

BIO-MIMETIC CLASSIFICATION ON MODERN PARALLEL HARDWARE: REALIZATIONS ON NVIDIA® CUDA™ AND OPENMP™

THOMAS NOWOTNY¹, MEHMET K. MUEZZINOGLU² AND RAMON HUERTA²

¹School of Informatics
University of Sussex
Brighton, BN1 9QJ, United Kingdom
t.nowotny@sussex.ac.uk

²BioCircuits Institute
University of California, San Diego
La Jolla, CA 92093-0402, USA
mmuezzin@ucsd.edu; rhuerta@ucsd.edu

Received February 2010; revised June 2010

ABSTRACT. *Both the brain and modern digital architectures rely on massive parallelism for efficient solutions to demanding computational tasks, such as pattern recognition. In this paper, we implement a parallel classification scheme inspired by the insect brain in two popular parallel computing frameworks, namely as an NVidia® CUDA™ implementation on a Tesla™ device and a brute force OpenMP™ parallel implementation on a quad-core CPU. When evaluating the systems on the MNIST data-set of handwritten digits, we can report that, compared with a standard serial implementation on a single CPU core, CUDA™ implementations of the bio-inspired classification provide a 7-to-11 fold speed-up, whereas the OpenMP™ implementation is 2-to-4 times faster. Our results are a proof of concept that suggests that modern parallel computing architectures and bio-mimetic algorithms are compatible and that the CUDA™ solution on an NVidia® Tesla™ C870 device at the time of writing has a small edge over an OpenMP solution on a recent quad core processor (3 GHz AMD® Phenom™ II X4 940).*

Keywords: Parallel hardware, Graphical processing unit, CUDA, OpenMP, Bio-mimetic systems, Classification, Insect brain, Olfactory system, MNIST data set, Multi-layer perceptron, Random connections

1. Introduction. The remarkable speed and accuracy of information processing in nature has been a strong motivation for computing research. A key characteristic of biological computing is architectural: the brain is organized in layers of neurons processing information in a highly parallel fashion. Bio-inspired algorithms, therefore, should naturally have the right structure to take advantage of highly parallel computing architectures. Although one can, in principle, implement such methods on conventional desktop computers in a serial manner, their full benefits can only be revealed on suitable parallel devices.

The practice of parallel computing dates back to the late 50s and early 60s, with the first parallel computer arguably being the D825 of Burroughs Corporation featuring 4 processors. Already in the early 80s, massively parallel computers like the “Connection Machine” with – depending on configuration – up to 65,536 processors were built. However, only recently, the domain of parallel computing has begun to expand from exotic super-computers towards the broader market in the form of multi-core processors, first