

Chapter 6

Introduction to Dataflow Computing

Nenad Korolija

School of Electrical Engineering, University of Belgrade, Serbia

Jovan Popović

Microsoft Development Center, Serbia

Miroslav M. Bojović

School of Electrical Engineering, University of Belgrade, Serbia

ABSTRACT

This chapter presents the possibilities for obtaining significant performance gains based on advanced implementations of algorithms using the dataflow hardware. A framework built on top of the dataflow architecture that provides tools for advanced implementations is also described. In particular, the authors point out to the following issues of interest for accelerating algorithms: (1) the dataflow paradigm appears as suitable for executing certain set of algorithms for high performance computing, namely algorithms that work with big data, as well as algorithms that include a lot of repetitions of the same set of instructions; (2) dataflow architecture could be configured using appropriate programming tools that can define hardware by generating VHDL files; (3) besides accelerating algorithms, dataflow architecture also reduces power consumption, which is an important security factor with edge computing.

BACKGROUND

The 1906 Nobel Prize in Physiology and Medicine was awarded to Camillo Golgi and Ramón y Cajal for having visualized and identified the neuron, the structural and functional unit of the nervous system (Grant, 2007). Since then, it has been discovered that the human brain contains roughly 100 billion neurons and 1000 trillion synapses. Neurons interact through electrochemical signals, also known as action potentials (AP) or spikes, transmitted from one neuron to the next through synaptic junctions, forming functional and definable circuits which can be organized into larger ‘neuronal’ networks and anatomical

DOI: 10.4018/978-1-7998-7156-9.ch006

structures. These networks integrate information from multiple brain regions as well as incoming information about the external environment (e.g., sound, light, smell, taste). The result is how we perceive the world, and produce complex behavior and cognitive processes including decision-making and learning (Kandel, 2012); also, with time, these processes modify the structure and function of networks through a process called neuroplasticity (Fuchs & Flugge, 2014).

Understanding how the brain works with the ultimate goal of developing treatments for neurological disease remains one of the greatest scientific challenges of this century. In fact, there is a substantial social and economic burden associated with neurological diseases (Wynford-Thomas & Robertson, 2017). In the US alone, the overall cost of neurological diseases (e.g., stroke, dementia, movement disorders, traumatic brain injury) amounts to nearly \$1T, and will dramatically increase in the next few years due to population ageing. Alarming, the cost of just dementias and stroke is expected to exceed \$600B by 2030 (Gooch, Pracht, & Borenstein, 2017). To tackle this challenge, neuroscientists have developed a battery of increasingly complex tools, which have amplified data storage and computational speed requirements to an unprecedented level, making the use of big data techniques and HPC such as supercomputers a necessity.

In the remainder of this chapter we will briefly review current and future applications of supercomputers in neuroscience, with a focus on computational neural models, brain imaging, and models for brain stimulation. These areas were chosen not only because their advances have been particularly driven by computational approaches, but also because they are highly interconnected. Thereby, understanding how each area is evolving aids prediction of future research trends in the other areas. We also briefly discuss how the next generation of supercomputers might enable further advancements in these areas.

Computational Neural Models

Computational models of neurons and neural networks represent one of the most essential tools that have contributed to the progress of neuroscience. For example, they are used to guide the design of experiments, to quantify relationships between anatomical and physiological data, to investigate the dynamics of systems that cannot be accessed via analytical methods, and to validate estimates made during theoretical derivations. Since the early days of neuronal simulations, a wide range of computational models have been developed ranging from models aimed at describing low-level mechanisms of neural function (e.g., molecular dynamics of ion channels in the neuron) to models of large-scale neuronal networks (Ippen, Eppler, Plesser, & Diesmann, 2017) (see (Fan & Markram, 2019) for a review).

The chosen level of abstraction for a model is based on the scientific question. If focused on sub-cellular processes (e.g., the transfer of ions underlying changes in membrane voltages that lead to APs), a neuron(s) would be described with detailed multi-compartment models (M. Hines, 1984). Instead, if the question addressed large scale network dynamics, many neurons would be described with one-compartment or few-compartment models that communicate electrically via spikes (Helias et al., 2012; Ippen et al., 2017). Simulators exist for many of these levels, including NEURON (M. L. Hines & Carnevale, 1997), SPLIT (Hammarlund & Ekeberg, 1998), PCSIM (Pecevski, Natschläger, & Schuch, 2009), the NEural Simulation Tool (NEST) (Gewaltig, 2007; van Albada, Kunkel, Morrison, & Diesmann, 2014), and C2 (Ananthanarayanan & Modha, 2007) (see (Helias et al., 2012; Tikidji-Hamburyan, Narayana, Bozkus, & El-Ghazawi, 2017) for a review).

A major challenge when developing such models include the total number of network elements that must be represented. To put the problem into perspective, a 1 mm³ of brain tissue modeled at the neuronal

8 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/introduction-to-dataflow-computing/273396

Related Content

Rough Entropy Clustering Algorithm in Image Segmentation

Dariusz Malyszko and Jarosław Stepaniuk (2010). *Novel Developments in Granular Computing: Applications for Advanced Human Reasoning and Soft Computation* (pp. 285-304).

www.irma-international.org/chapter/rough-entropy-clustering-algorithm-image/44708

Supervised and Unsupervised Information Granulation: A Study in Hyperbox Design

Andrzej Bargiela and Witold Pedrycz (2010). *Novel Developments in Granular Computing: Applications for Advanced Human Reasoning and Soft Computation* (pp. 48-68).

www.irma-international.org/chapter/supervised-unsupervised-information-granulation/44699

Resource Management

Valentin Cristea, Ciprian Dobre, Corina Stratan and Florin Pop (2010). *Large-Scale Distributed Computing and Applications: Models and Trends* (pp. 75-90).

www.irma-international.org/chapter/resource-management/43103

A Grid-Aware Emergency Response Model (G-AERM) for Disaster Management

Eleana Asimakopoulou, Chimay J. Anumba and Dino Bouchlaghem (2010). *International Journal of Distributed Systems and Technologies* (pp. 40-55).

www.irma-international.org/article/grid-aware-emergency-response-model/46049

Internet of Things Applications: Current and Future Development

Mahmoud Elkhodr, Seyed Shahrestani and Hon Cheung (2016). *Innovative Research and Applications in Next-Generation High Performance Computing* (pp. 397-427).

www.irma-international.org/chapter/internet-of-things-applications/159055