

Preface

Evolutionary algorithms (EAs) are generic heuristics that learn from natural collective behavior and can be applied to solve very difficult optimization problems. Applications include engineering problems, scheduling problems, routing problems, assignment problems, finance and investment problems, analyzing problems of complex social behavior, to state a few.

One of the most important features of EAs is that they are population-based searching algorithms, i.e., they work using one or more sets of candidate solutions (individuals). As a result, EAs deal well with the EvE (Exploration vs. Exploitation) paradox and are ideally suited for massively parallel runs, at the operator level, individual level, and population level and can therefore exploit efficiently general-purpose graphics processing units (GPGPUs) as well as massively parallel supercomputers. Because EAs are well suited to be applied to complex problems, they may require a long time to obtain good solutions or may require a large amount of computational resources to obtain high-quality solutions in a given time. Thus, since the earliest studies, researchers have attempted to use parallel EAs in order to accelerate the execution of EAs.

Many of the traditional parallel EAs run on multi-core machines, massively parallel cluster machines, or grid computing environments. However, recent advances in scientific computing have made it possible to use GPGPUs for parallel EAs. GPGPUs are low-cost, massively parallel, many-core processors. They can execute thousands of threads in parallel in single-program multiple-data (SPMD) manner. With the low-cost GPGPUs found in ordinary PCs, it is becoming possible to use parallel EAs to solve optimization problems of all sizes.

Decreasing execution time by orders of magnitude opens new possibilities for many researchers and EA users who were until now limited by the limited computing power of CPUs. At the time of writing, current GPGPU cards computing power are in the order of 10 TFlops, turning any PC into a high-performance parallel processing computer. Essentially, this provides anyone with the processing power previously available only to those with access to expensive supercomputers. Not only is there benefit to solve large problems, but also many small problems can be run many times in order to optimize parameter tuning.

The impact on optimization and inverse problem solving should be huge and this book should be of interest to anyone involved with EAs, including researchers, scientists, professionals, teachers, and students in the field.

Although the parallelism of EAs is well suited to SPMD-based GPGPUs, there are many issues to be resolved in the design and implementation of EAs on GPUs. For example, EAs run with a certain degree of randomness which can bring divergence that could degrade performance on SIMD cores. Some chapters of this book propose solutions to these kinds of issues in solving large-scale optimization problems.

This book provides a collection of 19 chapters ranging from general tutorials to state-of-the-art parallel EAs on GPUs to real-world applications that will hopefully help those who would like to parallelize their work on GPGPU cards.

The three chapters of Part I (Tutorials) introduce the problem, and present the hardware specificities of GPGPU cards as well as an easy way to start using GPGPUs for evolutionary computation, using the EASEA platform that was developed in Strasbourg.

Part II presents more advanced implementations of various kinds of EAs on GPGPUs, in order to show what can be done and serve as inspiration for new developments.

Part III is more concerned with real-world applications, showing that evolutionary GPGPU computing has come of age and is not confined to solving toy problems.

We would like to end this preface with our deepest thanks to all authors for their effort, help, and their pioneering contributions to the topics they have written about. Thanks to their commitment, we believe that this volume can have a great impact on our respective fields. Also, we would like to thank Prof. A. E. Eiben and Prof. Hans-Paul Schwefel for encouraging the publication of this volume, and Ronan Nugent of Springer for his continuous help throughout the publishing stages of this book.

Osaka, Japan
Strasbourg, France
January 2013

Shigeyoshi Tsutsui
Pierre Collet

Massively Parallel Evolutionary Computation on GPGPUs

Tsutsui, S.; Collet, P. (Eds.)

2013, XII, 453 p. 199 illus., 95 illus. in color., Hardcover

ISBN: 978-3-642-37958-1