Thesis Overview:

Performance analysis and optimization of parallel Best-First Search algorithms on multicore and cluster of multicore.

Victoria María Sanz.
School of Computer Science, National University of La Plata, Argentina
PhD Thesis in Computer Science
Advisors: Armando E. De Giusti, Marcelo Naiouf
{vsanz, degiusti, mnaiouf}@lidi.info.unlp.edu.ar

Since the emergence of clusters, multicore processors and clusters of multicore machines, researchers and developers have faced the challenge of parallelizing different kinds of applications in order to take advantage of the computing power and/or the accumulated memory that these architectures provide.

In the area of Artificial Intelligence, heuristic search algorithms are used as the basis to solve combinatorial optimization problems, such as: optimal route planning, robot navigation, optimal sequence alignments, among others (Russel & Norvig, 2003). One of the most widely used heuristic search algorithms for that purpose is A* (Hart, et al., 1968), a variant of *Best-First Search,* which requires high computing power and a large amount of memory. Consequently, during the last decades, the development of the parallel *Best-First Search* algorithms has been promoted which, in particular, may benefit from the parallel architectures mentioned formerly.

The contribution of the thesis is the development of two parallel *Best-First Search* algorithms, one that is suitable for execution on shared-memory machines (multicore), and another one that is suitable for execution on distributed memory machines (cluster). The former is based on the adaptation of the HDA* (*Hash Distributed A*) algorithm for multicore machines proposed by (Burns et al., 2010), while the latter is based on the HDA* (*Hash Distributed A*)* algorithm proposed by (Kishimoto, et al., 2013). The implemented algorithms incorporate parameters and/or techniques that improve their performance, with respect to the original algorithms proposed by the authors mentioned above.

Additionally, a comparison of the performance achieved and the memory consumed by both algorithms, when they are run on a multicore machine, is presented. These results show that a benefit would be obtained by converting HDA* into a hybrid application, that uses programming tools for shared and distributed memory, when the underlying architecture is a cluster of multicore, and therefore the bases for this hybrid algorithm are established.

This thesis (Sanz, 2015) is organized into nine chapters.

Chapter 1 presents an introduction to the topic of the thesis, its motivation, the main objective, and the contributions.

Chapter 2 describes the steps to follow in order to transform a real problem into a search problem. Also, it explores the most used sequential search algorithms and compares them so as to select the most suitable to solve the study case.

Chapter 3 provides a characterization of the *15-puzzle* problem, selected as study case, and specifies an algorithm for determining whether a final configuration of the problem can be reached from a given initial configuration. Then, it focuses on describing some heuristics developed for the *15-puzzle* problem by different authors, and mentions real-world problems that can be solved with the algorithms presented in this thesis.

Chapter 4 presents a classification of parallel architectures, describes modern parallel platforms (clusters and multicore machines), introduces the most commonly used parallel programming tools and discusses the challenges the programmer must face in order to take advantage of these architectures.

Chapter 5 introduces the reader to the basic metrics used to assess the performance of parallel search algorithms and to the common sources of overhead in parallel programs. Then, the chapter provides an overview of previous research on parallel search and focuses on the search of optimal solutions through A*. Additionally, the chapter presents simple strategies to detect termination of a distributed computation, analyses different dynamic memory allocators and considers possible improvements to existing parallel search algorithms.

Chapter 6 covers the implementation of the algorithms proposed in this thesis that allow solving the study case: the Sequential A* algorithm, the optimized HDA* algorithm for distributed-memory architectures (HDA* MPI), and the optimized HDA* algorithm for shared-memory architectures (HDA* Pthreads).

Chapter 7 describes the experimental work carried out and analyses the results obtained by running the algorithms proposed in Chapter 6 on different platforms. Firstly, Sequential A* was run using different dynamic memory allocators and heuristics; this test allows determining the configuration that is more beneficial to performance. Secondly, HDA* Pthreads was run on a multicore machine and HDA* MPI was run on a multicore machine and on a cluster of multicore machines; in each case, the impact on performance of the proposed techniques and parameter values is documented. The experimental results obtained by both algorithms show that they scale well on the mentioned architectures. Finally, the performance achieved and memory consumed by HDA* Pthreads and HDA* MPI, when they are run on multicore machine, is contrasted; the results show that potential benefits would be obtained by converting HDA* into a hybrid application, when the underlying architecture is a cluster of multicore machines.

Chapter 8 covers the hybrid programming models and establishes a base for the hybrid HDA* algorithm.

Chapter 9 summarizes the obtained results and discusses future lines of work.

Victoria María Sanz
vsanz@lidi.info.unlp.edu.ar

References

Burns, E., Lemons, S., Ruml, W. & Zhou, R., 2010. Best-First Heuristic Search for Multicore Machines. *Journal of Artificial Intelligence Research,* Vol. 39, pp. 689-743.

Hart, P., Nilsson, N. & Raphael, B., 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics,* 4(2), pp. 100-107.

Kishimoto, A., Fukunaga, A. & Botea, A., 2013. Evaluation of a simple, scalable, parallel Best-First Search strategy. *Artificial Intelligence,* Vol. 195, pp. 222–248.

Russel, S. & Norvig, P., 2003. *Artificial Intelligence: A Modern Approach.* Second edition. NJ: Prentice Hall.

Sanz, V. M., 2015. *Análisis de rendimiento y optimización de algoritmos paralelos Best-First Search sobre multicore y cluster de multicore.*
Available at: http://sedici.unlp.edu.ar/handle/10915/44478