = B_.ex_spikes_.get_value(
15     lag );
16 ...
17 // Ring buffer
18 V_.weighted_spikes_in_ = B_.in_spikes_.get_value(
19     lag );
20 ...
21 ...
22 ...
23 if ( S_.y3_ >= P_.Theta_ ) {
24     S_.r_ = V_.RefractoryCounts_;
25     S_.y3_ = P_.V_reset_;
26     // Firing spikes
27     set_spike_time( Time::step( origin.get_steps() +
28         lag + 1 ) );
29     SpikeEvent se;
30     kernel().event_deliv_mgr.send( *this, se, lag );
31 }
32 // Ring buffer
33 S_.y0_ = B_.currents_.get_value( lag );
34 ...
35 ...
36 ...

```

Listing 2: Generated OpenCL code

```

1  if ( S__r__[tid] == 0 ) {
2      S__y3__[tid] = V__P30__[tid] * ( S__y0__[tid] +
3          P__I_e__[tid] ) + ...;
4
5      S__y3__[tid] = ( S__y3__[tid] < P__LowerBound__[tid] ?
6          P__LowerBound__[tid] : S__y3__[tid] );
7  } else {
8      --S__r__[tid];
9  }
10 S__I_ex__[tid] = V__P21_ex__[tid] * S__dI_ex__[tid] +
11     ...;
12 ...
13 // Ring buffer
14 V__weighted_spikes_ex__[tid] = ring_buffer_get_value(
15     ex_spikes_, ring_buffer_size, num_nodes, tid,
16     lag, time_index);
17 ...
18 // Ring buffer
19 V__weighted_spikes_in__[tid] = ring_buffer_get_value(
20     in_spikes_, ring_buffer_size, num_nodes, tid,
21     lag, time_index);
22 ...
23 if ( S__y3__[tid] >= P__Theta__[tid] ) {
24     S__r__[tid] = V__RefractoryCounts__[tid];
25     S__y3__[tid] = P__V_reset__[tid];
26     // Firing spikes
27     spike_count[tid]++;
28
29 }
30
31 // Ring buffer
32 S__y0__[tid] = ring_buffer_get_value(currents_,
33     ring_buffer_size, num_nodes, tid,
34     lag, time_index);
35 ...
36 ...

```

Figure 6: Update function of the *iaf_psc_alpha* neuron model in the C++ and OpenCL implementation

5.2 Transformation of neuron models

Each neuron model defines the neuron’s behavior when interacting with other neurons through incoming and outgoing spikes. In NEST, the behavior is implemented in the Update function of the neuron class. The model further defines the internal state variables, which are modified in the Update function. Algorithm 1 shows pseudo code of the steps performed by the Update function: using the current state of the neuron and the incoming spikes from other neurons as input, the Update function executes state changes caused by the incoming spikes in a super step. The computation is divided into multiple small time steps (lines 2-3). In the loop, the neuron executes the procedure *ModelBehavior*, which is model-specific. At each time step, depending on the current state, the neuron may emit spikes to be transmitted to neighboring neurons. The *SimulationManager* invokes the *SimulationUpdate* function, which iterates through the neurons of the local VP and invokes the Update function on each of them. The above procedure may be executed by multiple VPs on disjunct sets of neurons concurrently.

To enable execution of the Update function on heterogeneous hardware, the computation scheme is modified as shown in Algorithm 2. Now, the *SimulationManager* invokes *MassUpdate*, which processes all neurons of the given type. In *MassUpdate*, if the neuron data has not been initialized on the device, we copy the neuron

data from host to device memory (lines 2-3). Hence, the data is transferred to the device only once. Then, at each time step, all neurons are updated in parallel on the device. During the update, neurons may emit spikes. Thus, we invoke *CopyDataDtoH* (line 6), which copies counters from the device to the host that indicate to the host the number of spikes generated by each neuron. The number of counters is equal to the number of neurons. Currently, we represent the counters as 4-byte variables. However, since in practice the values are small, depending on the scenario it may be sufficient to represent them using even fewer bytes. In our experiments, this data did not incur substantial data transfers overhead.

An important step in the transformation is the generation of a device kernel for *ModelBehaviorGPU*. To this end, we implemented a parsing tool based on *LibTooling* from the *CLang/LLVM* toolkit³ that parses the C++ code of the *ModelBehavior* function of the neuron model class. The parsing tool allows us to obtain the following required pieces of information:

- (1) The list of state variables used by the *ModelBehavior* function to identify the variables to be transferred to the device.
- (2) The type and dimensionality of each variable. This information is used to allocate device memory.

³<https://clang.llvm.org/docs/LibTooling.html>

- (3) Function calls performed during the update. The update function may invoke utility functions to 1. retrieve incoming spikes from a neuron’s local storage and 2. emit spikes.

Listing 1 of Figure 6 shows an example of the C++ code of the *ModelBehavior* routine for the *iaf_psc_alpha* model. This code executes a number of instructions on the neuron’s state variables such as *S_r_*, *S_y3_*, *S_I_ex_*. These variables are of primitive types (**int** or **double**), often as members of a C++ struct. The parsing tool gathers the type information from the class header file. Further, the accesses to the ring buffer holding incoming spikes (lines 14, 19, and 33) and the transmission of spikes (lines 27-30) are detected as function calls to *get_value* and *event_delivery_manager.send*.

Listing 2 of Figure 6 shows the generated OpenCL kernel code. All accesses to neuron state variables have been transformed from an array-of-struct (AoS) to a GPU-friendly struct-of-array (SoA) pattern. For instance, the variable *S_r_* is of type **int**, where *S_* is an instance of the *State_* structure defined in the neuron class. In the OpenCL implementation, we define an **int** array *S__r_* of size identical to the number of neurons. Now, the device thread with index *tid* accesses the *S__r_[tid]* entry in the array. The utility function for accessing the ring buffer of incoming spikes is replaced with a kernel call. For sending spikes, we maintain a counter *spike_count* of the number of spikes sent by each neuron. The code portion for triggering spikes in neuron *tid* is replaced by an incrementation of *spike_count[tid]*. This counter is later copied back to the host (Algorithm 2, line 6) so that the sending of spikes can rely on the existing facilities of NEST (Listing 1, lines 27-30).

6 EVALUATION

The experiments are performed with respect to the *hpc_benchmark* scenario used in several existing works [21, 28, 29, 32, 41]. The benchmark relies on the *iaf_psc_alpha* model for all neurons. The indegree of each neuron is fixed at 11 250. The total number of neurons in the network is controlled by a scale factor: $N = \text{scale} \times 11\,250$. There are two types of neurons in the network: *excitatory* and *inhibitory* neurons. The synapses connecting excitatory neurons rely on the *stdp_pl_synapse_hom* model, while all other synapses use the *static_synapse* model. The time step size is 0.1ms. The synaptic delay is 1.5ms. The simulation terminates after 250ms.

We executed the simulations on a system comprised of two nodes using Intel Xeon Gold 6148 2.4G CPUs with 20 physical cores and 40 threads each, and four NVIDIA Tesla V100 SXM2 GPUs with 16 GiB of RAM. The tool chain is comprised of GCC 7.2.0, CUDA toolkit 9.0, and OpenMPI 1.10.7. The GPU implementation is developed from version 2.14 of the NEST simulator. However, the CPU results are obtained using the development branch of NEST 2.14, which in our experiment provided better performance than NEST 2.14.

The correctness of our implementation was verified by comparing the average spiking rate per neuron between CPU and GPU runs for a number of scenarios with different network size and simulation time. Since a random number generator (*poisson_generator*) is associated with each virtual process, CPU and GPU runs were compared using the same number of CPU threads and GPU streams respectively. In all tested cases, the results produced by the CPU and GPU variants were identical.




Figure 7: Speedup with one stream per GPU compared to execution on a single CPU core.

6.1 Performance

In the first experiment, we measure the speedup obtained by our GPU version. Figure 7 shows the throughput improvement using a single and multiple GPUs compared to single-threaded CPU execution at different SNN scales. In the CPU version, we run the NEST 2.14 with only 1 virtual process, i.e., in a single thread. Our GPU program runs in a single process with one OpenMP thread controlling each GPU. This setup does not involve MPI communication. We observe that in most cases, the throughput scales almost linearly with the number of GPUs. For example, at a network size of 56 250 neurons, the speedup of 2, 3, and 4 GPUs is 6.0, 9.6, and 12.6 over a single CPU; or 1.7, 2.8, and 3.7 over a single GPU. The results also show that the GPU implementation achieves better performance as the network size increases. At 22 500 neurons, the speedup factors over a single CPU core using a single and 4 GPUs are 1.92 and 8.10. Meanwhile, at 56 250 neurons, these ratios are 2.86 and 12.58. The reason for this result is the improved utilization of the GPU at larger network sizes.

In the previous experiment, each GPU was controlled by a single thread. To better exploit the GPU’s computing resources, we vary the number of CPU threads (or GPU streams) sharing a single GPU. We compare the performance of multi-stream single-GPU runs with the performance of a single-stream single-GPU execution. The results are shown in Figure 8. Note that, since the remaining CPU portion of the simulator still performs minor amounts of work, e.g., for data serialization, larger numbers of threads also reduce the time spent on the CPU portion. Further, the performance is improved by overlapping CPU-GPU data transfers in one thread with kernel computation in other threads. Generally, by increasing the number of threads (GPU streams) sharing a single GPU, the performance improves. However, at 6 threads, the performance starts decreasing slightly. A likely cause is the overhead of the additional kernel calls in relation to the smaller amount of computation per call.

We also compare the performance between multi-core CPU and GPU execution. Figure 9 shows the speedup over single-threaded CPU execution at different network sizes. Overall, the performance gain yielded on the multi-core CPU is consistent across network sizes. For example, a run using 4 threads is about 2.7 to 3.7 times faster than a single-threaded run. Meanwhile, as mentioned above, the GPU performance increases with the network size. When using a single stream to control the GPU, the GPU performance is comparable to 2 to 3 CPU cores. When relying on 5 threads to control the GPU, substantially higher performance is achieved: even at 22 500,




Figure 8: Speedup when executing multiple streams on a single GPU over a single stream.






Figure 9: Performance comparison between a single multi-core CPU and a single GPU. The results are given as speedup factors over one CPU core. With multiple GPU streams, the GPU achieves performance equivalent to about 9 CPU cores.

the GPU outperforms 5 CPU cores. At 56 250 neurons, the speedup over a single CPU core is 9.4, or 2.4 times faster than a run using 5 CPU cores.


We also evaluate the effect of the MPI communication on the performance of our NEST GPU implementation. Instead of executing multiple threads per process as in the previous experiments, we execute multiple MPI processes, each of which is single-threaded and uses one GPU. To be able to clearly quantify the proportion of runtime spent on different phases of the simulation, each process controls the GPU using only one stream, although we have previously seen that employing multiple threads per GPU increases the performance dramatically. Inter-node communication relies on the existing MPI communicating scheme of NEST 2.14. We measure the running time of each of the following: *update*, *spike delivery* and *MPI communication*. The results are illustrated in Figure 10. It is evident that the running time is dominated by the *spike delivery*. In our experiments, the proportion of time spent on *spike delivery* is about 81% to 87%. Compared with the original NEST implementation on a single CPU, our GPU method reduces the absolute running time of the *spike delivery* by a factor of about 2. The results show the importance of the performance of the *spike delivery* for the overall running time. The *update* step accounts for about 10% of the running time, while *MPI communication* accounts for 2% to 8%. From this experiment, we conclude that in our small-scale setup, MPI communication does not hinder performance increases through the use of GPUs. In future work, experiments in large-scale supercomputing environments could provide further scalability results.



(a) 2 MPI processes



(b) 3 MPI processes



(c) 4 MPI processes

Figure 10: Breakdown of the running time with different numbers of MPI processes on a single execution node, each using one GPU.

Finally, we conducted preliminary experiments of a CPU-GPU co-execution, i.e., neurons and synapses are divided into partitions, each of which is processed either by the original CPU-based NEST or our OpenCL implementation running on a GPU. We used up to 6 CPU cores and up to 6 streams on a single GPU to simulate the benchmark network at a scale of 67 500 neurons. The main challenge lies in balancing the workload between the CPU and the GPU, since the GPU can process each time step substantially faster than a CPU core. Currently, we are relying on an even assignment of the same number of neurons to each CPU core or GPU stream. With this simple assignment, the best performance was achieved when the entire workload is processed by the GPU. To achieve a performance benefit with the co-execution, the number of neurons assigned to a CPU core or GPU stream should thus be chosen separately, which is part of our future work.

6.2 Memory consumption

Since accelerators such as GPUs are typically equipped with lower amounts of memory than is accessible to the host CPU, efficient use of the available memory is critical to be able to process large segments of the network concurrently. Due to the large degree of connectivity in SNNs, the synapse data typically constitutes the largest part of the memory consumption of an SNN simulation. For example, in a network of 22 500 neurons, the synapse data accounts

for 97% of the total memory consumptions. In Section 4.3, we presented our strategy to efficiently represent the network in memory. The memory consumption of the benchmark from our performance measurements is roughly proportional to the number of neurons in the network. A simulation of a network of 11 250 neurons required 3.7 GiB of host memory and 0.8 GiB of graphics memory. At 56 250 neurons, about 18 GiB of host memory and 2.9 GiB of graphics memory are required. Our experiments consumed about a factor of 4.6 to 6.2 more host memory than graphics memory.

7 DISCUSSION

One of the goals of our work is to support heterogeneous hardware environments. Compared to other GPU-accelerated simulators such as Nemo [15], CARLsim [10], GENN [48], NCS [25], and HRL-Sim [40], which use NVIDIA's CUDA platform, our implementation relies on OpenCL, allowing the simulation to make use of CPUs as well as a variety of accelerators (e.g., GPUs, FPGAs, DSPs, Intel Xeon Phi). The OpenCL code generated by the source-to-source transformation is hardware-agnostic, so that executable code can be compiled for any platform supported by OpenCL. Future work could explore the performance achievable on platforms beyond the CPU and GPU environments of the experiments in the present paper. By extending the transformation tool with hardware-specific optimizations, it may be possible to achieve higher performance on certain devices while still allowing for models to be formulated without regards for hardware specifics.

The main factor limiting the maximum size of the simulated network is the available memory on the accelerator. Existing work on GPU-accelerated SNN simulations rarely focused on scalability. Often, the considered networks contained only 100 to 1,000 connections per neuron so that the entire network could be stored in GPU memory. Thus, the applicability of these simulators to the large-scale networks frequently encountered in practice is limited. By batching the computational work (cf. Section 4.3), the accelerator memory consumption in our implementation is largely decoupled from the network size. In our experiments, we executed the *hpc_benchmark* scenarios in configurations requiring up to 18 GiB of GPU memory per execution node.

Our implementation targets NEST version 2.14. More recently, NEST 2.16 [36] has been released with improvements in network build time and scalability over the NEST 2.14 [29] and comparable performance in small-scale and medium-scale systems. Extending our work to NEST 2.16 could allow us to support large networks at even higher performance. Further, our transformation approach could be applied to established CPU-based simulators such as NEURON [24] and PCSIM [43].

Since recent supercomputers frequently rely on combinations of CPU and GPU devices, a CPU-GPU co-execution should be used to make full use of such environments. In our performance measurements, we have seen that while a CPU-GPU co-execution is already possible, a mechanism for workload balancing among these devices should be introduced to achieve high performance. Moreover, even when relying purely on a single accelerator device, the number of threads controlling the device must be chosen carefully to exploit the hardware fully. Thus, future work should consider automated workload balancing, e.g., using runtime profiling and autotuning.

Finally, the OpenCL-accelerated NEST stores neuron state data permanently in accelerator memory. To support the data collection requirements of real-world studies, it may be necessary to periodically transfer parts of this data to host memory. Support for such transfers can easily be added to our implementation. The effects on performance depend on the amount of data required and the frequency of data collection. In contrast to neuron and synapse state data, spikes are regularly transferred to the host. Thus, our current implementation already allows for data collection required to explore spiking patterns over time. In future work, it would be interesting to explore methods for *in situ* data analysis in accelerator memory, transferring only the analysis results to the host.

8 CONCLUSIONS

We presented an approach to transform CPU-based spiking neural network simulators to enable their execution on heterogeneous hardware. Our approach relies on manual porting of static core simulator functionalities, whereas neuron model code is analyzed and transformed automatically. We demonstrated our approach by transforming the well-known NEST simulator to support OpenCL, enabling its acceleration using hardware platforms such as GPUs, FPGAs, and DSPs. Since the transformed code supports co-execution on multiple device types, better hardware utilization and lower runtimes can be achieved on modern GPU-dominated supercomputers. Our performance measurements show that at sufficient utilization, a single GPU achieves the performance of about 9 CPU cores. An important avenue for future work is the workload balancing between CPU and GPU during co-execution. Applying the transformation to the recent NEST 2.16 may provide further performance increases. Finally, the OpenCL code generated by our code transformation can also exploit accelerators such as FPGAs and DSPs, which may further accelerate certain experiments. We hope the community will benefit from our publicly available implementation.

ACKNOWLEDGMENT

This work was financially supported by the Singapore National Research Foundation under its Campus for Research Excellence And Technological Enterprise (CREATE) programme.

REFERENCES

- [1] Rajagopal Ananthanarayanan, Steven K Esser, Horst D Simon, and Dharmendra S Modha. 2009. The Cat is out of the Bag: Cortical Simulations with 10^9 Neurons, 10^{13} Synapses. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. ACM, 63.
- [2] Philipp Andelfinger and Hannes Hartenstein. 2014. Exploiting the Parallelism of Large-Scale Application-Layer Networks by Adaptive GPU-based Simulation. In *Winter Simulation Conference (WSC)*. IEEE, 3471–3482.
- [3] Philipp Andelfinger, Jens Mittag, and Hannes Hartenstein. 2011. GPU-based Architectures and Their Benefit for Accurate and Efficient Wireless Network Simulations. In *International Symposium on Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS)*. IEEE, 421–424.
- [4] Soufiane Baghdadli, Armin Gröflinger, and Albert Cohen. 2010. Putting Automatic Polyhedral Compilation for GPGPU to Work. In *Workshop on Compilers for Parallel Computers (CPC)*.
- [5] Muthu Manikandan Baskaran, Jj Ramanujam, and P Sadayappan. 2010. Automatic C-to-CUDA Code Generation for Affine Programs. In *International Conference on Compiler Construction*. Springer, 244–263.
- [6] Trevor Bekolay, James Bergstra, Eric Hunsberger, Travis DeWolf, Terrence C Stewart, Daniel Rasmussen, Xuan Choo, Aaron Voelker, and Chris Eliasmith. 2014. Nengo: a Python Tool for Building Large-Scale Functional Brain Models. *Frontiers in Neuroinformatics* 7 (2014), 48.

- [7] Mohammad A Bhuiyan, Vivek K Pallipuram, Melissa C Smith, Tarek Taha, and Rommel Jalasutram. 2010. Acceleration of Spiking Neural Networks in Emerging Multi-Core and GPU Architectures. In *International Symposium on Parallel & Distributed Processing, Workshops and PhD Forum (IPDPSW)*. IEEE, 1–8.
- [8] Zhenshan Bing, Claus Meschede, Florian Röhrbein, Kai Huang, and Alois C Knoll. 2018. A Survey of Robotics Control Based on Learning-Inspired Spiking Neural Networks. *Frontiers in Neuroinformatics* 12 (2018).
- [9] Uday Bondhugula, Vinayaka Bandishti, Albert Cohen, Guillaín Potron, and Nicolas Vasilache. 2014. Tiling and Optimizing Time-Iterated Computations over Periodic Domains. In *International Conference on Parallel Architecture and Compilation Techniques (PACT)*. IEEE, 39–50.
- [10] Ting-Shuo Chou, Hirak J Kashyap, Jinwei Xing, Stanislav Listopad, Emily L Rounds, Michael Beyeler, Nikil Dutt, and Jeffrey L Krichmar. 2018. CARLsim 4: An Open Source Library for Large Scale, Biologically Detailed Spiking Neural Network Simulation using Heterogeneous Clusters. In *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [11] Biagio Cosenza, Nikita Popov, Ben Juurlink, Paul Richmond, Mozghan Kabiri Chimeh, Carmine Spagnuolo, Gennaro Cordasco, and Vittorio Scaranò. 2018. OpenABL: A Domain-Specific Language for Parallel and Distributed Agent-Based Simulations. In *European Conference on Parallel Processing*. Springer, 505–518.
- [12] Zachary DeVito, Niels Joubert, Francisco Palacios, Stephen Oakley, Montserrat Medina, Mike Barrientos, Erich Elsen, Frank Ham, Alex Aiken, Karthik Duraisamy, et al. 2011. Liszt: a Domain Specific Language for Building Portable Mesh-based PDE Solvers. In *International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 1–12.
- [13] Peng Di, Hui Wu, Jingling Xue, Feng Wang, and Canqun Yang. 2012. Parallelizing SOR for GPGPUs using Alternate Loop Tiling. *Parallel Comput.* 38, 6-7 (2012), 310–328.
- [14] Christian Feichtinger, Johannes Habich, Harald Köstler, Ulrich Rüdè, and Takayuki Aoki. 2015. Performance Modeling and Analysis of Heterogeneous Lattice Boltzmann Simulations on CPU–GPU Clusters. *Parallel Comput.* 46 (2015), 1–13.
- [15] Andreas K Fijelnd, Etienne B Roesch, Murray P Shanahan, and Wayne Luk. 2009. NeMo: a Platform for Neural Modelling of Spiking Neurons using GPUs. In *International Conference on Application-specific Systems, Architectures and Processors*. IEEE, 137–144.
- [16] Enrico Franchi. 2012. A Domain Specific Language Approach for Agent-based Social Network Modeling. In *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*. IEEE Computer Society, 607–612.
- [17] Marc-Oliver Gewaltig and Markus Diesmann. 2007. NEST (Neural Simulation Tool). *Scholarpedia* 2, 4 (2007), 1430.
- [18] Dan FM Goodman and Romain Brette. 2009. The Brian Simulator. *Frontiers in Neuroscience* 3 (2009), 26.
- [19] Qingfeng Guan, Xuan Shi, Miaoqing Huang, and Chenggang Lai. 2016. A Hybrid Parallel Cellular Automata Model for Urban Growth Simulation over GPU/CPU Heterogeneous Architectures. *International Journal of Geographical Information Science* 30, 3 (2016), 494–514.
- [20] KA Hawick and DP Playne. 2013. Simulation Software Generation using a Domain-Specific Language for Partial Differential Field Equations. In *International Conference on Software Engineering Research and Practice (SERP)*. The Steering Committee of The World Congress in Computer Science, Computer Engineering and Applied Computing (WorldComp), 69–75.
- [21] Moritz Helias, Susanne Kunkel, Gen Masumoto, Jun Igarashi, Jochen Martin Eppler, Shin Ishii, Tomoki Fukai, Abigail Morrison, and Markus Diesmann. 2012. Supercomputers Ready for use as Discovery Machines for Neuroscience. *Frontiers in Neuroinformatics* 6 (2012), 26.
- [22] Suzanaerculano-Houzel. 2012. The Remarkable, yet not Extraordinary, Human Brain as a Scaled-up Primate Brain and its Associated Cost. *Proceedings of the National Academy of Sciences* 109, Supplement 1 (2012), 10661–10668.
- [23] Suzanaerculano-Houzel and Jon H Kaas. 2011. Gorilla and Orangutan Brains Conform to the Primate Cellular Scaling Rules: Implications for Human Evolution. *Brain, Behavior and Evolution* 77, 1 (2011), 33–44.
- [24] Michael L Hines and Nicholas T Carnevale. 1997. The NEURON Simulation Environment. *Neural Computation* 9, 6 (1997), 1179–1209.
- [25] Roger V Hoang, Devyani Tanna, Laurence C Jayet Bray, Sergiu M Dascalu, and Frederick C Harris Jr. 2013. A Novel CPU/GPU Simulation Environment for Large-Scale Biologically Realistic Neural Modeling. *Frontiers in Neuroinformatics* 7 (2013), 19.
- [26] Sungpack Hong, Hassan Chafi, Edic Sedlar, and Kunle Olukotun. 2012. GreenMarl: a DSL for Easy and Efficient Graph Analysis. In *ACM SIGARCH Computer Architecture News*, Vol. 40. ACM, 349–362.
- [27] Kaixi Hou, Hao Wang, Wu-chun Feng, Jeffrey S Vetter, and Seyong Lee. 2018. Highly Efficient Compensation-based Parallelism for Wavefront Loops on GPUs. In *International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 276–285.
- [28] Tammo Ippen, Jochen M Eppler, Hans E Plesser, and Markus Diesmann. 2017. Constructing Neuronal Network Models in Massively Parallel Environments. *Frontiers in Neuroinformatics* 11 (2017), 30.
- [29] Jakob Jordan, Tammo Ippen, Moritz Helias, Itaru Kitayama, Mitsuhsa Sato, Jun Igarashi, Markus Diesmann, and Susanne Kunkel. 2018. Extremely Scalable Spiking Neuronal Network Simulation Code: From Laptops to Exascale Computers. *Frontiers in Neuroinformatics* 12 (2018), 2.
- [30] Gordon L Kindlmann, Charisee Chiu, Nicholas Seltzer, Lamont Samuels, and John H Reppy. 2016. Diderot: a Domain-Specific Language for Portable Parallel Scientific Visualization and Image Analysis. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 867–876.
- [31] Alois Knoll and Marc-Oliver Gewaltig. 2016. Neuroinformatics: a Strategic Pillar of the Human Brain Project. *Science Robotics* (2016).
- [32] Susanne Kunkel, Maximilian Schmidt, Jochen M Eppler, Hans E Plesser, Gen Masumoto, Jun Igarashi, Shin Ishii, Tomoki Fukai, Abigail Morrison, Markus Diesmann, et al. 2014. Spiking Network Simulation Code for Petascale Computers. *Frontiers in Neuroinformatics* 8 (2014), 78.
- [33] Christian Lengauer. 1993. Loop Parallelization in the Polytope Model. In *International Conference on Concurrency Theory*. Springer, 398–416.
- [34] Da Li, Hancheng Wu, and Michela Becchi. 2015. Nested Parallelism on GPU: Exploring Parallelization Templates for Irregular Loops and Recursive Computations. In *International Conference on Parallel Processing (ICPP)*. IEEE, 979–988.
- [35] Peilong Li, Yan Luo, Ning Zhang, and Yu Cao. 2015. Heterospark: A Heterogeneous CPU/GPU Spark Platform for Machine Learning Algorithms. In *International Conference on Networking, Architecture and Storage (NAS)*. IEEE, 347–348.
- [36] Charl Linssen, Mikkel Elle Lepperød, Jessica Mitchell, Jari Pronold, Jochen Martin Eppler, Chrisitan Keup, Alexander Peyser, Susanne Kunkel, Philipp Weidel, Yannick Nodem, et al. 2018. NEST 2.16.0. (Aug 2018).
- [37] Wolfgang Maass. 1997. Networks of Spiking Neurons: the Third Generation of Neural Network Models. *Neural Networks* 10, 9 (1997), 1659–1671.
- [38] Liam P Maguire, T Martin McGinnity, Brendan Glackin, Arfan Ghani, Ammar Belatreche, and Jim Harkin. 2007. Challenges for Large-Scale Implementations of Spiking Neural Networks on FPGAs. *Neurocomputing* 71, 1-3 (2007), 13–29.
- [39] Richard Membarth, Oliver Reiche, Frank Hannig, Jürgen Teich, Mario Körner, and Wieland Eckert. 2016. HIPA^{cc}: A Domain-Specific Language and Compiler for Image Processing. *IEEE Transactions on Parallel and Distributed Systems* 27, 1 (2016), 210–224.
- [40] Kirill Minkovich, Corey M Thibeault, Michael John O'Brien, Aleksey Nogin, Youngkwan Cho, and Narayan Srinivasa. 2014. HRLSim: a High Performance Spiking Neural Network Simulator for GPGPU Clusters. *IEEE Transactions on Neural Networks and Learning Systems* 25, 2 (2014), 316–331.
- [41] Abigail Morrison, Ad Aertsen, and Markus Diesmann. 2007. Spike-Timing-Dependent Plasticity in Balanced Random Networks. *Neural Computation* 19, 6 (2007), 1437–1467.
- [42] David M Nicol. 1993. The Cost of Conservative Synchronization in Parallel Discrete Event Simulations. *Journal of the ACM (JACM)* 40, 2 (1993), 304–333.
- [43] Dejan Pecevski, Thomas Natschläger, and Klaus Schuch. 2009. PCSIM: a Parallel Simulation Environment for Neural Circuits Fully Integrated with Python. *Frontiers in Neuroinformatics* 3 (2009), 11.
- [44] Jan-Phillip Tiesel and Anthony S Maida. 2009. Using Parallel GPU Architecture for Simulation of Planar I/F Networks. In *International Joint Conference on Neural Networks*. IEEE, 3118–3123.
- [45] Sven Verdoolaege, Juan Carlos Juega, Albert Cohen, Jose Ignacio Gomez, Christian Tenllado, and Francky Catthoor. 2013. Polyhedral Parallel Code Generation for CUDA. *ACM Transactions on Architecture and Code Optimization (TACO)* 9, 4 (2013), 54.
- [46] Jiajian Xiao, Philipp Andelfinger, David Eckhoff, Wentong Cai, and Alois Knoll. 2018. Exploring Execution Schemes for Agent-Based Traffic Simulation on Heterogeneous Hardware. In *Proceedings of the International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*. IEEE, 243–252.
- [47] Jiajian Xiao, Philipp Andelfinger, David Eckhoff, Wentong Cai, and Alois Knoll. 2019. A Survey on Agent-based Simulation using Hardware Accelerators. *ACM Computing Surveys (CSUR)* 52, 1 (Jan. 2019).
- [48] Esin Yavuz, James Turner, and Thomas Nowotny. 2016. GeNN: a Code Generation Framework for Accelerated Brain Simulations. *Scientific Reports* 6 (2016), 18854.