

A Survey on Agent-based Simulation using Hardware Accelerators

JIAJIAN XIAO, TUM Create Ltd.& Technische Universität München,
 PHILIPP ANDELFINGER, TUM Create Ltd. & Nanyang Technological University,
 DAVID ECKHOFF, TUM Create Ltd. & Technische Universität München,
 WENTONG CAI, Nanyang Technological University,
 ALOIS KNOLL, Technische Universität München,

Due to decelerating gains in single-core CPU performance, computationally expensive simulations are increasingly executed on highly parallel hardware platforms. Agent-based simulations, where simulated entities act with a certain degree of autonomy, frequently provide ample opportunities for parallelisation. Thus, a vast variety of approaches proposed in the literature demonstrated considerable performance gains using hardware platforms such as many-core CPUs and GPUs, merged CPU-GPU chips as well as FPGAs. Typically, a combination of techniques is required to achieve high performance for a given simulation model, putting substantial burden on modellers. To the best of our knowledge, no systematic overview of techniques for agent-based simulations on hardware accelerators has been given in the literature. To close this gap, we provide an overview and categorization of the literature according to the applied techniques. Since at the current state of research, challenges such as the partitioning of a model for execution on heterogeneous hardware are still a largely manual process, we sketch directions for future research towards automating the hardware mapping and execution. This survey targets modellers seeking an overview of suitable hardware platforms and execution techniques for a specific simulation model, as well as methodology researchers interested in potential research gaps requiring further exploration.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Computing methodologies** → **Distributed simulation**; *Agent / discrete models*; • **Hardware** → **Hardware accelerators**;

Additional Key Words and Phrases: Agent-based simulation, Heterogeneous Computing

ACM Reference Format:

Jiajian Xiao, Philipp Andelfinger, David Eckhoff, Wentong Cai, and Alois Knoll. 2018. A Survey on Agent-based Simulation using Hardware Accelerators. *ACM Comput. Surv.* X, Y, Article Z (March 2018), 38 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

Since around 2005, it can be observed that due to the breakdown of Dennard scaling, clock frequencies of single CPUs are no longer increasing significantly, even though the transistor counts are

Authors' addresses: Jiajian Xiao, TUM Create Ltd.& Technische Universität München, Area-Interlinking Design Analysis, Singapore, jiajian.xiao@tum-create.edu.sg; Philipp Andelfinger, TUM Create Ltd. & Nanyang Technological University, Area-Interlinking Design Analysis, Singapore, philipp.andelfinger@tum-create.edu.sg; David Eckhoff, TUM Create Ltd. & Technische Universität München, Area-Interlinking Design Analysis, Singapore, david.eckhoff@tum-create.edu.sg; Wentong Cai, Nanyang Technological University, School of Computer Science and Engineering, Singapore, david.eckhoff@tum-create.edu.sg; Alois Knoll, Technische Universität München, Department of Informatics, Munich, knoll@in.tum.de.

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.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery. 0360-0300/2018/3-ARTZ \$15.00
<https://doi.org/0000001.0000001>

still growing [152]. Instead, CPU manufacturers have more and more focused on developing multi-core processors. This in turn calls for parallel computing techniques, as programmes (including simulations) that cannot be run in parallel can no longer simply be sped up by incorporating a newer and faster CPU. Performance can be increased further when the workload of a programme is efficiently distributed to heterogeneous hardware such as Graphics Processing Units (GPUs) or Field Programmable Gate Arrays (FPGAs) [27].

Some types of hardware are better suited for certain tasks than others, for example, tasks with large amounts of fine-grained parallelism can benefit greatly from the massively parallel architecture of modern GPUs with its thousands of cores. Tasks that are largely sequential or characterised by unpredictable data accesses and control flow lend themselves better to CPUs with out-of-order execution, long pipelines and large caches. Similarly, if offloading a task to a GPU requires copying large amounts of data to and from graphics memory, execution on a CPU may be preferable even if substantial parallelism is available. This issue can be addressed by an Accelerated Processing Unit (APU), where CPU and an integrated graphics core (of lower performance compared to stand-alone GPUs) share the same memory. Lastly, compute-intensive and memory-light tasks can be outsourced to FPGAs which can be programmed to carry out specific computations in hardware.

One field that has always sought after more performance is the field of simulation. Faster computers allow an increase in complexity of the incorporated simulation models, allowing researchers to obtain more accurate results in a faster manner. Agent-based simulations have received broad attention as they can be employed to study various domains, such as road traffic [39], social networks [42], pedestrian movement [156], military [29], biology [4], economics [157] and so on. The main characteristic of agent-based simulation is that autonomous agents (e.g., individuals or entities) act and interact to create effects of emergence on the entire system. The complex decision-making of agents and the huge scale of many simulated systems can lead to enormous runtimes, motivating the need for employing high-performance computing platforms.

Agent-based simulations are a promising target for parallel computing techniques as agents are autonomous and in some cases carry out independent computations. In mobility simulations, interactions between agents usually only take place between close-by agents in a somewhat regular 2D or 3D environment, allowing researchers to employ space partitioning without inducing too much synchronisation overhead. Moreover, many ABS are time-stepped and agents are often updated at the same logical time, providing inherent independence and thus potentials for parallelised execution. Unfortunately, being able to partition a problem and execute it in parallel is not a guarantee that it can be accelerated using heterogeneous hardware.

To enable ABS on heterogeneous hardware, some general challenges have to be overcome. First, the simulation has to be partitioned with heterogeneity in mind to decide which part of the program lends itself best to a specific hardware device, considering the resulting overhead from data transfers between the different devices. From this it follows that depending on the used hardware, the mapping of simulation parts to hardware devices will likely be different. Complex simulations typically also exhibit scattered and unpredictable memory access and control flow as the model state develops dynamically over time. This further complicates an efficient distribution to heterogeneous hardware. Lastly, in order to make heterogeneous accelerators available to modellers even without having in-depth knowledge of the specific hardware platforms, there is a need for frameworks that abstract away from hardware specifics. Some of the common frameworks provide variants supporting parallel and distributed execution, e.g., MASON [101], Repast-HPC [32], EcoLab [149], and GridABM [154]. However, these frameworks only support traditional CPU-based environments. Some frameworks such as FLAME GPU [31] and MCMAS [92] have been proposed that focus on the execution on specific accelerators such as graphics cards.

In this survey, we structure the complex landscape of agent-based simulation on heterogeneous hardware. We give an overview of existing types of hardware that have been employed to accelerate agent-based simulations and discuss past developments and current trends. While some surveys exist that present generic high-performance computing techniques using heterogeneous hardware [43, 107, 160], we highlight the specific challenges of ABS on heterogeneous hardware and categorize an ample body of related work along these challenges. For each challenge, we discuss in detail how existing literature has contributed to solving them. This overview allows us to identify research gaps that need to be filled in order to establish heterogeneous accelerators in the simulation domain and making them applicable to a wider range of problems – ideally by providing an automated process to support the modeller.

The remainder of this survey is structured as follows: in Section 2, we characterise the main classes of hardware accelerators for general-purpose computations. Section 3 provides an overview of agent-based simulation concepts and outlines the computational challenges of executing agent-based simulations on hardware accelerators. In Section 4, we systematise and survey the existing works according to the identified challenges and according to the techniques used to do so. In Section 5, we discuss unresolved challenges and outline how a system tackling these challenges could look like, thus sketching avenues for future work. Section 6 summarizes our findings and concludes the survey.

2 HARDWARE PLATFORMS

In this section, we describe the technical characteristics, the benefits as well as the limitations of hardware platforms that have been used to accelerate agent-based simulations. We focus on many-core CPUs, Graphics Processing Units (GPUs), Accelerated Processing Units (APUs), and Field-Programmable Gate Arrays (FPGAs). Readers familiar with these hardware platforms may skip this section and continue to Section 3.

2.1 Many-Core CPUs

Architecture: A many-core (or many integrated core, MIC) CPU contains a group of CPU cores on a single chip. One of the well-known many-core CPUs, the Intel Xeon Phi, is equipped with up to 72 x86-compatible CPU cores communicating via an internal Network-on-Chip which enables fast and parallel data transfer between the cores.

A many-core CPU can be connected to the host machine via PCI-E or can be a standalone CPU with direct access to the system memory. Figure 1 shows an overview of the second generation Intel Xeon Phi 7290F (Knights Landing) processor with its 72 cores that are grouped into 36 tiles interconnected by a 2D mesh channel. Each tile has 2 cores sharing 1MB of L2 cache. All L2 caches

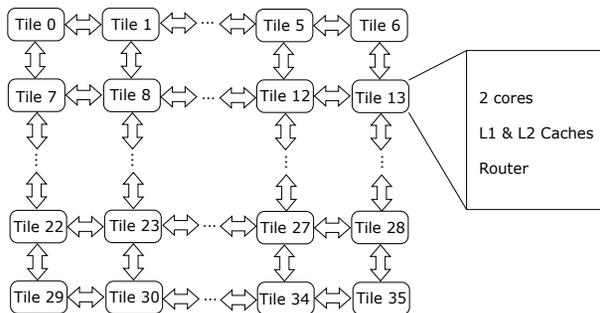


Fig. 1. The tile-based architecture of the Intel Xeon Phi 7290F Processor based on Knights Landing [147].

are kept fully coherent by a distributed tag directory. The processor supports a maximum of 384GB of DDR4 RAM. In addition, 16GB of 3D-stacked multi-channel DRAM can be used for transparent caching or managed manually.

In the past years, a number of non-x86 many-core CPUs have emerged, such as the Parallel Board [3], the Epiphany-V [119], and the Kalray MPPA (Massively Parallel Processor Array) series [36].

Benefits: A notable advantage of some many-core CPUs over GPUs and FPGAs is their capability to execute largely unmodified code written for regular CPUs [90]. This makes migration to these platforms easier, given the code is highly parallelisable. Since the individual cores support out-of-order execution, employ deep instruction pipelines, and have access to comparatively large caches, the need to adapt a program's control flow to the hardware is less pressing than with, e.g., GPUs [22]. Still, as some many-core processors support vector operations through instruction set extensions such as AVX-512 [76], a single-instruction, multiple-data (SIMD) style of programming can extract further parallelism.

Recent work showed that many-core CPUs can substantially accelerate Discrete Event Simulation (DES) [79, 167]. A number of authors have also evaluated the acceleration of various types of simulations such as fluid dynamics and seismic wave propagation using non-x86 many-cores [26, 131].

Limitations: In light of the comparatively high cost of recent many-core CPUs (\approx US\$3368.00 as of 03/2018 for an Intel Xeon Phi Processor 7290F) compared to other accelerators, the performance gains compared to traditional multi-core CPUs have frequently been relatively low. Even when optimising scientific code for a many-core CPU, there may only be a single-digit speedup over an execution on a traditional multi-core CPU, while in some cases there may even be an increase in runtime [12]. Further, since the performance depends strongly on parameters such as the number of threads and on employing the different types of memory available on a many-core CPU, it is necessary to tune these aspects to the given problem and hardware [99].

2.2 Graphics Processing Units (GPUs)

Architecture: GPUs utilise a massively parallel architecture, which makes them considerably more efficient than general purpose CPUs when large volumes of data can be processed in parallel. Their original purpose was to accelerate the processing of three-dimensional scenes to be displayed on two-dimensional screens. However, modern GPUs have evolved to support a wide range of computational tasks.

The evolution of GPUs (and with that their applicability for simulation) is characterised by three essential steps. In the 1990s, GPUs followed a fixed-function architecture, which processed a scene's geometry to produce the colour and transparency values for each of the screen's pixels in a pipelined fashion. In 2001, Nvidia released the GeForce 3, a new GPU generation which marks the second stage of GPU evolution. The GeForce 3 included so-called shader units, which execute programs applied to large numbers of pixel RGBA values or vertices of the objects in a 3D scene.

The flexibility of shader programming made the idea of General-Purpose Computing on Graphics Processing Units (GPGPU) practical, with early GPGPU work mapping raw data to pixels or vertices to achieve GPU-based parallel programming. Finally, in 2006, the shader architecture was unified by no longer distinguishing between vertex and pixel shaders. Now, with the Compute Unified Device Architecture (CUDA) programming framework, it became possible for GPUs to seamlessly perform general-purpose computational tasks [117, 133].

We sketch the GPU architecture and programming model on the basis of Nvidia's terminology. AMD hardware follows a similar design. A modern GPU consists of a scalable number of Streaming Multiprocessors (SMs), which contain a number of Streaming Processors (SPs) that perform most

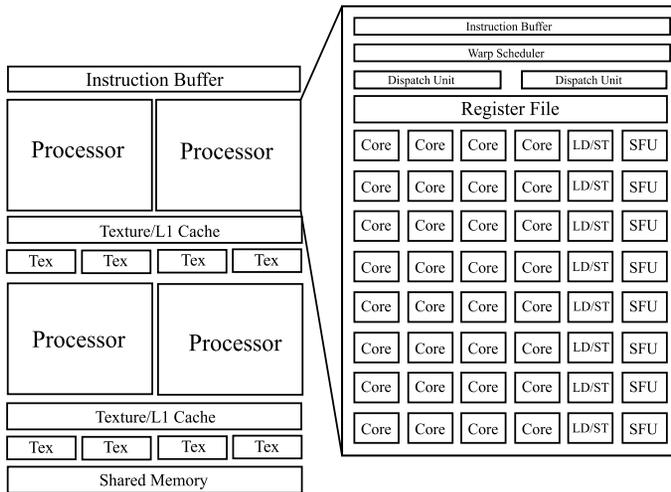


Fig. 2. A Streaming Multiprocessor (SM) in a GP105 GPU based on Nvidia's Pascal architecture [114].

of the computations, Special Function Units (SFUs) to efficiently perform special operations such as executing trigonometric functions, on-chip memory, cache and registers. In addition, modern GPUs have L2 cache and off-chip RAM, both of which are shared among all SMs [114]. Nvidia's GeForce GTX1080, for example, has 20 SMs, each of them containing 128 SPs, 32 SFUs, 256KB of registers, 8 texture units, 96KB of low-latency memory, and 48KB of L1 cache Figure 2 . There are 2 MB of L2 cache and 8 GB of off-chip RAM, referred to as global memory.

CUDA [117] (supporting Nvidia GPUs), OpenACC [120] and OpenCL [150] (the latter two supporting Nvidia, AMD GPUs, and Intel CPUs and integrated GPU) are common programming frameworks for GPGPU. Both CUDA and OpenCL follow a similar programming model, with some differences in terminology.

The work to be performed by a GPU program is organized in a hierarchical fashion, aligned with the properties of the underlying hardware: at the lowest level, there are threads representing a sequential control flow. On a logical level, all threads execute the same GPU program in parallel. Threads are grouped into *warps* of a hardware-specific size (32 threads on Nvidia hardware). Within each warp, threads execute in lockstep, i.e., if the control flow among threads within a warp diverges, the different branches are serialised. Thus, although the serialisation is transparent to the programmer, it is important to minimise intra-warp divergence. A configurable number of warps forms a block. Warps inside a block have access to a limited amount of low-latency shared memory and can synchronize efficiently. Blocks are assigned to an SM persistently, i.e., the required registers and shared memory are assigned to the block until all warps have finished execution. Per-SM warp schedulers dynamically assign runnable warps to the available SPs to minimise stalling on high-latency memory accesses. Typically, there are many more threads than physical SPs, providing ample opportunities for this type of memory latency hiding [117].

A key aspect when programming GPUs is the optimisation of memory access patterns. The GPU hardware prescribes certain rules according to which memory accesses multiple threads can be *coalesced*, i.e., executed in aggregate [117]. Generally, the number of physical memory transactions required is minimised when threads with adjacent logical indexes access adjacent locations in memory. Since many applications require scattered or even unpredictable memory access, achieving memory coalescing is a common focus of works in parallel programming on GPUs (e.g., [45, 169]).

The recent Nvidia Volta architecture provides individual threads with their own execution context, enabling a more fine-grained control over the intra-warp control flow [116].

Benefits: The hardware and programming model of GPUs lend itself well to problems that can be expressed so that large numbers of similar or identical operations are performed on different data. Commonly, GPUs accelerate fine-grained data-parallel tasks by one to two orders of magnitude compared to an implementation on multi-core CPUs.

Mature GPU programming frameworks such as CUDA, OpenCL and OpenACC enable relatively simple programming compared to other accelerators such as FPGAs [44]. Libraries such as Thrust [17] and CUBLAS [113] supply the programmer with high-performance parallel implementations for common tasks such as parallel reduction, sorting, and linear algebra operations. Programming frameworks are available even for more specialised tasks such as agent-based simulation [31].

In addition to the benefits of GPU in terms of computations, Richmond and Romano [136] emphasise the opportunities for efficient visualisation of simulation. Since the agent data is already stored in graphics memory, visualisation can be achieved easily by passing the agents data to vertex or texture buffers.

Limitations: Most current GPUs are connected to their host CPU via the PCI-E bus. Thus, the GPU does not have direct access to the host memory. Data transfer between CPU and GPU is expensive in terms of latency and should therefore be reduced as much as possible. For instance, according to its specification, a PCI-E 3.0 x16 connection allows an Nvidia Titan X card to transfer data between host and graphics memory at up to 16 GB/s, while the graphics memory of the card can achieve a throughput of 336.5 GB/s. However, on recent GPU architectures, interconnects such as Nvidia's NVLink [115] and AMD's Infinity Fabric [2] achieve throughputs of up to 300 GB/s, alleviating the impact of data transfers.

High performance on a GPU requires the given task to be expressed in a way that fits the GPU's hardware properties. The main requirements are a large degree of parallelism and the possibility to achieve coalesced memory access as well as a common control flow among the threads within a warp. Thus, memory-intensive tasks with complex data dependencies are typically difficult to execute efficiently on GPUs [25, 82].

Compared to many-core CPUs, programming for GPUs still requires profound knowledge of the GPU architecture [168]. As with many-core CPUs, the large number of configurable parameters render the performance tuning of GPU programs an important but challenging task [158].

2.3 Accelerated Processing Units (APUs)

Architecture: APUs integrate CPU and GPU on a single die. Although the term APUs has been coined by AMD, recent Intel CPUs with Intel HD Graphics follow a similar architecture. Unlike stand-alone GPUs, the fused GPU of an APU has direct access to the host memory through a low-latency and high-bandwidth memory bus. Figure 3 sketches the high-level architecture of an APU.

Benefits: The main benefit of APUs is the opportunity for zero-copy memory access: since all memory is accessible both from the CPU and the GPU, costly data transfers over a relatively low-bandwidth bus like PCI-E can be avoided. Zero-copy memory access also provides savings in memory space, as only one copy of an object in memory is required. Further, scattered memory accesses which could only be executed inefficiently by the GPU can instead be performed by the CPU.

Limitations: Existing APU products have focused more on energy efficiency than high performance. They typically contain fewer processing units than stand-alone CPUs and GPUs of the same hardware generation. For example, the Ryzen 5 2400G APU by AMD has 704 Vega-based stream

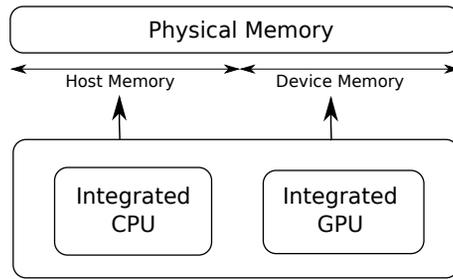


Fig. 3. In an APU, memory is shared between the fused CPU and GPU.

processors, while the stand-alone graphics card AMD RX Vega 64 has 4096 stream processors. As a consequence, compared to high-end stand-alone CPUs and GPUs, their computational power is relatively low. Still, as will be discussed in Section 4, some works have considered APUs for accelerating agent-based simulations.

2.4 Field-Programmable Gate Arrays (FPGAs)

Architecture: A Field-Programmable Gate Array is an integrated circuit made of an array of interconnected Configurable Logic Blocks (CLBs). Additionally, FPGAs are equipped with input and output pads and Digital Signal Processing (DSP) blocks. FPGAs often provide various communication interfaces such as PCI-E, UART, and Ethernet.

A CLB consists of several slices (sometimes also called logic cells), each slice containing a set of storage elements and Look-Up Tables (LUTs). A LUT has a number of inputs and outputs as well as flip flops that store a mapping between possible inputs and outputs. The mapping between inputs and outputs is defined by the users [63]. The number of slices is one of the most important benchmarks to determine the computational power of an FPGA and can range from several thousand to several million. For instance, the XCVU37P Virtex UltraScale FPGA from Xilinx has 2,851,800 slices [170]. In addition, the FPGA may have access to several GB of off-chip DRAM.

To describe the logic to be placed on an FPGA, typically a Hardware Description Language (HDL) is used. The two most widely used HDLs are VHDL [65] and Verilog [121]. In recent years, there have been intensive efforts to enable High-Level Synthesis, i.e., to generate FPGA layouts directly from high-level programming languages such as C, C++, or Java. Recently, Intel released a dedicated SDK to use OpenCL to program FPGAs [75].

FPGAs are sometimes used as a prototyping tool for the development of Application Specific Integrated Circuits (ASICs), which requires an extensive and costly design process. To the best of our knowledge, there is no work so far that employs custom ASICs for ABS. In the field of DES, a number of works have considered offloading of specific simulation tasks to ASICs [19, 30, 51, 102, 134]. Notably, some of the envisioned components were fabricated physically [19]. Since the works on ASICs have only limited relevance to the field of ABS, we exclude them from our survey.

Benefits: Due to the flexibility and high energy efficiency of FPGAs, they are frequently used for computationally intensive and highly parallelisable tasks. For instance, FPGAs can be three orders of magnitude faster than GPUs when conducting specialised tasks such as encrypting a single 64-bit block by the Data Encryption Standard (DES) [27]. In contrast to CPU or GPU, on which data paths are fixed, FPGAs provide flexible and customised data paths [132]. In the past years, FPGAs have received more attention in the field of simulation. FPGAs are particularly suited for

simulations in Electronic Design Automation, since hardware designs can be naturally expressed as FPGA layouts.

Limitations: As with GPUs, FPGAs are connected to a host CPU without direct access to system memory. The resulting need for data transfers can reduce the potential for performance gains.

FPGAs are regarded as lacking in programmability when compared to CPUs and GPUs [27, 44]. Although recent efforts towards high-level synthesis alleviate this limitation, manual tuning is still necessary to achieve the best performance [49, 110].

Finally, FPGAs are configured for a specific task. Since reconfiguration can require multiple hours [180], FPGAs do not facilitate development processes that require fast iteration. This may limit the applicability of FPGAs in early phases of simulation model development, where changes to the simulation model frequently occur and require immediate feedback for evaluation.

3 AGENT-BASED SIMULATION

Agent-based modelling and simulation (ABMS) is a common approach [112] used for evaluating complex systems in domains such as traffic, crowds, economics, information propagation, biology, etc. In the following, we characterise the modelling approach and discuss the properties of ABS with execution on heterogeneous hardware in mind.

3.1 Modelling Approach

In ABS, the simulated entities are agents that perform actions autonomously and interact with other agents based on certain rules. ABS typically follows a *Sense-Think-Act* cycle (e.g., [138]): in the *Sense* stage, an agent detects and analyses its neighbours as well as the environment in which it resides. In the *Think* stage, an agent makes judgement based on the information collected during the *Sense* stage. The update of states takes place in the *Act* stage. The simulation time is typically advanced in fixed time steps at which all agents update their states. However, if a model requires agents to update their states at variable points in simulation time, time advancement using a *discrete-event simulation* (DES) approach may be more appropriate. In DES, state updates are performed through events scheduled for execution at discrete points in simulation time. The simulation proceeds by iteratively executing the earliest of the remaining events, potentially scheduling new events in the process.

Independent of the time advancement mechanism, the defining characteristic of ABS distinguishing it from other simulation techniques is the *autonomy* of agents, i.e., “agents are endowed with behaviours that allow them to make independent decisions” [104]. Since the focus of this survey paper is on ABS, we exclude simulation domains such as physics and chemistry, which usually consider sets of entities that are passively affected by their environment. However, we do discuss a number of methods proposed outside of the ABS domain with direct applications to ABS, e.g., GPU-based priority queues in the context of DES.

3.2 Computational Aspects

When considering models with complex agent decision-making and behaviour at large scales, ABS can be computationally intensive. In addition, due to the stochastic nature of ABS, the simulation of a given scenario is usually repeated multiple times in order to generate meaningful results, further increasing computational demands [85]. However, when considering the *Sense-Think-Act* cycle described above, there are ample opportunities for parallel execution. Since the *Sense* and *Think* stages are performed on a per-agent basis and do not modify the simulation state, these stages can be executed in parallel across all agents. To achieve a consistent view of the simulation state for all agents, the state changes performed in the *Act* stage must then be propagated to other processing elements.

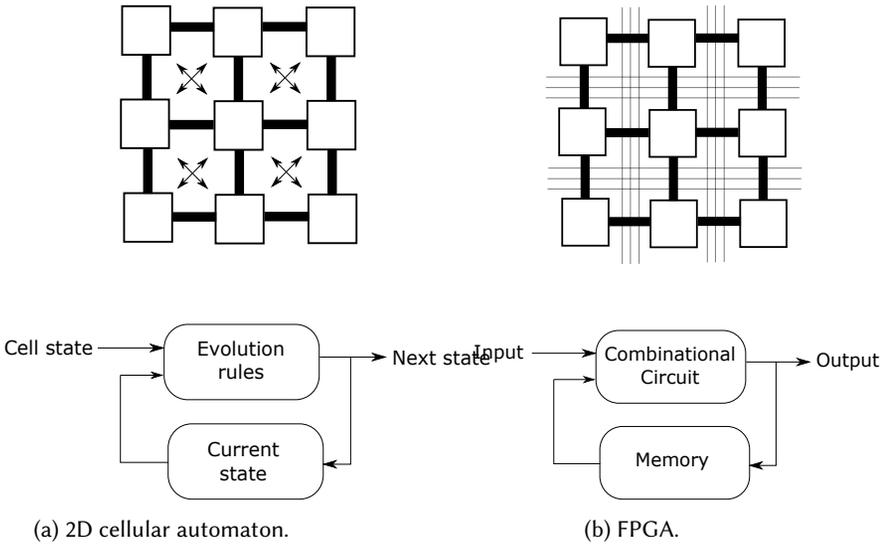


Fig. 4. Structural similarity between a two-dimensional cellular automaton and an FPGA [162].

When parallelising across multiple traditional CPUs or CPU cores, each processing element can execute the state updates for a subset of agents. A well-known challenge in parallel and distributed ABS lies in partitioning the simulation workload among the processing elements. Generally, there are two dimensions according to which a simulation can be partitioned [109]: *domain decomposition* partitions according to the simulation space (e.g., different roads in a traffic simulation), while *functional decomposition* partitions according to different models (e.g., different layers of the network stack in a computer network simulation). High-quality partitionings are characterised by low amounts of workload imbalance and communication among the processing elements. When targeting heterogeneous hardware environments, the partitioning problem is complicated by the differences in the suitability of each hardware device for certain types of computations. Thus, to achieve high performance, a key challenge is to find a suitable **hardware assignment** of the simulation tasks according to characteristics such as the instruction mix and the available degree of parallelism.

Since typically, some communication between the partitions cannot be avoided, techniques for the **minimisation of data transfers** are required to reduce the performance impact of the communication (e.g., [80]).

On CPUs, the ABS performance benefits from long instruction pipelines, large caches and effective branch prediction. Beyond traditional parallel and distributed simulation, many-core CPUs enable high degrees of parallelism while supporting unmodified x86 code. The key difference between a CPU execution in a multi-core and many-core setting is the interconnect through which the CPUs communicate. Since the architecture of each core still closely follows a traditional CPU core, no major code adaptations are required to execute the agent update logic efficiently.

In contrast, since both GPUs and FPGAs achieve highest performance with computational problems of a highly regular structure, another challenge of executing ABS on hardware accelerators lies in dealing with the **scattered memory accesses** resulting from the largely unpredictable runtime behaviour of the simulation. Further, irregular control flows and fine-grained computations make it challenging to fully utilise high-performance many-core devices. Thus, methods for the **maximisation of parallelism** are required. As an example, consider a model where the simulation

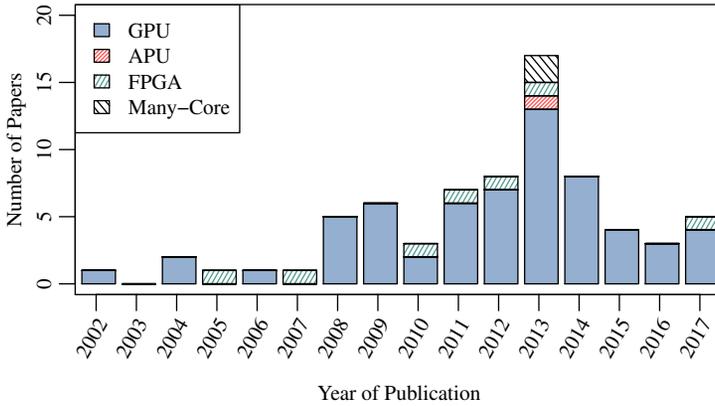


Fig. 5. Publications on agent-based simulation on heterogeneous hardware by year and hardware type.

space is represented by a rectangular grid of cells, each cell being occupied by at most one agent. Here, a simple hardware assignment is a one-to-one mapping of arithmetic units to cells. On a GPU, due to its heritage in highly regular data-parallel tasks on pixel values, such a hardware assignment tends to enable high cache locality, minimisation of memory transactions, and high utilisation of the arithmetic units. In fact, prior to the general-purpose programmability of GPUs, a number of works proposed translating grid-based simulations to operations on graphic textures (e.g., [62]). The Brook language developed at Stanford [24] automates the translation process to graphics operations. Similarly, there is a correspondence between the structure of an FPGA and cellular grids [162]. The basic function of a circuit in an FPGA can be seen as analogous to the function of a cell in a cellular automaton (cf. Figure 4).

However, in many models such as road traffic or social network simulations, the simulation space is a graph. Graph representations adapted to the architectural properties of the available hardware are required to efficiently support sensing an agent's neighbourhood and updating the simulation state while fully exploiting the available hardware. The general trend in the literature is moving towards supporting increasingly irregular types of simulations on accelerators.

The vast majority of literature on ABS using hardware accelerators has focused on GPUs (see Figure 5 for an overview of the number of publications since 2002) We identify three reasons for the popularity of GPUs as accelerators for ABS: first, they are comparatively inexpensive. Second, in the recent years, the ease of programming of GPUs is slowly approaching that of CPUs. Third, well-established programming frameworks such as OpenCL enable the formulation of models in a less hardware-specific manner.

In comparison, the use of FPGAs poses substantial challenges to modellers: only comparatively costly high-end FPGAs run at clock rates close to GPUs. Thus, enormous degrees of parallelism are required to match a GPU's performance. Further, while there exist some frameworks enabling high-level programmability, the achievable performance is limited compared to a more low-level specification of the desired logic in a hardware description language such as VHDL or Verilog. As with GPUs, there is a need for libraries and frameworks that provide a higher degree of **abstraction from hardware specifics**. Finally, the long runtimes of synthesis steps to generate an FPGA layout make model development and adaptation a cumbersome process. Nevertheless, some works consider the use of FPGAs for ABS with promising results [53, 162].

4 ADDRESSING THE CHALLENGES OF AGENT-BASED SIMULATION ON ACCELERATORS

In the following, we discuss the techniques from the literature applicable to the key challenges in ABS on accelerators as identified in Section 3: **hardware assignment**, **data transfer overheads**, **scattered memory accesses**, **maximisation of parallelism**, and **abstraction from hardware specifics**. Table 1 summarises the systematisation of knowledge presented in this survey. It contains our classification of challenges, techniques, and publications, as well as the considered types of accelerators. For the publications that considered specific simulation models, Table 2 shows the simulation domains and hardware platforms, providing researchers with pointers to relevant works in their respective domain.

4.1 Hardware assignment

One of the main challenges in parallel and distributed computations in heterogeneous hardware environments lies in finding a suitable partitioning, i.e., assignment of a given problem to the available hardware [50]. We discuss techniques that have been used to address this problem according to two different, yet interrelated, aspects: first, we consider techniques to select suitable hardware for sub-tasks according to their ability to efficiently execute certain types of computations. The minimisation of data transfers among the partitions running on separate devices will be considered in the next subsection. The existing approaches can be roughly categorized as follows:

- 1. Static assignment:** if the simulation model involves different types of computations that clearly suggest a certain hardware mapping, it may be sufficient to partition the model prior to a simulation run without any adaptation during runtime. For instance, model segments involving large numbers of independent floating point operations may be well-suited for execution on a GPU, whereas segments with highly data-dependent control flow suggest the execution on a CPU.
- 2. Dynamic assignment:** frequently, the dynamic behaviour of a simulated system at runtime translates to unpredictable computational patterns. In such cases, maintaining high performance may require an adaptation of the hardware mapping based on performance measurements at runtime. An inherent challenge of dynamic assignment is the trade-off

Challenge	Technique	Publications
Hardware assignment	Static assignment by type of computation	Many-Core [90], GPU [6, 15, 21, 68, 72, 106, 125, 148, 172] [67, 69, 176], APU [163], FPGA [34, 53, 159, 162]
	Dynamic assignment based on runtime measurements	GPU [18, 56, 59, 86, 165, 176], FPGA [18]
Data transfer overheads	Overlapping of communication and computation	GPU [15, 16, 89]
	Computation replication at partition boundaries	GPU [1, 181]
Scattered memory accesses	Manual caching in shared memory	GPU [96, 135, 181]
	Heuristics for agent update order	GPU [7, 61, 83, 84]
	Representation of irregular data structures by arrays and grids	APU [163], GPU [40, 62, 88, 103, 127–129, 136, 151, 161] [5, 13, 85, 98, 123, 124, 142, 153, 166, 179], FPGA [108, 132]
Maximisation of parallelism	Multiple replications in parallel	GPU [89, 93, 97, 125, 143, 175]
	Window-based event execution	GPU [5, 24, 122, 124, 126, 140, 155, 179]
	Speculative execution	GPU [95, 98], FPGA [108]
	Computation sorting	GPU [85, 89, 155]
Abstraction from hardware specifics	Frameworks to support simulation development	Many-Core [92], GPU [71, 92, 103, 135, 137]
	Unified memory access	GPU [77, 78, 94, 173]

Table 1. A classification of the challenges in agent-based simulation on accelerators along the relevant works addressing them.

Domain/Hardware	Many-Core CPU	GPU	APU	FPGA
Mobility		[127] [129] [151] [143] [164] [172] [148] [71] [72] [7]	[163]	[159]
Biology	[92]	[136] [40] [1] [128] [135] [181] [155] [67] [69] [85] [137] [96]		[34]
Ecology		[93] [175]		[162]
Social	[90]	[83] [84] [179] [161] [97] [168] [95]		[53]
Physics and Chemistry		[88] [142] [106] [16] [126] [62] [175]		[108]
Network		[89] [21] [6] [122] [124] [140] [155] [5]		

Table 2. Simulation model domains considered in the works covered in the survey.

between the performance increase through an improved assignment and the costs of runtime measurements and re-assignment.

An ample body of research has considered the parallelisation of general programs onto heterogeneous platforms, which is an enormous challenge due to the arbitrary control flows and memory access patterns that can be present in general programs. Thus, typically, the approaches limit themselves to program portions that are particularly amenable to parallelisation on accelerators. In the case of ABS, constraints such as the separation of data into a per-agent state and the limited sensing range of agents somewhat simplify the problem of parallelisation, potentially enabling a higher degree of automation in the hardware mapping. In Section 5, we outline the vision of an automated approach and the required building blocks towards an automated hardware mapping for heterogeneous ABS.

4.1.1 Static assignment. Hirabayashi et al. [72] compare a fully GPU-based execution to a hybrid CPU-GPU scheme where the CPU controls the progress of the simulation and calls the GPU for specific tasks. In a traffic simulation based on the Optimal Velocity model [11], the fully GPU-based acceleration clearly outperforms the hybrid scheme, although the lack of synchronisation operations across blocks introduces errors into the simulation results.

A similar categorisation is presented by Pavlov and Müller [125], who conclude that a CPU-GPU approach where both the CPU and the GPU hold duplicated or partial agent and environment information is the most promising. Their approach employs a resource manager on both CPU and GPU to synchronise the data. Scalability is achieved by increasing the number of GPUs or CPU/GPU nodes.

Andelfinger et al. [6] compare four GPU/CPU simulator architectures in the context of discrete-event simulations. In a basic CPU/GPU hybrid scheme (cf. Figure 6a), the CPU offloads each event to the GPU individually. Input data is transferred to the GPU at the beginning of the cycle. After the computation is completed, the output data is transferred back to the CPU. By aggregating independent events and executing them in parallel in a single step, data transfers are reduced (cf. Figure 6b). A further optimisation in data transfers is achieved by exploiting situations where an event relies on results computed by a previous event. Such results can remain in graphics

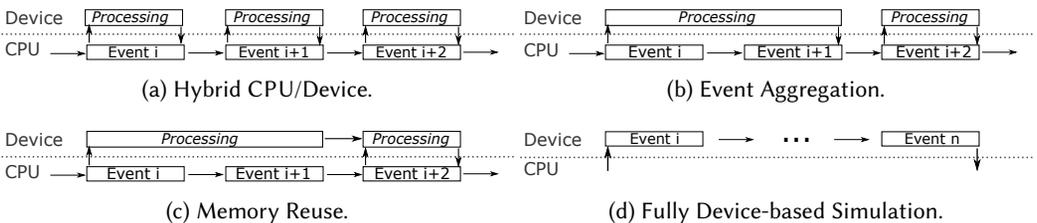


Fig. 6. Four CPU-Device simulation schemes [6]. Devices in the figure can be GPUs or many-cores.

memory and read directly from graphics memory by subsequent events (cf. Figure 6c). In the last architecture, the entire simulation is ported to the GPU. Data transfers are only required at the start of the simulation and once the simulation terminates (cf. Figure 6d). While GPU utilisation and simulation performance increase with each of the above optimisations, the overall simulator architecture becomes more and more dispersed among the CPU and GPU portions, complicating development and reducing maintainability.

Lai et al. [90] implement the four geo-spatial applications of Kriging interpolation [144], ISO-DATA [10], Game of Life [52] and an urban sprawl simulation using cellular automata [169]. The authors compare the performance achieved when using one CPU per execution node, one GPU per node and 60 cores per CPU-based many-core accelerator, using MPI for inter-node communication in each instance. Two approaches for parallelisation are explored: the offloading approach and the direct host approach, which correspond to the hybrid CPU/device (cf. Figure 6a) and fully device-based simulation (cf. Figure 6d) in [6], respectively.

The conclusion drawn from this study is that the use of GPUs and CPU-based many-core accelerators both provide a performance benefit over the purely CPU-based execution. Given a sufficiently large number of assigned processors, using the CPU-based many-core accelerator with the direct host approach achieves similar performance as the GPU-based acceleration.

Another approach to static hardware assignment is proposed by Bauer et al. [14], where in combined continuous-discrete simulation, the continuous part is assigned to the GPU while the discrete part is assigned to the CPU. The simulation is synchronised optimistically using the Time Warp synchronisation algorithm [81]. In Time Warp-based simulations, processing elements process events received from the other processing elements in the order of their logical timestamps reception time. A roll-back mechanism is invoked if an event with an earlier time stamp than a previously executed event is received. Bauer et al. employ the benchmark model PHOLD [46] to explore different GPU configurations such as the thread block size, the number of floating point instructions, the amount communication between CPU and GPU, and the communication pattern. The authors conclude that while keeping the GPU fully utilised poses a challenge, models in a combined simulation with a large number of floating point computations can benefit from GPU acceleration.

When the underlying simulation can be clearly separated into model computation and management tasks, a master-worker scheduling approach as shown by Bilel et al. [21] in the context of large-scale mobile networks simulation can be used. In the proposed design, the model is executed on the GPU, while the CPU orchestrates the event scheduling, simulation status monitoring and memory allocation. A node of the simulation network is partitioned into multiple processes, each process being executed by one GPU warp.

The nature of traffic simulation allows for a relatively straight-forward static hardware assignment according to different simulation aspects. Xu et al. [172] and Song et al. [148] introduce a mesoscopic traffic simulation in a hybrid CPU-GPU architecture. The computation of the agent mobility is assigned to the GPU whereas the route calculation, the agent generation, and file reading and writing remain on the CPU. The two parts run asynchronously to hide data transfer latencies.

Taking into account zero-copy memory access, Wang et al. [163] show how a road traffic simulation can be accelerated using an APU. In their simulator, sorting of agent states is required to locate each agent's neighbours. While the GPU resources of the APU are used to perform state updates and local sorting, the sorting across states assigned to different GPU blocks is performed using the CPU resources of the APU. The work separation between the CPU and GPU resources can be performed efficiently due to the APU's facilities for zero-copy access.

Using FPGAs, Tripp et al. [159] showed how the movement of agents on individual lanes can be computed on the FPGA, while the agents' transitions from one road to another as well as the

behaviour at intersections is computed on the CPU. However, most works on ABS on FPGAs focus on simulation models that allow for statically assigning the entire simulation to an FPGA. For instance, the representation of cellular grids on FPGAs is explored by Vourkas and Sirakoulis [162], who implement an environmental model simulation based on cellular automata (CA). The authors note the structural similarity between a two-dimensional cellular automaton and an FPGA (cf. Section 3.2). A lattice of cells is simulated, each Configurable Logic Block (CLB) simulating one cell. In case the number of cells exceeds the number of CLBs, the simulation lattice is partitioned into several layers, which are processed one after the other. Similarly, Cui et al. achieve high performance with grid-based cellular automata on an FPGA [34]. A pipeline comprised of address generation, reading from memory, data alignment, rule computing, and updating of memory is applied to maximise throughput. A similar method can be applied to cellular automaton-based crowd evacuation models as demonstrated by Georgoudas et al. [53]. They implement the simulation purely on an FPGA and split the overall simulation task into three parts: the closest exit identification circuit, the closest neighbourhood coordinates circuit, which computes the coordinates of the direct neighbours of a cell, and the shortest route tracking circuit, which identifies the closest neighbour towards the exit.

General guidelines for development of GPU-accelerated ABS starting from a CPU-based simulator implementation are proposed in [68, 106]. The methodology requires the decomposition of the simulation model into small task modules and the heuristic identification of modules suitable for execution on a GPU. As a heuristic, the authors state that loops and code segments with low amounts of conditional branching tend to be suited for execution on a GPU. Then, the original task modules are manually replaced with GPU-executable modules. Several case studies [67, 69, 106] show promising speedup by deploying this method.

In the ABS framework TurtleKit, the authors leave the simulation of agent behaviours to the CPU while environment dynamics are handled by the GPU [106]. With this approach, the authors aim to reduce the impact of the GPU acceleration facilities on the maintainability of the simulator code. To increase the performance, portions of the agent behaviour that do not depend on the agent state, e.g., perception of properties of the environment, are performed on the GPU independently of individual agents for all locations and time steps.

Generally, the approaches relying on static hardware assignment split the simulation workload into coarse-grained functional tasks so that some tasks are clearly suited for a certain hardware device. To minimise trial-and-error, heuristics may be applied to identify a suitable mapping of tasks to the hardware. For instance, in the literature, tasks involving large numbers of parallel floating point operations are among the most common tasks offloaded to accelerators. Further observations are made by Zhang et al. in the context of co-running programs on a CPU and GPU or in an APU [176]: 1. programs that are suitable to run in a CPU/GPU environment tend to have low memory bandwidth usage, 2. most programs suitable for the APU allow for a large amount of overlap between CPU and GPU computations.

4.1.2 Dynamic assignment. While a wide range of literature has considered the problem of dynamically adapting a partitioning of agent-based simulations to multiple CPUs (e.g., [33, 100, 171]), we are not aware of such works that specifically target heterogeneous hardware environments. In the following, we outline recent works on dynamic assignment of general computational workloads to heterogeneous hardware. Since these works are generic, they cannot rely on knowledge of the general structure of ABS simulators or on model knowledge.

Belviranlı et al. [18] propose a self-scheduling scheme for partitioning generic application workloads into blocks and assigning them to CPUs, GPUs and FPGAs. The proposed system consists of two phases: in the first phase, the system performs an online training with a small amount of data to estimate the maximum workload capacity size of each hardware device. Fast

convergence is achieved by fitting four sampled data points to a logarithmic function. Once the capacity is determined, the processing unit's performance can be inferred from the same data. When the change of processing speed between two samples drop below a threshold, it is used as the final estimated value. In the second phase, the remaining workload is partitioned based on the percentage of the individual processing speed to the total speed of all available processing units, enabling faster processing units to handle a larger portion of the workload.

Some authors use machine learning techniques such as support vector machines, artificial neural networks and decision trees to distribute the workload of OpenCL programs to CPUs and GPUs. For example, Grasso et al. [56, 86] and Zhang et al. [176] focus on translating a single-device OpenCL program to a multiple-device program, while Wen et al. [165] focus on scheduling multiple OpenCL functions to run in parallel on CPU/GPU. They train a machine learning algorithm according to a set of typical OpenCL programs and benchmarks. The prediction generated by the machine learning algorithm then guides the assignment of a portion of the computation to CPU or GPU. Their results show that the above three machine learning approaches outperform purely CPU- or GPU-based approaches. The scheduling scheme by Wen et al. achieves a performance improvement compared to both a first-come-first-server scheme and a scheme where computation-heavy task are executed on the GPU.

To automate the compilation of sequential programs for parallelised execution on heterogeneous hardware, Grosser and Groesslinger [59] present a compiler that generates CPU and GPU code. Regions with mostly static control flow and sufficient computational intensity are detected and transformed to a formal representation to facilitate program transformations [58]. After optimisations have been performed to increase memory access locality and parallelism, CUDA code for GPU is generated from the formal representation. A runtime library eliminates repeated memory allocations and unnecessary data transfers between CPU and GPU. The decision whether a region is compute-intensive enough for execution on the GPU is made statically or at runtime using heuristics based on metrics such as the number of instructions. The authors conclude that the compiler is able to translate CPU code into cross-platform code with no performance penalty. For some computations, such as the correlation benchmark from polybench [130], significant speedup of up to two orders of magnitude can be achieved.

The main difficulty in automated hardware mapping lies in determining the control flow and data dependencies of the original program. Current approaches either rely on the program code being formulated in languages such as OpenCL that express independent control flows explicitly, or only consider specific portions of programs such as loops with largely static control flow. In ABS, however, most of the available parallelism may exist across the update routines of separate agents. Thus, without semantic information describing the code structure, automatic detection of the parallelism is challenging. In Section 5, we sketch how the common structure of many ABS may be utilised to support the extraction of parallelism.

4.2 Minimisation of data transfer overheads

Since most hardware accelerators are equipped with their own memory, simulations making use of accelerators typically require data transfers between host and accelerator memory. Even with a high-quality partitioning of the simulation, these data transfers incur an overhead that reduces the speedup gained from the distributed computation. In this section, we survey works that focus on minimising the cost of such data transfers. The existing approaches can be roughly categorized according to the following techniques:

1. **Overlapping of communication and computation:** since some communication overhead between the processing elements involved in a simulation cannot be avoided, some authors

proposed techniques to hide communication overheads by transferring data while independent computations are performed. Sometimes, the technique has been referred to as *latency hiding* (e.g., [23]).

2. **Computation replication at partition boundaries:** another technique to address communication overheads is to increase the amount of computation performed before synchronisation between processing elements is required. This is achieved by duplicating some computations on multiple processing elements, thus delaying the need to resolve data dependencies across processing elements.

4.2.1 Overlapping of communication and computation. One way of mitigating the overhead from data transfers between the host and an accelerator is to execute computations at the same time as data is being transferred. In the approach described by Kunz et al. [89], event computations are overlapped with data transfers across the CPU-GPU boundary, thus hiding data transfer latencies in a pipelined fashion. Since events from multiple simulation instances are considered concurrently, there are substantial opportunities for overlapping these steps.

Bauer et al. [15, 16] propose a generic API to optimise the data transfer between global memory and shared memory of CUDA GPUs using so-called warp specialisation. The warps within one cooperative thread arrays are split into two groups: *Dedicated memory warps* are in charge of data transfer between the on-chip and off-chip memory. *Compute warps* process the data. The approach improves performance over thread-level separation between communication and computation since separate warps can follow divergent control flows without any performance penalty. While their general idea can be applied to other types of independent processing elements, the warp-based implementation is specific to GPUs.

4.2.2 Computation replication at partition boundaries. In time-stepped ABS, at model time t each agent updates its state based on the states of its neighbours at time $t - 1$. If the simulation is distributed across multiple processing elements, synchronisation and data transfers are required to provide this information at each time step. The associated communication latencies may make up a substantial portion of the simulation runtime. Thus, some authors have proposed methods to reduce synchronisation by replicating some computations on multiple processing elements, similarly to performance optimisations in numerical computing [38].

Aaby et al. [1] present a multi-level data partitioning scheme for cellular simulations on multi-CPU/GPU clusters. The simulation state is partitioned into blocks and each block is executed by a thread, a core, or a node, depending on the configured granularity. In contrast to the traditional data partitioning into blocks of $B \times B$ cells and synchronisation at each time step, their approach partitions the data into several overlapping $(B + 2R) \times (B + 2R)$ blocks where $((B + 2R)^2 - B^2)$ cells form the overlapping area (cf. Figure 7). The computation in the overlapping area is performed redundantly by multiple processing units. Thus, assuming that at each time step, a cell can only affect its immediate neighbours, R time steps are required for a cell in the inner block to be affected by cells in another processing element. Therefore, synchronisation is only required every R time steps. Between synchronisation points, an error propagates inwards within the overlapping areas, but does not affect the inner $B \times B$ cells before a new synchronisation occurs. This partitioning approach is further employed in multi-GPU clusters on the node-, GPU-, block-, and thread-level, and for multi-CPU clusters at the node-, socket-, core-, and thread-level.

While Aaby et al.'s illustrates the idea based on cellular grids, the approach applies to general ABS. The sensing range of agents is generally limited and provides an upper bound on the propagation of the effects of an agent's actions. As long as overlapping segments of the simulation space can be distributed to the processing elements in a manner so that an effect requires at least $R > 1$ time steps, some synchronisations can be avoided. The generality of the approach is illustrated by

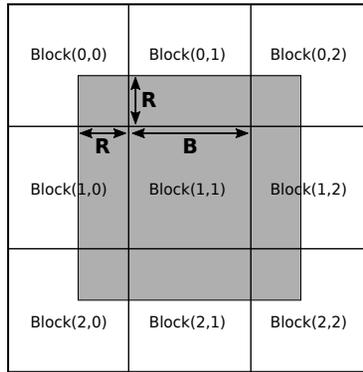


Fig. 7. In the partitioning scheme by Aaby et al. [1], R cells are duplicated among neighbouring processing elements so that each processing element handles $(B + 2R) \times (B + 2R)$ cells.

Zou et al. [181], who extend the idea of computation replication to graph-based topologies in a GPU-accelerated epidemic ABS.

4.3 Scattered memory accesses

Throughout the past decades, the increase in computational performance has outpaced the decrease in memory access latencies, leading to modern hardware designs towards ever-increasing cache sizes and deep memory hierarchies.

In the context of simulations, the issue of memory access latencies is particularly pressing: typically, a model's behaviour cannot be predicted before executing the simulation, significantly limiting the opportunities for a priori optimisation of data access patterns. However, commonalities between different simulation models can be exploited to propose data structures supporting efficient simulation of an entire range of models on a specific type of accelerator.

Since dynamic memory allocation on GPUs is costly [48], most GPU-based simulators allocate graphics memory for the main data structures such as the agent states statically (e.g., [96]). Another approach is to determine after each simulation step the required amount of memory and perform allocations accordingly [124].

The existing approaches to address the issue of scattered memory accesses can be roughly categorized as follows:

1. **Manual caching in shared memory:** although the support for transparent caching has improved in recent years, achieving highest performance frequently still requires manual caching in low-latency segments of an accelerator's memory hierarchy. Since typically the amount of low-latency memory is small, an iterative approach can be taken to limit the number of accesses to high-latency memory when accessing large amounts of data.
2. **Heuristics for agent update order:** since the data dependencies between agent state updates are typically not known prior to the execution of the simulation, minimising cache misses during the state updates is non-trivial. Heuristic have been proposed, aiming to favour sequences of computations acting on the same agent data.
3. **Representation of irregular data structures by arrays and grids:** the hardware architecture of GPUs and FPGAs is designed so that highest performance is achieved when acting on regular data structures such as arrays and grids. Thus, efforts are taken to represent highly irregular data structures in a regular fashion. When covering the techniques from the literature, we first cover *model-specific data structures* such as graph representations of

a simulated road network. Subsequently, we discuss works covering two generic building blocks commonly required as part of ABS engines: *priority queues and sorting*.

4.3.1 Manual caching in shared memory. Richmond et al. [135] propose to utilise the shared memory of the GPU as a manual cache. In their agent-based simulation framework for cellular models in biology based on FLAME GPU [31], they copy sets of messages to be transferred between agents into shared memory. Each thread within a block can then efficiently iterate through the messages and identify those pertaining to the local agent. Once all threads have iterated through the messages, the next sets of messages are loaded into shared memory.

Similarly, Zou et al. [181] implement a manual software cache in shared memory to increase the performance of their graph-based epidemic simulation on GPU clusters. Before the simulation commences on the GPU, the CPU sorts the edges of the directed graph according to the source vertex. Each thread block's shared memory stores edges originating from one specific node. Since each block processes only edges originating from this node, a cache hit rate of at least 50% is ensured.

In agent-based simulation, agents often influence and are influenced by their direct neighbours. This fact can be exploited when arranging the simulation data in memory, reducing high-latency memory accesses when updating agents. Li et al. [96] propose such a method for GPU-based ABS: Assuming a constant number of agents, each agent is assigned to a GPU thread and its state data is permanently kept in global memory. The simulation space is partitioned into a grid of rectangles. Once a search for the neighbours within a circle around an agent is required, a search rectangle that encloses the searching circle is created, so only agents inside the search rectangle have to be considered. Two approaches to utilise the GPU's shared memory are proposed: in the first approach, one block manages the searching process for a chunk C of close-by agents. Per-block shared-memory loads the data of the agent and the agent's neighbours. Each agent in C has a high probability of being in the other agents' neighbourhoods, so that these agents can frequently be accessed through the current block's low-latency shared memory. However, since the limited shared memory capacity allows only for small numbers of agents to be stored, it is still likely that some neighbours are managed by another block and thus have to be accessed through global memory. In the second approach, the shared memory loads the data of agents located in the union of all search rectangles of the agents' handled by the current block. If the shared memory is not sufficient to hold all agents' data, the data is loaded as a sequence of chunks. Of course, the increase in the search space given by the union of search rectangles leads to a higher number of unnecessary agent access through shared memory. To address this problem, the union rectangle can be constructed on the warp level instead of the block level.

4.3.2 Heuristics for agent update order. The order in which agent updates are performed must adhere to the causal dependencies between the agent states and behaviours, e.g., in road traffic simulation, vehicles in direct proximity must be at the same point in simulated time to be able to interact according to the model specification.

Typically, this is achieved by a strictly time-stepped scheme in which agents always reside at the same time step, after which conflicts in the resulting agent states are resolved [174]. However, since in a typical simulation not all agents interact at each point in time, some agents may be updated further into the simulated future than others without affecting the simulation results [7]. Harris and Scheutz have shown that distributed agent-based simulations can be accelerated by favouring agent updates that resolve dependencies across multiple processing elements [61]. This way, processing elements waiting for others to proceed can be unblocked, decreasing the amount of idle time. Their approach can be applied independently of the underlying hardware platform, but requires bounds on the agent movement per time step.

Jin et al. [83] present an information propagation simulation supporting execution on HPC systems and single GPUs and extend it to run on multiple GPUs [84]. Their focus lies on maximising the cache hit rate when traversing a graph according to rules defined by the simulation models. Two categories of approaches are developed for the cascade model [54] and the threshold model [55], which both simulate the propagation of information among nodes in a graph: vertex-oriented processing and edge-oriented processing. For the vertex-oriented approach, the authors further describe two agent update orders: one iterates starting from active vertices, i.e., those that already have the information, and the other from inactive vertices. Since the costs depend on the portion of active nodes, the simulation can switch dynamically between the two vertex-oriented approaches. Finally, the edge-oriented approach iterates over the connecting edges between two vertices. Since the number of edges is constant over a simulation run, the cost of the edge-oriented approach is less variable than that of the vertex-oriented approaches. The authors achieved the highest performance when dynamically switching between the two vertex-oriented approaches.

4.3.3 Representation of irregular data structures by arrays and grids. GPUs and FPGAs are particularly suited for operations on regularly structured data. However, many model types specify topologies that are more naturally expressed in terms of irregular structures such as graphs. Further, execution of the simulator core itself may require operations on irregular data structures.

A basic optimisation commonly applied in works on GPU-based computing to improve memory access patterns is the transformation of the data layout in memory from arrays of structures (AoS) to structures of arrays (SoA) (e.g., [135, 151]). Commonly, sequential programs represent data in an AoS representation. Since AoS bundles the properties associated with each object in object-oriented programming, or the states of agents in agent-based simulations, it is a natural way to represent data within these paradigms. However, with an AoS data layout, parallel operations on the same property across many objects results in scattered memory accesses. An SoA data layout bundles the same property across all objects, which can increase cache hits rates and opportunities for memory access coalescing, thus improving performance substantially.

Beyond this simple optimisation, the data representation can be specialised for a given model to further improve performance. In the following, we give an overview of methods used in agent-based simulation to achieve high performance by translating irregular data structures to a more regular form.

Model-specific data structures

Early works on executing ABS using GPUs frequently focused on cellular grids and translated the required computations into the graphics processing domain. In a pioneering work done by Harris et al. [62], GPU shaders are used for implementing computations on the RGBA values in a texture that holds the agents' states. The same idea is employed by Lysenko et al. [103], Perumalla and Aaby [128], and Kolb et al. [88].

Perumalla et al. [128] evaluate the performance of running agent-based simulation entirely on a GPU. They ported the cellular models Mood Diffusion [70, 111], Game of Life [52] and Schelling Segregation [141]. Through the Open Graphics Library (OpenGL), individual agent states are mapped to pixel colour values. The authors report a speedup of 15 to 40 compared to CPU-based sequential execution. Kolb et al. [88] develop a particle simulation and a GPU-based collision detection mechanism built on the authors' previous work [87]. Similarly, Richmond et al. [136] utilise the GPU's texture processing ability and map agent states onto texture data. To accelerate the neighbourhood detection, the simulation space is partitioned dynamically according to the agents' current states. The algorithm to generate partitions is borrowed from the particle pinning problem in rigid body particles physics [57, 60]. Identification of the start and end of the partition boundary is performed similarly to the method described in [118]. Textures are used to represent the agents'

states and vertex texture fetching enables the search for the start and end of the partition boundary by comparing the partition value to the previous agent's state.

To make traffic simulations on GPUs possible, Perumalla [127] (and Perumalla et Aaby [129]) proposes to model the road network as a grid made up of cells. A road network in Cartesian coordinates is translated to a grid representation overlaying the network: a cell in the grid is marked as occupied when an edge of the original road network starts in the cell, passes the cell or ends in the cell. In graphics memory, the cells' properties such as turning probabilities and length are stored in texture buffers. Simulation is carried out by performing operations on the texture buffers.

A different method for traffic simulation on GPUs is presented by Strippgen and Nagel [151], who propose a queue-based approach using CUDA. Each road is represented as a single first-in, first-out (FIFO) queue stored in memory in the form of a ring buffer. With the ring buffer, insertion of a vehicle entering a road and removal of a vehicle exiting a road is achieved with constant time complexity. Coalesced memory access can be achieved by processing adjacent roads using adjacent threads. Since the vehicles' mobility is modelled by a fixed per-link velocity, their approach can be considered mesoscopic. Behaviours such as overtaking or lane-changing are not modelled and would require random insertions and removals from the ring buffers, which are associated with linear time complexity.

Other domains in which agent-based simulations have been successfully ported to GPUs using model-specific data structures include collision detection [161] and a simulation study of tuberculosis [40]. In the former, a grid is split into tiles and data at the boundary of the tiles is replicated so that a consecutive space is occupied in the global memory of the GPU. In the latter, the authors propose to use a sorted array according to the liveness status of agents, so that the state of a new agent can be stored in a memory location previously occupied by one of the dead agents.

Sorting and priority queues

Full or partial sorting are frequently required in agent-based simulations, e.g., for neighbourhood discovery or to implement priority queues (PQ) if time advancement is performed in a discrete-event manner. These operations can involve large amounts of data-dependent and scattered memory accesses and are therefore challenging to implement efficiently on hardware accelerators. Since this operation can occupy a substantial portion of the simulation runtime [139], a number of works have focused on memory layouts and algorithms for sorting and priority queues on accelerators.

As building blocks for time advancement in a discrete-event fashion, parallel reduction (cf. Figure 8) and bitonic sorting (cf. Figure 9) are commonly used in GPU- and FPGA-based simulation [85, 123, 142, 163, 179]. We discuss these two operations jointly due to their structural similarities. In both cases, an input array is split into chunks, each chunk being handled by one thread. At each cycle, the sorted arrays/minimum values of two threads are then merged to form a new input array. Thus, at each cycle, the number of chunks and active threads is cut into half. The algorithm iterates until only one thread is active, leaving a sorted array or the global minimum value, respectively.

He et al. [64] propose a parallel heap-based PQ on GPU based on a previous CPU-based design [37]. The data structure resembles a binary min-heap, but stores r items per heap node. Items are inserted and extracted in a joint bulk operation that inserts up to $k \leq 2r$ and extracts up to r elements. At any time the root node is guaranteed to hold the highest-priority elements, while elements of lower priority are gradually inserted into deeper levels of the tree over the course of multiple insert-extract operations. Parallelism can be exploited across the sorting operations on the items within a tree node, across the nodes on one level of the tree, and by processing all even-numbered and odd-numbered levels of the tree in parallel. The costs of the queue operations can be hidden by performing them in parallel with the processing of extracted items.

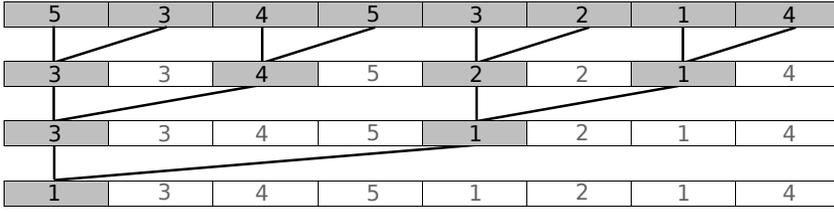


Fig. 8. An example of parallel reduction to compute the minimum among 8 inputs. At each step of the reduction, each active processing element determines the minimum among two of the remaining values.

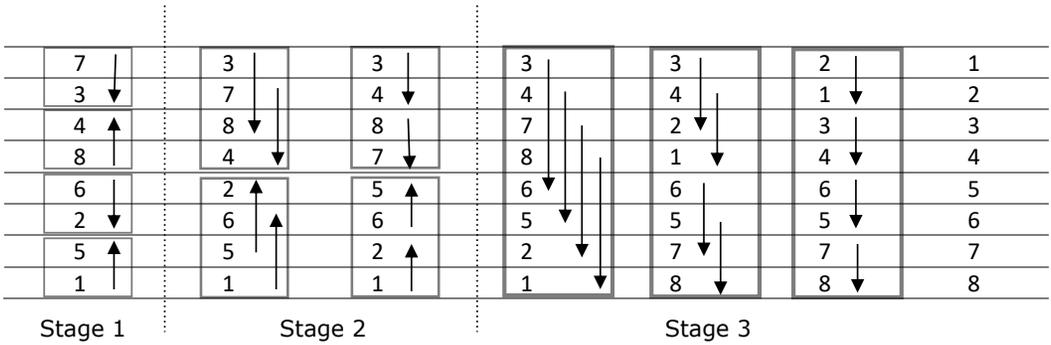


Fig. 9. An example of bitonic sorting with 8 inputs. Rectangles indicate sorting within one thread. Items at an arrow are swapped so that the arrow points to the larger value.

Similarly, the FPGA-based DES simulators by Rahman et al. [132] relies on a pipelined heap [20] for storing events. In contrast to the parallel heap by He et al., the pipelined heap is designed to achieve near-constant access times, but does not provide bulk operations.

A number of works avoid the need for a global PQ holding all future events. Instead, the set of events is considered jointly in an unsorted fashion [153], split by model segment [142] or simulated entity [5, 98, 179], split according to a fixed policy [124, 166], or split randomly [108]. To determine the events that can be executed without violating the simulation correctness, a parallel reduction is performed to determine the minimum timestamp among the events.

Baudis et al. [13] evaluate the performance of PQs on a GPU implemented as a single parallel heap or as a set of ring buffers, implicit binary heaps, and splay trees [146] in the context of DES and path finding on grids. The evaluation results indicate that for up to about 500 elements per PQ, ring buffers achieve the highest performance. At larger element counts, implicit heaps outperform the other approaches in their study. Their results suggest that higher performance is achieved by relying on multiple PQs, one for each agent or set of agents, compared to a single PQ holding all events.

4.4 Maximisation of Parallelism

The limited predictability of how the state of a simulated system evolves over time translates not only to scattered memory accesses, but also to an irregular control flow, which can negatively affect performance in two ways: first, variations in the computational intensity among the model segments may leave some processing elements idle. Second, the single-instruction multiple-thread execution model of GPUs requires divergent operations within a warp to be serialised.

The existing techniques to maximise the parallelism of ABS using accelerators can be roughly categorized as follows:

1. **Multiple replications in parallel:** full utilisation of a massively parallel accelerator requires large numbers of computations that are independent and can thus be executed concurrently. If a simulation involves a sequence of mostly dependent computations, the overheads for communication may outweigh the gains from parallel processing. Thus, techniques have been proposed to aggregate the computations from multiple simulation runs to execute them in parallel.
2. **Window-based event execution:** in simulations involving a discrete-event mechanism, only a proper subset of the simulated entities may require an update at a certain point in simulation time. Multiple authors have proposed gathering events across a window in simulated time, and executing these events in parallel. In effect, this approach forces a discrete-event approach into a time-stepped execution. A key difference among the techniques lies in whether the simulation correctness is strictly maintained.
3. **Speculative execution:** as in general optimistic parallel and distributed simulation [47], computations may be performed speculatively to improve hardware utilisation. A rollback mechanism is required to revert to a correct simulation state after erroneous computations.
4. **Computation sorting:** when assigning neighbouring threads of a GPU to individual agents or events, divergence occurs if required computations are inhomogeneous. Some authors have proposed sorting of computations to minimise divergence.

4.4.1 Multiple Replications in Parallel. If an individual simulation run does not provide sufficient parallelism to fully utilise the available hardware, a Multiple Replications in Parallel (MRIP) approach [125] can be applied, as shown by Shen et al. [143]: in their approach, multiple replications of a traffic simulation [164] are executed in parallel on a GPU. Thus, both the parallelism among agents as well as the parallelism across replications can be exploited. Laville et al. [93] implement a multi-agent simulation of microorganisms in soil for CPU/GPU in OpenCL. Each GPU thread manages one agent and each block is responsible for one simulation instance so that multiple simulation instances can run concurrently on one graphics card. The idea is applied to discrete-event simulations by Kunz et al. [89], focusing on executing parameter studies comprised of multiple replications on a GPU.

In addition to exploiting the parallelism across replications, Li et al. [97] aim to avoid unnecessary redundant computations common to multiple replications. They propose a cloning mechanism for ABS on the GPU: in an ensemble simulation run comprised of multiple simulation instances, the computations that are common to multiple instances are only performed once. When the behaviour of an agent diverges between two simulation instances, a clone of the agent is created. Since the agent may affect other agents, cloning is performed according to the propagation of the effects of the original change in agent behaviour. Across cloned simulation instances, neighbour detection can be aggregated to improve the utilisation of the GPU resources. The benefit of cloning is limited when simulation runs diverge strongly, e.g., across multiple runs of a stochastic simulation using different seeds for random number generation. Recently, the cloning approach has been applied to large-scale cellular simulations on GPU clusters [175].

4.4.2 Window-based event execution. On a GPU, all threads in a warp execute the same sequence of instructions on different elements of data. If no input data is available for some of the threads within a warp, the hardware utilisation is reduced. In ABS, this issue is particularly obvious when time advancement is performed in a discrete-event fashion to accommodate varying state update intervals among the agents. Then, events are scattered along the time axis, i.e., the probability

that many events share the same timestamp may be low. Thus, a simple parallelisation across the events at a certain point in model time may be insufficient. An approach to address this problem is to execute DES models in a time-stepped fashion: all events within a certain time interval are executed in parallel. The lower bound of this time interval is usually referred to as Lower Bound on Time Stamp (LBTS), which is similar to Global Virtual Time in optimistically synchronised parallel and distributed simulation [81]. With a sufficiently large time step size, hardware utilisation is increased. However, since dependencies between events are not considered, the simulation results may differ from a sequential execution.

A study comparing the performance of time advancement mechanisms for simulations on the CPU and the GPU is presented by Perumalla [126]. They study diffusion simulations running in a time-stepped, discrete-event, and hybrid fashion. The GPU variant is implemented in the GPU programming language Brook [24]. While the GPU outperforms the CPU in the time-stepped variant, it does not perform as well as the discrete-event implementation on the CPU. However, high speedup is achieved using the hybrid approach, where at each cycle, the minimum gap between two events is used as a time step. The simulation time then advances according to this time step. Fishwick. [122, 124] present a method for queuing network simulation that executes a DES model in a time-stepped fashion. The simulation time advances according to a fixed time step size, but skips periods where no events occur. All events within the current time step are executed in parallel without considering their potential dependencies. Although the simulation results may be affected by their approach, the authors show that for a queueing network simulation, error bounds can be given. Other works assume a minimum time delta between an event and its creation (*lookahead*) to guarantee the correctness of the simulation results [5, 140, 179]. If lookahead is available, a window can be determined within which events are independent, allowing for parallel execution without affecting the simulation correctness.

The current time window is extended dynamically in work by Tang and Yao [155] to allow more events to be executed in parallel. After executing all events within the current window, their algorithm evaluates the first event in the event queue with a timestamp larger than the LBTS that can still safely be executed according to the lookahead.

4.4.3 Speculative execution. To maintain the correctness of the simulation results when executing in parallel on an accelerator, the simulator must consider the dependencies between state updates. In some of the approaches described above, a time window is determined where state updates cannot affect each other. If it is difficult to determine a time window of sufficient size to extract substantial parallelism, a speculative (also referred to as *optimistic*) approach can be employed: state updates are performed without regard for correctness, and subsequently rolled back if errors are detected.

The possibility of speculative execution of simulations on FPGAs has been first demonstrated by Model and Herbordt [108]. They make use of an event predictor, which predicts the interaction between two particles and generates new events accordingly. Events may later be cancelled as a consequence of a false prediction.

Targeting GPUs, Li et al. [95] present an execution model that avoids divergent control flow by speculative event execution. In an initial step, all events that may occur in the simulation are created. Subsequently, all events are executed in parallel. A scanning process is performed to detect and revoke causally invalid event executions: if an event leaves the simulation in an incorrect state according to a model-specific criterion, the erroneous event and all events created by it are revoked recursively.

A more general approach for GPU-based discrete-event simulation is presented by Liu and Andelfinger [98]. An optimistic execution scheme based on the Time Warp algorithm [81] implemented in CUDA is shown to be beneficial at low event density in simulated time. To support rollbacks in case of erroneous computations, the authors show how the default random number generator in CUDA can be reversed computationally without storing additional data.

4.4.4 Computation sorting. On a GPU, threads within a warp following divergent branches of the control flow are serialised. For instance, if some threads in a warp execute the body of an `if` statement, whereas others execute the body of the corresponding `else` statement, the two sets of threads perform their actions one after another. Some approaches attempt to arrange the assignment of computations to the available threads so that branch divergence is minimised.

In their DES engine on the GPU, Tang and Yao [155] sort events by type before execution, i.e., by the code associated with the event.

The idea is applied to GPU-based execution of multiple simulation instances at the same time by Kunz et al. [89] (cf. Section 4.4.1). If the simulation instances do not diverge too strongly, many events of the same type are available across multiple instances, enabling efficient parallel execution.

Kofler et al. apply computation sorting to their ABS of mosquitoes [85]. In their simulator, a one-to-one mapping between agents and threads is used. Depending on their current state, agents may perform different operations, which can result in taking different control flow branches during the state updates. Thus, to reduce branch divergence among threads within a warp, agents are sorted according to their current state, so that the state updates of adjacent agents share the same control flow.

4.5 Abstraction from hardware specifics

Compared with model development in CPU-based environments, development for accelerators can be cumbersome and error-prone. To avoid the need for modellers to gain deep expertise in programming for specific accelerators, several frameworks have been proposed that enable the specification of parts of the model structure and behaviour in a hardware-agnostic fashion. The approaches to avoid the need for modellers to consider low-level aspects of accelerators can be classified as follows:

- 1. Frameworks to support simulation development:** some authors have proposed generating partial model code to be executed on accelerators from domain-specific languages or the reliance on a library of pre-defined implementations of common simulation tasks and models. However, in these approaches, developing a full ABS will typically still require manual implementation work using a comparatively low-level languages such as CUDA. Further, workload partitioning and assignment to different hardware devices is currently not considered by these approaches.
- 2. Unified memory access:** since in most cases, the CPU and hardware accelerators involved in a simulation operate on separate memory, resolving data dependencies may involve cumbersome explicit data transfers. A number of authors have proposed techniques to transparently access data in programs executed on heterogeneous hardware.

4.5.1 Frameworks to support simulation development. In the Flexible Large Scale Agent Modelling Environment (FLAME GPU) [135, 137], agent states are specified using the state machine model X-Machine [41, 74]. Modellers define agent states in an XML-based format, while state transitions, i.e., the code segments describing the state updates, have to be manually specified as CUDA code. Generic facilities for exchanging messages between agents are provided by the framework. Use cases of the FLAME GPU framework can be, e.g., traffic simulation [71].

Another framework called Many-Core Multi-Agent System (MCMAS) for GPU and other many-core architectures is introduced in [92]. The framework provides a high-level Java interface to OpenCL code as well as a set of pre-defined data structures and functions called plugins. To implement agent models, users either rely on plugins or define their own plugins as OpenCL code that can be called from Java code. The authors state that unlike FLAME GPU, in which models are targeted exclusively to the framework, the models defined in MCMAS can be reused by other agent-based simulators.

While FLAME and MCMAS both reduce the implementation work required to develop agent-based simulations targeting accelerators, these frameworks do not provide guidance or automation in distributing the simulation workload to the available hardware. Thus, manual experimentation is required to determine a suitable hardware mapping.

4.5.2 Unified memory access. GPGPU frameworks such as OpenCL or CUDA require the user to either explicitly trigger data transfers between host and device memory, to explicitly select certain variables or memory regions for access from both CPU and GPU code [117], or to annotate the program to manage data transfers [94, 173]. These manual steps complicate the development of agent-based simulations in heterogeneous environments. Some works aim to improve on this situation by transparently transferring required data between host and graphics memory. However, in languages based on C or C++, static alias analysis, i.e., determining which pointers refer to the same memory regions, is known to be undecidable [78].

Jablin et al. [77, 78] presented the first fully automated data management system based on compilation steps and a runtime library. The developer formulates his program and GPU code as if all data resides in host memory and can be accessed both from the CPU and GPU. The proposed approach then instruments the code to track the accesses to different memory regions using code instrumentation and trapping of system calls. To avoid the need for static pointer analysis, memory accesses through pointers are tracked by the runtime library. In addition to transparently handling data transfers, CPU-GPU communication is optimised during compile time by re-ordering the program flow to reduce the alternation between computations and data transfers. To avoid unnecessary data transfers, the runtime library leaves data in the GPU memory until it is accessed from the host.

While the work of Jablin et al. could be applied to automate data transfers in heterogeneous ABS, the detection of parallelism is not covered. In Section 5, we sketch research directions towards automation in porting ABS to accelerators.

5 TOWARDS AN AUTOMATED OFFLOADING PROCEDURE

From the observations in the previous section, we can state that there is a vast range of techniques covering the main challenges high-performance ABS on hardware accelerators. However, there exist only few ABS frameworks that support such accelerators. Since existing agent-based simulators and models have usually been implemented targeting purely CPU-based environments, there is a clear need for processes and tools to support the transition to an execution on accelerators. More specifically, modellers and simulationists should be supported in the parallelisation and hardware mapping as much as possible. While methodologies have been proposed to systematise the steps of porting a simulation to a GPU [69, 106], there is still a lack of automated tools to support this process.

The problem of automatic parallelisation of general programs is a broad and active field of research [48]. Substantial successes have been achieved with respect to parallelisation of computationally intensive loops with predictable and mostly static control flow [59], whereas the extraction of parallelism across complex and irregular programs is still a largely manual process.

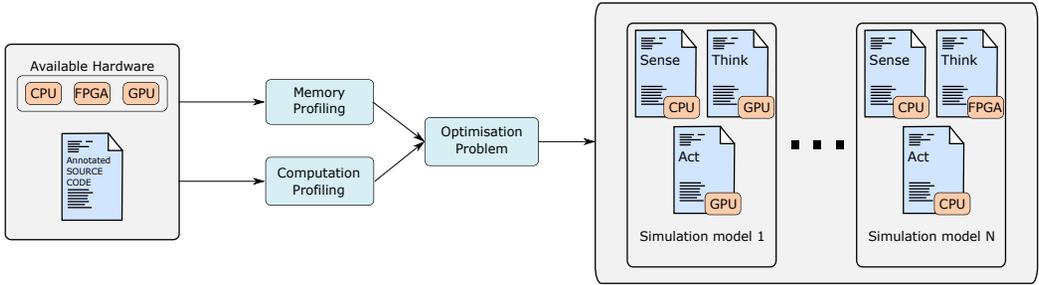


Fig. 10. Workflow of the envisioned automated offloading procedure.

Common approaches include specifying software systems using formalisms that express parallelism explicitly [73, 91, 105] or annotating programs with parallelisation hints [35]. In essence, these approaches provide the compiler or parallelisation middleware with a dependency graph of the statements or code blocks within the original program.

Fortunately, many agent-based simulators and models roughly follow a common set of properties that simplify the extraction of parallelism. We identify the following constraints that can be leveraged to support the parallelisation process:

- (1) Time-stepped execution: usually, the model time is advanced in fixed increments. At each time step, all agents update their states.
- (2) Two states per agent: to decouple the simulation results from the order in which agent state updates are performed, simulators commonly support storing each agent's old state at $t - 1$ and the new state at t separately. During an update from $t - 1$ to t , only read accesses are performed to the agents' states and the environment state at $t - 1$, and only write accesses to the states at t . Thus, within an update, there are no read-after-write dependencies across agents.
- (3) Sense-Think-Act cycle: we assume that agent updates follow the well-known Sense-Think-Act cycle (cf. Sec. 3.1), with one such cycle per model.

With these constraints, a natural approach to parallelisation is to offload individual stages of a model's Sense-Think-Act cycle to an accelerator. For instance, in crowd simulations using the social force model, the Think stage comprised of the computation of the force affecting an agent may be performed by one thread of a GPU per agent.

In the following, we sketch an envisioned workflow and the required tools to support users in porting an existing CPU-based ABS to a system equipped with hardware accelerators. For the targeted simulator architecture, we assume a traditional *master-worker* scheme, with the host CPU acting as the master and assigning work to the available accelerators at each time step.

5.1 Proposed Work Flow

The proposed semi-automated process is visualised in Figure 10. To facilitate the automatic partitioning of the simulation source code into segments that can be outsourced to various types of hardware, we suggest manually annotating the source code according to the Sense-Think-Act paradigm. From that it follows that the smallest unit that can be offloaded to a hardware accelerator in our proposed framework is one of these three stages. Each of the stages is profiled in terms of memory and computational requirements. According to the gathered requirements, an optimisation problem is solved to generate a hardware assignment (rightmost part of Figure 10).

Algorithm 1 Example for model code annotated with the stages of an agent update.

```

1: #pragma agent_begin
2: class Agent:
3:     Coord position;
4:     void executeOnTimeStep():
5:         #pragma sense_begin
6:         List agents = getNeighbouringAgents(position);
7:         #pragma sense_end
8:         #pragma think_begin
9:         Coord velocity = computeVelocity(agents);
10:        #pragma think_end
11:        #pragma act_begin
12:        position = position + velocity;
13:        #pragma act_end
14: #pragma agent_end

```

For simplicity, we assume that all data required by the stage fits into one of the accelerator’s memory entirely. Otherwise, agents could be distributed across multiple accelerators or processed in batches, both implying additional communication costs.

5.1.1 Input. The source code is annotated manually to signify the stages of the Sense-Think-Act cycle, e.g., in the form `#pragma sense_begin`, `#pragma sense_end`, and so forth. A simple example for a crowd simulation is given in Algorithm 1. In addition to the manual annotations, this clear separation may require refactoring of the simulation code. By parsing the annotated source code, the framework obtains a mapping between code and stages that will later be enriched with data from measurements.

The second input is a specification of the available hardware. Each hardware device is characterised by its available memory, computational performance, and host-device data transfer overhead. The computational performance can be stated in terms of single-threaded performance on CPUs, many-core CPUs, GPUs, and APUs. We assume that for an FPGA, only model stages for which implementations already exist are eligible for offloading. Thus, the computational performance of an FPGA is given with respect to specific model stages.

5.1.2 Memory access profiling. Now that the source code is partitioned into offloadable stages and the capabilities of all the hardware components are known, the data dependencies of each stage are determined. Assuming a node in a graph represents one stage, then an edge in this graph represents a data dependency between these stages. The dependency can refer to both agent or environment data. The weight of the edge is the volume of the data that is accessed in the CPU-based simulator, i.e., that has to be transferred during offloading. Usually, the Think stage only has a dependency on the Sense stage within the same model and agent (intra-agent dependency), whereas the Sense stage might depend on the environment and on other agents’ states (inter-agent dependency). Although we assume that an individual stage is not partitioned across multiple hardware devices, the amount of data gathered during the sense stage may vary over the course of the simulation. For instance, if agents form clusters in the simulation space, the number of neighbours per agent may increase over time. Thus, the data dependencies should be measured with respect to typical scenario conditions. To avoid exceeding the memory capacity of one of the considered hardware devices, the profiling can be repeated for a worst-case scenario.

Tools exist that are able to ascribe memory accesses performed during a program run to the source functions, data structures or threads [8]. For instance, the tool PinComm constructs a

dynamic data flow graph from instrumented program executions [66]. The annotations shown in Algorithm 1 allow us to map function names to the separate agent update stages. Thus, it is possible to obtain the amount of memory accessed within each stage. Once the graph describing the amount of memory accesses across stages is created, the implications in terms of memory copying of moving a certain stage to a hardware device can directly be evaluated. For example, if the Think stage is moved to the GPU and the Sense and Act stage remains on the host CPU, then the edges entering and leaving the Think node determine the data transfer overhead. The actual cost of this copy procedure can be obtained from the device specification or through measurements.

5.1.3 Computational profiling. In addition to the memory requirements of each stage, information about the computational characteristics of each stage is required. The estimated runtime could be inferred from hardware performance models [9, 28, 145, 178]. Approaches as those described in Section 4.1.2 can be applied to estimate the suitability of different agent update stages for execution on a certain accelerator. By characterising the workload incurred by each stage in terms of instruction mix and memory accesses as well as the number of agents, the performance of executing the full-scale simulation can be estimated [56, 86, 165, 176]. Alternatively, if the runtime of a stage is dominated by a sub-task that can easily be ported to an accelerator, measurements with respect to this task can be performed directly on the accelerator [18].

5.1.4 Optimisation problem. Building on the graph that represents data dependencies, an optimisation problem of assigning stages to hardware types can be formulated, similar to the approach targeting embedded systems by Zhang et al. [177]. In essence, constraints are formulated so that each stage is assigned to the host or a device, resulting in an overall simulation schedule. Importantly, the optimisation problem must reflect the data location after each stage or time step (e.g., [105]). For instance, to avoid data transfers, it may be more efficient to execute two subsequent stages on the same accelerator. The objective function of the optimisation problem is the overall simulation runtime, i.e., the sum of all estimated execution times on the respective device and the incurred communication costs by distributing nodes of the dependency graph that are connected by an edge.

5.1.5 Output. The output of the optimisation steps is a recommendation of which stages should be executed on which hardware device. It is then the task of the user to port the code of each stage so it can be executed on the assigned device. This might require specific knowledge, e.g., programming in VHDL or OpenCL and can therefore be an obstacle to some researchers. Given that some established simulation models are used by many researchers (e.g., a CSMA/CA model in network simulation or different car-following models in traffic simulation), a public repository of common simulation models could be created, similarly to the plugin approach used in MCMAS [92]. Researchers could download these crowd-sourced simulation models to enable parts of their own simulations to be run on heterogeneous hardware environments, and contribute their own model implementations. Such a repository would also reduce the need to estimate execution times and improve the optimisation results by allowing direct measurements on the potential target devices. Similarly, after porting a specific model stage, new measurements may be performed to provide the optimisation process with more accurate performance data.

5.1.6 Discussion. In our proposed approach, we take a pragmatic perspective: while the envisioned workflow is achievable based on existing building blocks, our assumptions may leave substantial performance potentials unexplored. In particular, by assuming that models and their stages are both executed as a series of dependent steps, we only exploit the inter-agent parallelism within each stage, while any parallelism across stages is not considered. In the following, we revisit the key challenges of ABS using hardware accelerators and sketch key techniques from the

literature that could be applied to maximise the performance benefits given the assumptions made in our approach.

The **hardware assignment** (cf. Section 4.1) is the main focus of the proposed work flow. Above, we describe a static assignment in the form of a functional decomposition. Nevertheless, the optimisation problem that determines the hardware mapping could be updated according to runtime measurements.

To minimise **data transfer overheads** that cannot be avoided (cf. Section 4.2), a bulk execution of multiple simulation runs would be feasible. The optimisation problem could be adapted so that the computational and memory requirements reflect those of each stage executed within multiple simulations runs at the same time. The output of the optimisation process would then be a schedule for an execution in a multiple replications in parallel (MRIP) fashion [89, 125] .

The technique of overlapping computations with data transfers seems challenging in our approach, since we assume a serialisation of the agent update stages. However, within each stage, pre-fetching across stages may be performed by commencing data transfers once some agents have finished a stage.

Scattered memory accesses and the **maximisation of parallelism** (cf. Sections 4.3 and 4.4) could be addressed by providing a library of optimised functions and data structures for operations such as inter-agent communication or neighbour search (e.g., [31, 92]).

A certain degree of **abstraction from hardware specifics** (cf. Section 4.5) is achieved by the automated profiling and hardware mapping of our proposed workflow. Since each stage is executed on a single accelerator, facilities for unified memory access across all devices are not required. Instead, all agent data is updated locally on the accelerator and transferred automatically according to the schedule determined in the optimisation process.

Overall, the envisioned workflow is intended to rely on existing tools and techniques to allow researchers to exploit the hardware at their disposal with reasonable performance gains, while avoiding the need for costly and time-consuming manual optimisation steps as much as possible.

6 CONCLUSIONS

We presented a survey of the literature on agent-based simulation using hardware accelerators. We categorized existing approaches according to the key challenges of hardware assignment, minimisation of data transfer overheads, scattered memory accesses, maximisation of parallelism, and the abstraction from hardware specifics. Our survey provides modellers with an overview of techniques to execute a certain class of models on the available hardware. Methodology researchers are given a summary of the existing work, pointing out research gaps where further exploration is required. Our main observations are two-fold: first, most of the literature in the past years has focused on GPUs. We expect a significant amount of work exploring agent-based simulations on FPGAs to appear in the near future. Second, while a vast amount of work has proposed techniques that allow for efficient execution of agent-based simulations, only a small number of techniques has found their way into a unified framework. Thus, the burden of developing a simulation that is executable in a heterogeneous environment is carried by the modeller. Aiming to reduce the need for expertise in the programming for accelerators, we sketched our vision of a framework to perform an automated hardware mapping and performance optimisation based on building blocks from the literature.

ACKNOWLEDGEMENT

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

REFERENCES

- [1] Brandon G. Aaby, Kalyan S. Perumalla, and Sudip K. Seal. 2010. Efficient Simulation of Agent-Based Models on Multi-GPU and Multi-Core Clusters. In *Proceedings of the International Conference on Simulation Tools and Techniques (SIMUTools '10)*. ICST, Torremolinos, Malaga, Spain, 29:1–29:10.
- [2] Advanced Micro Devices, Inc. 2017. *Radeon's next-generation Vega architecture*. Technical Report 061317_FINAL_V2. Radeon Technologies Group.
- [3] Spiros N. Agathos, Alexandros Papadogiannakis, and Vassilios V. Dimakopoulos. 2015. Targeting the parallella. In *Proceedings of the European Conference on Parallel Processing (Euro-Par '15)*. Springer, Vienna, Austria, 662–674.
- [4] Gary An, Qi Mi, Joyeeta Dutta-Moscato, and Yoram Vodovotz. 2009. Agent-Based Models in Translational Systems Biology. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine* 1, 2 (September 2009), 159–171.
- [5] Philipp Andelfinger and Hannes Hartenstein. 2014. Exploiting the Parallelism of Large-Scale Application-Layer Networks by Adaptive GPU-Based Simulation. In *Proceedings of the Winter Simulation Conference (WSC '14)*. IEEE, Savannah, GA, USA, 3471–3482.
- [6] Philipp Andelfinger, Jens Mittag, and Hannes Hartenstein. 2011. GPU-Based Architectures and Their Benefit for Accurate and Efficient Wireless Network Simulations. In *Proceedings of the International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS '11)*. IEEE, Singapore, 421–424.
- [7] Philipp Andelfinger, Yadong Xu, David Eckhoff, Wentong Cai, and Alois Knoll. 2018. Fast-Forwarding Agent States to Accelerate Microscopic Traffic Simulations. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '18)*. ACM, Rome, Italy, 113–124.
- [8] Imran Ashraf, Mottaqiallah Taouil, and Koen Bertels. 2015. Memory Profiling for Intra-Application Data-Communication Quantification: A Survey. In *Proceedings of the International Design & Test Symposium (IDT '15)*. IEEE, Dead Sea, Jordan, 32–37.
- [9] Sara S. Baghsorkhi, Matthieu Delahaye, Sanjay J. Patel, William D. Gropp, and Wen-mei W. Hwu. 2010. An adaptive performance modeling tool for GPU architectures. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '10)*. ACM, Bangalore, India, 105–114.
- [10] Geoffrey H. Ball and David J. Hall. 1965. *ISODATA, a Novel Method of Data Analysis and Pattern Classification*. Technical Report AD0699616. Stanford Research Institute, Menlo Park, CA.
- [11] Masako Bando, Katsuya Hasebe, Ken Nakanishi, Akihiro Nakayama, Akihiro Shibata, and Yūki Sugiyama. 1995. Phenomenological Study of Dynamical Model of Traffic Flow. *Journal de Physique I* 5, 11 (November 1995), 1389–1399.
- [12] Taylor Barnes, Brandon Cook, Jack Deslippe, Douglas Doerfler, Brian Friesen, Yun (Helen) He, Thorsten Kurth, Tuomas Koskela, Mathieu Lobet, Tareq Malas, Leonid Oliker, Andrey Ovsyannikov, Abhinav Sarje, Jean-Luc Vay, Henri Vincenti, Samuel Williams, Pierre Carrier, Nathan Wichmann, Marcus Wagner, Paul Kent, Christopher Kerr, and John Dennis. 2016. Evaluating and Optimizing the NERSC Workload on Knights Landing. In *Proceedings of the International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computing Systems (PMBS '16)*. IEEE, Salt Lake City, UT, USA, 43–53.
- [13] Nikolai Baudis, Florian Jacob, and Philipp Andelfinger. 2017. Performance Evaluation of Priority Queues for Fine-Grained Parallel Tasks on GPUs. In *Proceedings of the International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS '17)*. IEEE, Banff, Canada, 1–11.
- [14] David W. Bauer, Matthew McMahon, and Ernest H. Page. 2008. An Approach for the Effective Utilization of GP-GPUs in Parallel Combined Simulation. In *Proceedings of the Conference on Winter Simulation (WSC '08)*. IEEE, Miami, FL, USA, 695–702.
- [15] Michael Bauer, Henry Cook, and Bruce Khailany. 2011. CudaDMA: Optimizing GPU Memory Bandwidth via Warp Specialization. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '11)*. ACM, Seattle, DC, USA, 12:1–12:11.
- [16] Michael Bauer, Sean Treichler, and Alex Aiken. 2014. Singe: Leveraging Warp Specialization for High Performance on GPUs. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '14)*. ACM, Orlando, FL, USA, 119–130.
- [17] Nathan Bell and Jared Hoberock. 2011. Thrust: A Productivity-Oriented Library for CUDA. In *GPU Computing Gems Jade Edition*. Elsevier, 359–371.
- [18] Mehmet E. Belviranlı, Laxmi N. Bhuyan, and Rajiv Gupta. 2013. A Dynamic Self-scheduling Scheme for Heterogeneous Multiprocessor Architectures. *ACM Transactions on Architecture and Code Optimization - Special Issue on High-Performance Embedded Architectures and Compilers* 9, 4 (January 2013), 57:1–57:20.
- [19] Jacob L. Berlin. 1993. *Design of a Parallel Discrete Event Simulation Coprocessor*. Master's thesis.
- [20] Ranjita Bhagwan and Bill Lin. 2000. Fast and Scalable Priority Queue Architecture for High-Speed Network Switches. In *Proceedings of the IEEE INFOCOM Conference on Computer Communications (INFOCOM '00)*. IEEE, Tel Aviv, Israel, 538–547.

- [21] Ben Romdhanne Bilel, Nikaein Navid, and Mohamed Said Mosli Bouksiaa. 2012. Hybrid CPU-GPU Distributed Framework for Large Scale Mobile Networks Simulation. In *Proceedings of the IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (DS-RT '12)*. IEEE, Dublin, Ireland, 44–53.
- [22] André R. Brodtkorb, Trond R. Hagen, and Martin L. Sætra. 2013. Graphics Processing Unit (GPU) Programming Strategies and Trends in GPU Computing. *Elsevier Journal of Parallel and Distributed Computing* 73, 1 (January 2013), 4–13.
- [23] Ulrich Brüning, Wolfgang K. Giloi, and Wolfgang Schroeder-Preikschat. 1994. Latency Hiding in Message-Passing Architectures. In *Proceedings of the International Parallel Processing Symposium (IPPS '94)*. IEEE, Cancun, Mexico, 704–709.
- [24] Ian Buck, Tim Foley, Daniel Horn, Jeremy Sugerman, Kayvon Fatahalian, Mike Houston, and Pat Hanrahan. 2004. Brook for GPUs: Stream Computing on Graphics Hardware. In *Proceedings of the International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '04)*. ACM, Los Angeles, CA, USA, 777–786.
- [25] Martin Burtscher, Rupesh Nasre, and Keshav Pingali. 2012. A Quantitative Study of Irregular Programs on GPUs. In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC '12)*. IEEE, La Jolla, CA, USA, 141–151.
- [26] Márcio Castro, Emilio Franceschini, Fabrice Dupros, Hideo Aochi, Philippe Navaux, and Jean-François Mehaut. 2016. Seismic Wave Propagation Simulations on Low-Power and Performance-Centric Manycores. *Elsevier Journal of Parallel Computing* 54 (May 2016), 108–120.
- [27] Shuai Che, Jie Li, Jeremy W. Sheaffer, Kevin Skadron, and John Lach. 2008. Accelerating Compute-Intensive Applications With GPUs and FPGAs. In *Proceedings of the Symposium on Application Specific Processors (SASP '08)*. IEEE, Anaheim, CA, USA, 101–107.
- [28] Xi E. Chen and Tor M. Aamodt. 2009. A First-Order Fine-Grained Multithreaded Throughput Model. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA '09)*. IEEE, Raleigh, NC, USA, 329–340.
- [29] Thomas M. Cioppa, Thomas W. Lucas, and Susan M. Sanchez. 2004. Military Applications of Agent-Based Simulations. In *Proceedings of the Winter Simulation Conference (WSC '04)*. IEEE, Washington, DC, USA, 171–180.
- [30] John G. Cleary, Murray Pearson, and Husam Kinawi. 1995. The Architecture of an Optimistic CPU: The WarpEngine. In *Proceedings of the Hawaii International Conference on System Sciences (HICSS '95)*. IEEE, Wailea, HI, USA, 163–172.
- [31] Simon Coakley, Paul Richmond, Marian Gheorghe, Shawn Chin, David Worth, Mike Holcombe, and Chris Greenough. 2016. Large-Scale Simulations with FLAME. In *Intelligent Agents in Data-intensive Computing*. Joanna Kolodziej, Luis Correia, and José Manuel Molina (Eds.). Springer International Publishing, 123–142.
- [32] N.T. Collier and M.J. North. 2011. Repast SC++: A Platform for Large-scale Agent-based Modeling. In *Large-Scale Computing Techniques for Complex System Simulations*. W. Dubitzky, K. Kurowski, and B Schott (Eds.). Wiley.
- [33] Biagio Cosenza, Gennaro Cordasco, Rosario De Chiara, and Vittorio Scarano. 2011. Distributed Load Balancing for Parallel Agent-Based Simulations. In *Proceedings of the International Euromicro Conference on Parallel, Distributed and Network-Based Processing (PDP '11)*. IEEE, Ayia Napa, Cyprus, 62–69.
- [34] Lintao Cui, Jing Chen, Yu Hu, Jinjun Xiong, Zhe Feng, and Lei He. 2011. Acceleration of Multi-Agent Simulation on FPGAs. In *Proceedings of the International Conference on Field Programmable Logic and Applications (FPL '11)*. IEEE, Chania, Greece, 470–473.
- [35] Leonardo Dagum and Ramesh Menon. 1998. OpenMP: An Industry Standard API for Shared-Memory Programming. *IEEE Journal of Computational Science and Engineering* 5, 1 (January 1998), 46–55.
- [36] B. D. de Dinechin, R. Ayrignac, P. E. Beauamps, P. Couvert, B. Ganne, P. G. de Massas, F. Jacquet, S. Jones, N. M. Chaisemartin, F. Riss, and T. Strudel. 2013. A Clustered Manycore Processor Architecture for Embedded and Accelerated Applications. In *Proceedings of the IEEE High Performance Extreme Computing Conference (HPEC '13)*. IEEE, Waltham, MA, USA, 1–6.
- [37] Narsingh Deo and Sushil Prasad. 1992. Parallel Heap: An Optimal Parallel Priority Queue. *Springer Journal of Supercomputing* 6, 1 (March 1992), 87–98.
- [38] Chris Ding and Yun He. 2001. A Ghost Cell Expansion Method for Reducing Communications in Solving PDE Problems. In *Proceedings of the ACM/IEEE Conference on Supercomputing (SC '01)*. IEEE, Denver, CO, USA, 55–55.
- [39] Arnaud Doniec, René Mandiau, Sylvain Piechowiak, and Stéphane Espié. 2008. A Behavioral Multi-Agent Model for Road Traffic Simulation. *Elsevier Journal of Engineering Applications of Artificial Intelligence* 21, 8 (December 2008), 1443–1454.
- [40] Roshan M. D'Souza, Mikola Lysenko, Simeone Marino, and Denise Kirschner. 2009. Data-Parallel Algorithms for Agent-Based Model Simulation of Tuberculosis on Graphics Processing Units. In *Proceedings of the Spring Simulation Multiconference (SpringSim '09)*. SCS, San Diego, CA, USA, 21:1–21:12.
- [41] Samuel Eilenberg. 1974. *Automata, Languages, and Machines*. Academic Press.

- [42] Joshua M. Epstein. 1999. Agent-Based Computational Models and Generative Social Science. *Complexity* 4, 5 (May 1999), 41–60.
- [43] Fernando A. Escobar, Xin Chang, and Carlos Valderrama. 2016. Suitability Analysis of FPGAs for Heterogeneous Platforms in HPC. *IEEE Transactions on Parallel and Distributed Systems* 27, 2 (February 2016), 600–612.
- [44] Babak Falsafi, Bill Dally, Desh Singh, Derek Chiou, J Yi Joshua, and Resit Sendag. 2017. FPGAs Versus GPUs in Data Centers. *IEEE Micro* 37, 1 (January 2017), 60–72.
- [45] Naznin Fauzia, Louis-Noël Pouchet, and P. Sadayappan. 2015. Characterizing and Enhancing Global Memory Data Coalescing on GPUs. In *Proceedings of the Annual IEEE/ACM International Symposium on Code Generation and Optimization (CGO '15)*. IEEE, San Francisco, CA, USA, 12–22.
- [46] Richard M. Fujimoto. 1990. Performance of Time Warp Under Synthetic Workloads. In *Proceedings of the Distributed Simulation Conference (DSC '90)*. SCS, San Diego, CA, USA, 23–28.
- [47] Richard M. Fujimoto. 2000. *Parallel and Distributed Simulation Systems*. Wiley New York.
- [48] Richard M. Fujimoto. 2016. Research Challenges in Parallel and Distributed Simulation. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 26, 4 (May 2016), 22:1–22:29.
- [49] Richard M. Fujimoto, Conrad Bock, Wei Chen, Ernest Page, and Jitesh H. Panchal. 2017. *Research Challenges in Modeling and Simulation for Engineering Complex Systems*. Springer.
- [50] Richard M. Fujimoto, Christopher Carothers, Alois Ferscha, David Jefferson, Margaret Loper, Madhav Marathe, and Simon J.E. Taylor. 2017. Computational Challenges in Modeling & Simulation of Complex Systems. In *Proceedings of the Winter Simulation Conference (WSC '17)*. IEEE, Las Vegas, NV, USA, 431–445.
- [51] Richard M. Fujimoto, Jya-Jang. Tsai, and Ganesh Gopalakrishnan. 1988. Design and Performance of Special Purpose Hardware for Time Warp. In *Proceedings of the Annual International Symposium on Computer Architecture (SCA '88)*. IEEE, Honolulu, HI, USA, 401–409.
- [52] Martin Gardner. 1970. Mathematical Games: The Fantastic Combinations of John Conway’s New Solitaire Game “Life”. *Scientific American* 223, 4 (October 1970), 120–123.
- [53] Ioakeim G. Georgoudas, Panagiotis Kyriakos, G. Ch. Sirakoulis, and I. Th. Andreadis. 2010. An FPGA Implemented Cellular Automaton Crowd Evacuation Model Inspired by the Electrostatic-Induced Potential Fields. *Elsevier Journal of Microprocessors and Microsystems* 34, 7 (November 2010), 285–300.
- [54] Jacob Goldenberg, Barak Libai, and Eitan Muller. 2001. Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth. *Springer Marketing letters* 12, 3 (August 2001), 211–223.
- [55] Mark Granovetter. 1978. Threshold Models of Collective Behavior. *University of Chicago Press American Journal of Sociology* 83, 6 (1978), 1420–1443.
- [56] Ivan Grasso, Klaus Kofler, Biagio Cosenza, and Thomas Fahringer. 2013. Automatic Problem Size Sensitive Task Partitioning on Heterogeneous Parallel Systems. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '13)*. ACM, Shenzhen, China, 281–282.
- [57] Simon Green. 2010. Particle Simulation Using Cuda. *NVIDIA Whitepaper* 6 (2010), 121–128.
- [58] Tobias Grosser, Armin Groesslinger, and Christian Lengauer. 2012. Polly—Performing Polyhedral Optimizations on a Low-Level Intermediate Representation. *World Scientific Parallel Processing Letters* 22, 04 (December 2012), 1250010.
- [59] Tobias Grosser and Torsten Hoefer. 2016. Polly-ACC Transparent Compilation to Heterogeneous Hardware. In *Proceedings of the International Conference on Supercomputing (ICS '16)*. ACM, Istanbul, Turkey, 1:1–1:13.
- [60] Takahiro Harada. 2007. Real-Time Rigid Body Simulation on GPUs. *NVIDIA GPU Gems* 3 (December 2007), 123–148.
- [61] Jack Harris and Matthias Scheutz. 2012. New Advances in Asynchronous Agent-based Scheduling. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '12)*. Las Vegas, NV, USA, 1–7.
- [62] Mark J. Harris, Greg Coombe, Thorsten Scheuermann, and Anselmo Lastra. 2002. Physically-Based Visual Simulation on Graphics Hardware. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware (HWS '02)*. ACM, Saarbrücken, Germany, 109–118.
- [63] Scott Hauck and Andre DeHon. 2010. *Reconfigurable Computing: The Theory and Practice of FPGA-Based Computation*. Morgan Kaufmann.
- [64] Xi He, Deborah Agarwal, and Sushil K. Prasad. 2012. Design and Implementation of a Parallel Priority Queue on Many-Core Architectures. In *Proceedings of the International Conference on High Performance Computing (HiPC '12)*. IEEE, Pune, India, 1–10.
- [65] Ulrich Heinkel, Martin Padeflke, Werner Haas, Thomas Buerner, Herbert Braisz, Thomas Gentner, and Alexander Grassmann. 2000. *The VHDL Reference: A Practical Guide to Computer-Aided Integrated Circuit Design (Including VHDL-AMS)*. John Wiley & Sons, Inc.
- [66] Wim Heirman, Dirk Stroobandt, Narasinga Rao Miniskar, Roel Wuyts, and Francky Catthoor. 2010. PinComm: Characterizing Intra-Application Communication for the Many-Core Era. In *Proceedings of the IEEE International Conference on Parallel and Distributed Systems (ICPADS '10)*. IEEE, Shanghai, China, 500–507.

- [67] Emmanuel Hermellin and Fabien Michel. 2015. GPU Environmental Delegation of Agent Perceptions: Application to Reynolds's Boids. In *Proceedings of the International Workshop on Multi-Agent Systems and Agent-Based Simulation (MABS '15)*. Springer, Istanbul, Turkey, 71–86.
- [68] Emmanuel Hermellin and Fabien Michel. 2016. Defining a Methodology Based on GPU Delegation for Developing MABS Using GPGPU. In *Proceedings of the International Workshop on Multi-Agent Systems and Agent-Based Simulation (MABS '16)*. Springer, Singapore, 24–41.
- [69] Emmanuel Hermellin and Fabien Michel. 2016. GPU Delegation: Toward a Generic Approach for Developing MABS using GPU Programming. In *Proceedings of the International Conference on Autonomous Agents & Multiagent Systems (AAMAS '16)*. IFAAMAS, Singapore, 1249–1258.
- [70] James D. Hess, Jacqueline J. Kacen, and Junyong Kim. 2006. Mood-Management Dynamics: The Interrelationship Between Moods and Behaviours. *Wiley Online Library British Journal of Mathematical and Statistical Psychology* 59, 2 (November 2006), 347–378.
- [71] Peter Heywood, Paul Richmond, and Steve Maddock. 2015. Road Network Simulation Using FLAME GPU. In *Proceedings of the European Conference on Parallel Processing (Euro-Par '15)*. Springer, Vienna, Austria, 430–441.
- [72] Manato Hirabayashi, Shinpei Kato, Masato Edahiro, and Yuki Sugiyama. 2012. Toward GPU-Accelerated Traffic Simulation and Its Real-Time Challenge. In *Proceedings of the International Workshop on Real-time and Distributed Computing in Emerging Applications (REACTION '12)*. Universidad Carlos III de Madrid, San Juan, Puerto Rico, 45–50.
- [73] Charles Antony Richard Hoare. 1978. Communicating Sequential Processes. *Commun. ACM* 21, 8 (August 1978), 666–677.
- [74] Mike Holcombe. 1988. X-Machines as a Basis for Dynamic System Specification. *IET Software Engineering Journal* 3, 2 (March 1988), 69–76.
- [75] Intel Corporation. 2017. *Intel® FPGA SDK for OpenCL – Programming Guide*. Technical Report UG-OCL002.
- [76] Intel Corporation. 2018. Intel® Architecture Instruction Set Extensions and Future Features – Programming Reference. (January 2018).
- [77] Thomas B. Jablin, James A. Jablin, Prakash Prabhu, Feng Liu, and David I. August. 2012. Dynamically Managed Data for CPU-GPU Architectures. In *Proceedings of the International Symposium on Code Generation and Optimization (CGO '12)*. ACM, San Jose, CA, USA, 165–174.
- [78] Thomas B. Jablin, Prakash Prabhu, James A. Jablin, Nick P. Johnson, Stephen R. Beard, and David I. August. 2011. Automatic CPU-GPU Communication Management and Optimization. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '11)*. ACM, San Jose, CA, USA, 142–151.
- [79] Deepak Jagtap, Ketan Bahulkar, Dmitry Ponomarev, and Nael Abu-Ghazaleh. 2012. Characterizing and Understanding PDES Behavior on Tiler Architecture. In *Proceedings of the ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation (PADS '12)*. IEEE, Zhangjiajie, China, 53–62.
- [80] Myeong-Wuk Jang and Gul Agha. 2006. Agent Framework Services to Reduce Agent Communication Overhead in Large-Scale Agent-Based Simulations. *Elsevier Journal of Simulation Modelling Practice and Theory* 14, 6 (August 2006), 679–694.
- [81] David R. Jefferson. 1985. Virtual Time. *ACM Transactions on Programming Languages and Systems* 7, 3 (July 1985), 404–425.
- [82] Y. Jiao, H. Lin, P. Balaji, and W. Feng. 2010. Power and Performance Characterization of Computational Kernels on the GPU. In *Proceedings of the IEEE/ACM International Conference on Green Computing and Communications & International Conference on Cyber, Physical and Social Computing (GREENCOM-CPSCOM '10)*. IEEE, Hangzhou, China, 221–228.
- [83] Jiangming Jin, Stephen John Turner, Bu-Sung Lee, Jianlong Zhong, and Bingsheng He. 2012. HPC Simulations of Information Propagation Over Social Networks. In *Proceedings of the International Conference on Computational Science (ICCS '12)*. Elsevier, Omaha, NE, USA, 292–301.
- [84] Jiangming Jin, Stephen John Turner, Bu-Sung Lee, Jianlong Zhong, and Bingsheng He. 2013. Simulation of Information Propagation Over Complex Networks: Performance Studies on Multi-GPU. In *Proceedings of the IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (DS-RT '13)*. IEEE, Delft, Netherlands, 179–188.
- [85] Klaus Kofler, Gregory Davis, and Sandra Gesing. 2014. Sampo: An Agent-Based Mosquito Point Model in OpenCL. In *Proceedings of the Symposium on Agent Directed Simulation (ADS '14)*. SCS, Tampa, FL, USA, 5:1–5:10.
- [86] Klaus Kofler, Ivan Grasso, Biagio Cosenza, and Thomas Fahringer. 2013. An Automatic Input-Sensitive Approach for Heterogeneous Task Partitioning. In *Proceedings of the International ACM Conference on International Conference on Supercomputing (ICS '13)*. ACM, Eugene, OR, USA, 149–160.
- [87] Andreas Kolb, Lars John, et al. 2001. Volumetric Model Repair for Virtual Reality Applications. In *EUROGRAPHICS Short Presentation (2001)*. The Eurographics Association, Manchester, England, 249–256.
- [88] Andreas Kolb, Lutz Latta, and Christof Rezk-Salama. 2004. Hardware-Based Simulation and Collision Detection for Large Particle Systems. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware*

- (*HWWS '04*). ACM, Grenoble, France, 123–131.
- [89] Georg Kunz, Daniel Schemmel, James Gross, and Klaus Wehrle. 2012. Multi-Level Parallelism for Time- and Cost-Efficient Parallel Discrete Event Simulation on GPUs. In *Proceedings of the ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation (PADS '12)*. IEEE, Zhangjiajie, China, 23–32.
- [90] Chenggang Lai, Miaoqing Huang, Xuan Shi, and Haihang You. 2013. Accelerating Geospatial Applications on Hybrid Architectures. In *Proceedings of the IEEE International Conference on High Performance Computing and Communications, IEEE International Conference on Embedded and Ubiquitous Computing (HPCC & EUC '13)*. IEEE, Singapore, 1545–1552.
- [91] Leslie Lamport. 2002. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley Longman Publishing Co., Inc.
- [92] Guillaume Laille, Kamel Mazouzi, Christophe Lang, Nicolas Marilleau, Bénédicte Herrmann, and Laurent Philippe. 2013. MCMAS: A Toolkit to Benefit From Many-Core Architecture in Agent-Based Simulation. In *Proceedings of the European Conference on Parallel Processing (Euro-Par '13)*. Springer, Aachen, Germany, 544–554.
- [93] Guillaume Laille, Kamel Mazouzi, Christophe Lang, Laurent Philippe, and Nicolas Marilleau. 2013. Using GPU for Multi-Agent Soil Simulation. In *Proceedings of the Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP '13)*. IEEE, Belfast, UK, 392–399.
- [94] Seyong Lee, Seung-Jai Min, and Rudolf Eigenmann. 2009. OpenMP to GPGPU: A Compiler Framework for Automatic Translation and Optimization. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '09)*. ACM, Raleigh, NC, USA, 101–110.
- [95] Xiaosong Li, Wentong Cai, and Stephen John Turner. 2013. GPU Accelerated Three-Stage Execution Model for Event-Parallel Simulation. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '13)*. ACM, Montreal, Canada, 57–66.
- [96] Xiaosong Li, Wentong Cai, and Stephen John Turner. 2014. Efficient Neighbor Searching for Agent-Based Simulation on GPU. In *Proceedings of the IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (DS-RT '14)*. IEEE, Toulouse, France, 87–96.
- [97] Xiaosong Li, Wentong Cai, and Stephen John Turner. 2015. Cloning Agent-based Simulation on GPU. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '15)*. ACM, London, UK, 173–182.
- [98] Xinhui Liu and Philipp Andelfinger. 2017. Time Warp on the GPU: Design and Assessment. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '17)*. ACM, Singapore, 109–120.
- [99] Xu Liu, Langshi Chen, Jesun S Firoz, Judy Qiu, and Lei Jiang. 2017. Performance Characterization of Multi-Threaded Graph Processing Applications on Intel Many-Integrated-Core Architecture. (August 2017). arXiv:CoRR/1708.04701
- [100] Qingqi Long, Jie Lin, and Zhixun Sun. 2011. Agent Scheduling Model for Adaptive Dynamic Load Balancing in Agent-Based Distributed Simulations. *Elsevier Journal of Simulation Modelling Practice and Theory* 19, 4 (April 2011), 1021–1034.
- [101] Sean Luke, Claudio Cioffi-Revilla, Liviu Panait, Keith Sullivan, and Gabriel Balan. 2005. MASON: A Multiagent Simulation Environment. *SCSI Journal of Simulation* 81, 7 (July 2005), 517–527.
- [102] Elizabeth Whitaker Lynch. 2011. *Hardware Acceleration for Conservative Parallel Discrete Event Simulation on Multi-Core Systems*. Ph.D. Dissertation. School of Electrical and Computer Engineering, Georgia Institute of Technology.
- [103] Mikola Lysenko and Roshan M. D'Souza. 2008. A Framework for Megascale Agent Based Model Simulations on Graphics Processing Units. *JASSS Journal of Artificial Societies and Social Simulation* 11, 4 (October 2008), 10.
- [104] Charles M. Macal and Michael J. North. 2010. Tutorial on Agent-Based Modeling and Simulation. *Springer Journal of Simulation* 4, 3 (September 2010), 151–162.
- [105] Deepak Majeti, Kuldeep S. Meel, Rajkishore Barik, and Vivek Sarkar. 2016. Automatic Data Layout Generation and Kernel Mapping for CPU+GPU Architectures. In *Proceedings of the International Conference on Compiler Construction (CC '16)*. ACM, Barcelona, Spain, 240–250.
- [106] Fabien Michel. 2013. Translating Agent Perception Computations Into Environmental Processes in Multi-Agent-Based Simulations: A Means for Integrating Graphics Processing Unit Programming Within Usual Agent-Based Simulation Platforms. *Wiley Online Library Journal of Systems Research and Behavioral Science* 30, 6 (November 2013), 703–715.
- [107] Sparsh Mittal and Jeffrey S. Vetter. 2015. A Survey of CPU-GPU Heterogeneous Computing Techniques. *Comput. Surveys* 47, 4 (July 2015), 69:1–69:35.
- [108] Josh Model and Martin C. Herbordt. 2007. Discrete Event Simulation of Molecular Dynamics With Configurable Logic. In *Proceedings of the International Conference on Field Programmable Logic and Applications (FPL '07)*. IEEE, Amsterdam, Netherlands, 151–158.
- [109] Kai Nagel and Marcus Rickert. 2001. Parallel Implementation of the TRANSIMS Micro-Simulation. *Elsevier Journal of Parallel Computing* 27, 12 (November 2001), 1611–1639.
- [110] Razvan Nane, Vlad-Mihai Sima, Christian Pilato, Jongsok Choi, Blair Fort, Andrew Canis, Yu Ting Chen, Hsuan Hsiao, Stephen Brown, Fabrizio Ferrandi, Jason Anderson, and Koen Bertels. 2016. A Survey and Evaluation of FPGA High-Level Synthesis Tools. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 35, 10

- (October 2016), 1591–1604.
- [111] Roland Neumann and Fritz Strack. 2000. 'Mood Contagion': The Automatic Transfer of Mood Between Persons. *American Psychological Association Journal of Personality and Social Psychology* 79, 2 (August 2000), 211.
- [112] Michael J. North and Charles M. Macal. 2007. *Managing Business Complexity: Discovering Strategic Solutions With Agent-Based Modeling and Simulation*. Oxford University Press.
- [113] NVIDIA Corporation. 2012. *CUDA Toolkit 4.2 – CUBLAS Library*. Technical Report PG-05326-041_v01. Santa Clara, CA, USA.
- [114] NVIDIA Corporation. 2016. Whitepaper – NVIDIA GeForce GTX 1080 – Gaming Perfected. (2016).
- [115] NVIDIA Corporation. 2016. *Whitepaper – NVIDIA Tesla P100 – The Most Advanced Datacenter Accelerator Ever Built Featuring Pascal GP100, the World's Fastest GPU*. Technical Report WP-08019-001_v01.1.
- [116] NVIDIA Corporation. 2017. *NVIDIA TESLA V100 GPU Architecture – The World's Most Advanced Data Center GPU*. Technical Report WP-08608-001_v1.1.
- [117] NVIDIA Corporation. 2018. *NVIDIA CUDA C Programming Guide*. Version 9.1.85, NVIDIA Corporation.
- [118] Christopher Oat, Joshua Barczak, and Jeremy Shopf. 2009. *Efficient Spatial Binning on the GPU*. Technical Report. Advanced Micro Devices, Inc.
- [119] Andreas Olofsson. 2016. Epiphany-V: A 1024 processor 64-bit RISC System-On-Chip. (October 2016). arXiv:CoRR/1610.01832
- [120] OpenACC Working Group. 2015. *The OpenACC Application Programming Interface*. Technical Report Version 2.5.
- [121] Samir Palnitkar. 2003. *Verilog®Hdl: A Guide to Digital Design and Synthesis, Second Edition*. Prentice Hall Press.
- [122] Hyungwook Park and Paul A. Fishwick. 2008. A Fast Hybrid Time-Synchronous/Event Approach to Parallel Discrete Event Simulation of Queuing Networks. In *Proceedings of the Winter Simulation Conference (WSC '08)*. IEEE, Miami, FL, USA, 795–803.
- [123] Hyungwook Park and Paul A. Fishwick. 2010. A GPU-Based Application Framework Supporting Fast Discrete-Event Simulation. *SCSI Simulation* 86, 10 (October 2010), 613–628.
- [124] Hyungwook Park and Paul A. Fishwick. 2011. An Analysis of Queuing Network Simulation Using GPU-Based Hardware Acceleration. *ACM Transactions on Modeling and Computer Simulation* 21, 3 (March 2011), 18:1–18:22.
- [125] Roman Pavlov and Jörg P. Müller. 2013. Multi-Agent Systems Meet GPU: Deploying Agent-Based Architectures on Graphics Processors. In *Proceedings of the Doctoral Conference on Computing, Electrical and Industrial Systems (DoCEIS '13)*. Springer, Costa de Caparica, Portugal, 115–122.
- [126] Kalyan S. Perumalla. 2006. Discrete-Event Execution Alternatives on General Purpose Graphical Processing Units (GPGPUs). In *Proceedings of the Workshop on Principles of Advanced and Distributed Simulation (PADS '06)*. IEEE, Singapore, 74–81.
- [127] Kalyan S. Perumalla. 2008. Efficient Execution on GPUs of Field-Based Vehicular Mobility Models. In *Proceedings of the Workshop on Principles of Advanced and Distributed Simulation (PADS '08)*. IEEE, Roma, Italy, 154–154.
- [128] Kalyan S. Perumalla and Brandon G. Aaby. 2008. Data Parallel Execution Challenges and Runtime Performance of Agent Simulations on GPUs. In *Proceedings of the Spring Simulation Multiconference (SpringSim '08)*. SCSi, Ottawa, Canada, 116–123.
- [129] Kalyan S. Perumalla, Brandon G. Aaby, Srikanth B. Yognath, and Sudip K. Seal. 2009. GPU-Based Real-Time Execution of Vehicular Mobility Models in Large-Scale Road Network Scenarios. In *Proceedings of the ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation (PADS '09)*. IEEE, Lake Placid, NY, USA, 95–103.
- [130] Louis-Noël Pouchet. 2012. Polybench: the Polyhedral Benchmark suite. (2012). Retrieved May 11, 2018 from <http://web.cs.ucla.edu/~pouchet/software/polybench/>
- [131] Sebastian Raase and Tomas Nordström. 2015. On the Use of a Many-Core Processor for Computational Fluid Dynamics Simulations. *Proceedings of the International Conference On Computational Science (ICCS '15)*, 1403–1412.
- [132] Shafiqur Rahman, Nael Abu-Ghazaleh, and Walid Najjar. 2017. PDES-A: A Parallel Discrete Event Simulation Accelerator for FPGAs. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '17)*. ACM, Singapore, 133–144.
- [133] Ashu Rege. 2008. An Introduction to Modern GPU Architecture. (2008).
- [134] Paul F. Reynolds, Carmen M. Pancerella, and Sudhir Srinivasan. 1993. Design and Performance Analysis of Hardware Support for Parallel Simulations. *Elsevier Journal of Parallel and Distributed Computing* 18, 4 (August 1993), 435–453.
- [135] Paul Richmond, Simon Coakley, and Daniela Romano. 2009. Cellular Level Agent Based Modelling on the Graphics Processing Unit. In *Proceedings of the International Workshop on High Performance Computational Systems Biology (HIBI '09)*. IEEE, Trento, Italy, 43–50.
- [136] Paul Richmond and Daniela Romano. 2008. Agent Based GPU, a Real-time 3D Simulation and Interactive Visualisation Framework for Massive Agent Based Modelling on the GPU. In *Proceedings of the International Workshop on Supervisualisation (IWSV '08)*. Island of Kos, Aegean Sea, Greece.

- [137] Paul Richmond and Daniela Romano. 2011. Template-Driven Agent-Based Modeling and Simulation With CUDA. *Elsevier GPU Computing Gems Emerald Edition, Applications of GPU Computing Series* (February 2011), 313–324.
- [138] Patrick F. Riley and George F. Riley. 2003. Next Generation Modeling III-Agents: Spades—a Distributed Agent Simulation Environment With Software-in-the-Loop Execution. In *Proceedings of the Conference on Winter Simulation (WSC '03)*. IEEE, New Orleans, LA, USA, 817–825.
- [139] Robert Rönnngren and Rassul Ayani. 1997. A Comparative Study of Parallel and Sequential Priority Queue Algorithms. *ACM Transactions on Modeling and Computer Simulation* 7, 2 (April 1997), 157–209.
- [140] Janche Sang, Che-Rung Lee, Vernon Rego, and Chung-Ta King. 2013. A Fast Implementation of Parallel Discrete-Event Simulation on GPGPU. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '13)*. Las Vegas, NV, US, 501.
- [141] Thomas C. Schelling. 2006. *Micromotives and Macrobehavior*. WW Norton & Company.
- [142] Moon Gi Seok and Tag Gon Kim. 2012. Parallel Discrete Event Simulation for DEVS Cellular Models Using a GPU. In *Proceedings of the Symposium on High Performance Computing (HPC '12)*. SCSi, Orlando, FL, USA, 11:1–11:7.
- [143] Zhen Shen, Kai Wang, and Fenghua Zhu. 2011. Agent-Based Traffic Simulation and Traffic Signal Timing Optimization With GPU. In *Proceedings of the International IEEE Conference on Intelligent Transportation Systems (ITSC '11)*. IEEE, Washington, DC, USA, 145–150.
- [144] Xuan Shi and Fei Ye. 2013. Kriging Interpolation Over Heterogeneous Computer Architectures and Systems. *Taylor & Francis Journal of GIScience & Remote Sensing* 50, 2 (2013), 196–211.
- [145] Jaewoong Sim, Aniruddha Dasgupta, Hyesoon Kim, and Richard Vuduc. 2012. A Performance Analysis Framework for Identifying Potential Benefits in GPGPU Applications. In *Proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming (PPoPP '12)*. ACM, New Orleans, LA, USA, 11–22.
- [146] Daniel Dominic Sleator and Robert Endre Tarjan. 1985. Self-Adjusting Binary Search Trees. *J. ACM* 32, 3 (July 1985), 652–686.
- [147] Avinash Sodani, Roger Gramunt, Jesus Corbal, Ho-Seop Kim, Krishna Vinod, Sundaram Chinthamani, Steven Hutsell, Rajat Agarwal, and Yen-Chen Liu. 2016. Knights Landing: Second-Generation Intel Xeon Phi Product. *IEEE Micro* 36, 2 (April 2016), 34–46.
- [148] Xiao Song, Ziping Xie, Yan Xu, Gary Tan, Wenjie Tang, Jing Bi, and Xiaosong Li. 2017. Supporting Real-World Network-Oriented Mesoscopic Traffic Simulation on GPU. *Elsevier Journal of Simulation Modelling Practice and Theory* 74 (May 2017), 46–63.
- [149] Russell K. Standish and Richard Leow. 2004. EcoLab: Agent based modeling for C++ programmers. (January 2004). arXiv:CoRR/0401026
- [150] John E. Stone, David Gohara, and Guochun Shi. 2010. OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems. *IEEE Computing in Science & Engineering* 12, 3 (May 2010), 66–73.
- [151] David Strippen and Kai Nagel. 2009. Using Common Graphics Hardware for Multi-Agent Traffic Simulation With CUDA. In *Proceedings of the International Conference on Simulation Tools and Techniques (Simutools '09)*. ICST, Rome, Italy, 62:1–62:8.
- [152] Herb Sutter. 2005. The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software. *Dr. Dobbs' Journal* 30, 3 (March 2005), 202–210.
- [153] Brian Paul Swenson. 2015. *Techniques to Improve the Performance of Large-Scale Discrete-Event Simulation*. Dissertation. Georgia Institute of Technology.
- [154] G. Szemes, L. Gulyás, G. Kampis, and W. de Back. 2010. GridABM-templates for Distributed Agent Based Simulation. In *Open Grid Forum*, Vol. 28. Munich, Germany.
- [155] Wenjie Tang and Yiping Yao. 2013. A GPU-Based Discrete Event Simulation Kernel. *SCSI Journal of Simulation* 89, 11 (October 2013), 1335–1354.
- [156] Kardi Teknomo, Yasushi Takeyama, and Hajime Inamura. 2016. Review on Microscopic Pedestrian Simulation Model. (September 2016). arXiv:CoRR/1609.01808
- [157] Leigh Tesfatsion. 2006. Agent-Based Computational Economics: A Constructive Approach to Economic Theory. *Elsevier Handbook of Computational Economics* 2 (May 2006), 831–880.
- [158] Y. Torres, A. Gonzalez-Escribano, and D.R. Llanos. 2011. Understanding the Impact of CUDA Tuning Techniques for Fermi. In *Proceedings of the International Conference on High Performance Computing and Simulation (HPCS '11)*. IEEE, Istanbul, Turkey, 631–639.
- [159] Justin L. Tripp, Henning S. Mortveit, Anders A. Hansson, and Maya Gokhale. 2005. Metropolitan Road Traffic Simulation on FPGAs. In *Proceedings of the Annual IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM '05)*. IEEE, Napa, CA, USA, 117–126.
- [160] Mario Vestias and Horácio Neto. 2014. TRENDS OF CPU, GPU AND FPGA FOR HIGH-PERFORMANCE COMPUTING. In *Proceedings of the International Conference on Field Programmable Logic and Applications (FPL '14)*. IEEE, Munich, Germany, 1–6.

- [161] Guillermo Viguera, Juan M. Orduña, Miguel Lozano, José M. Cecilia, and José M. García. 2014. Accelerating Collision Detection for Large-Scale Crowd Simulation on Multi-Core and Many-Core Architectures. *Sage Publications The International Journal of High Performance Computing Applications* 28, 1 (February 2014), 33–49.
- [162] Ioannis Vourkas and Georgios Ch. Sirakoulis. 2012. FPGA Based Cellular Automata for Environmental Modeling. In *Proceedings of the International Conference on Electronics, Circuits and Systems (ICECS '12)*. IEEE, Seville, Spain, 93–96.
- [163] Jin Wang, Norman Rubin, Haicheng Wu, and Sudhakar Yalamanchili. 2013. Accelerating Simulation of Agent-Based Models on Heterogeneous Architectures. In *Proceedings of the Workshop on General Purpose Processor Using Graphics Processing Units (GPGPU '06)*. ACM, Houston, TX, USA, 108–119.
- [164] Kai Wang and Zhen Shen. 2012. A GPU Based Trafficparallel Simulation Module of Artificial Transportation Systems. In *Proceedings of the IEEE International Conference on Service Operations and Logistics, and Informatics (SOLI '12)*. IEEE, Suzhou, China, 160–165.
- [165] Yuan Wen, Zheng Wang, and Michael F.P. O'Boyle. 2014. Smart Multi-Task Scheduling for OpenCL Programs on CPU/GPU Heterogeneous Platforms. In *Proceedings of the International Conference on High Performance Computing (HiPC '14)*. IEEE, Dona Paula, India, 1–10.
- [166] Tang Wenjie, Yao Yiping, and Zhu Feng. 2013. An Expansion-Aided Synchronous Conservative Time Management Algorithm on GPU. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '13)*. ACM, Montreal, Canada, 367–372.
- [167] Barry Williams, Dmitry Ponomarev, Nael Abu-Ghazaleh, and Philip Wilsey. 2017. Performance Characterization of Parallel Discrete Event Simulation on Knights Landing Processor. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '17)*. ACM, Singapore, Singapore, 121–132.
- [168] Was, Jaroslaw and Mróz, Hubert and Topa, Paweł. 2016. GPGPU Computing for Microscopic Simulations of Crowd Dynamics. *Slovak Academy of Sciences Computing and Informatics* 34, 6 (February 2016), 1418–1434.
- [169] Fulong Wu. 2002. Calibration of Stochastic Cellular Automata: The Application to Rural-Urban Land Conversions. *Taylor & Francis International Journal of Geographical Information Science* 16, 8 (2002), 795–818.
- [170] Xilinx. 2017. XilinxVirtex UltraScale+ Product Table. (December 2017). Retrieved May 11, 2018 from <https://www.xilinx.com/products/silicon-devices/fpga/virtex-ultrascale-plus.html#productTable>
- [171] Yadong Xu, Wentong Cai, David Eckhoff, Suraj Nair, and Alois Knoll. 2017. A Graph Partitioning Algorithm for Parallel Agent-Based Road Traffic Simulation. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '17)*. ACM, Singapore, 209–219.
- [172] Yan Xu, Gary Tan, Xiaosong Li, and Xiao Song. 2014. Mesoscopic traffic simulation on CPU/GPU. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '14)*. ACM, Denver, CO, USA, 39–50.
- [173] Yonghong Yan, Max Grossman, and Vivek Sarkar. 2009. JCUDA: A Programmer-Friendly Interface for Accelerating Java Programs With CUDA. In *Proceedings of the European Conference on Parallel Processing (Euro-Par '09)*. Springer, Delft, The Netherlands, 887–899.
- [174] Mingyu Yang, Philipp Andelfinger, Wentong Cai, and Alois Knoll. 2018. Evaluation of Conflict Resolution Methods for Agent-Based Simulation on the GPU. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS '18)*. ACM, Rome, Italy, 129–132.
- [175] Srikanth B. Yogninath and Kalyan S. Perumalla. 2018. Scalable Cloning on Large-Scale GPU Platforms With Application to Time-Stepped Simulations on Grids. *ACM Transactions on Modeling and Computer Simulation* 28, 1 (January 2018), 5:11–5:26.
- [176] Feng Zhang, Jidong Zhai, Bingsheng He, Shuhao Zhang, and Wenguang Chen. 2017. Understanding Co-Running Behaviors on Integrated CPU/GPU Architectures. *IEEE Transactions on Parallel and Distributed Systems* 28, 3 (March 2017), 905–918.
- [177] Tao Zhang, Xin Zhao, Xinqi An, Haojun Quan, and Zhichun Lei. 2017. Using Blind Optimization Algorithm for Hardware/Software Partitioning. *IEEE Access* 5 (February 2017), 1353–1362.
- [178] Yao Zhang and John D Owens. 2011. A Quantitative Performance Analysis Model for GPU Architectures. In *Proceedings of the 2011 IEEE International Symposium on High Performance Computer Architecture (HPCA '11)*. IEEE, San Antonio, TX, USA, 382–393.
- [179] Li Zhen, Qiuxiao Gang, Guo Gang, and Chen Bin. 2014. A GPU-Based Simulation Kernel within Heterogeneous Collaborative Computation on Large-Scale Artificial Society. *IACSIT Press International Journal of Modeling and Optimization* 4, 3 (Jun 2014), 205–210.
- [180] Hamid Reza Zohouri, Naoya Maruyama, Aaron Smith, Motohiko Matsuda, and Satoshi Matsuoka. 2016. Evaluating and Optimizing OpenCL Kernels for High Performance Computing With FPGAs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '16)*. IEEE, Salt Lake City, UT, USA, 35:1–35:12.
- [181] Peng Zou, Ya-shuai Lü, Li-li Chen, and Yi-ping Yao. 2013. Epidemic Simulation of Large-Scale Social Contact Network on GPU Clusters. *Sage Publications Journal of Simulation* 89, 10 (2013), 1154–1172.