



Título: GPU-Based Backtracking Strategies for Solving Permutation Combinatorial Problems

Data: 27/09/2017 Horário: 08:30h Local: Sala de Seminários - Bloco 952 - Campus do Pici

Resumo:

New GPGPU technologies, such as CUDA Dynamic Parallelism (CDP), can help dealing with recursive patterns of computation, such as divide-and-conquer, used by backtracking algorithms. In this paper, we propose a GPU-accelerated backtracking algorithm using CDP that extends a well-known parallel backtracking model. The search starts on CPU processing the search tree until a first cutoff depth. Based on this partial backtracking tree, the algorithm performs an analysis of the memory required by the subsequent kernel generations. The proposed algorithm performs no dynamic allocation of memory on GPU, unlike related works from the literature. The required memory is preallocated based on an analysis of the partial backtracking tree. The proposed algorithm has been extensively tested using the N-Queens Puzzle problem and instances of the Asymmetric Traveling Salesman Problem (ATSP) as test-cases. The proposed CDP algorithm may, under some conditions, outperform its non-CDP counterpart by a factor up to 25. But, it may also be up to twice slower. The CDP-based implementation has much better worst case execution times and makes algorithm's performance less dependent on the tuning of parameters. Compared to other CDP-based strategies from the literature, the proposed algorithm is in average 8 times faster. We have identified that some works in the literature cannot solve problems with huge memory requirement. The main reason is that they do not provide means to setup CUDA runtime variables. In the second part of this work, we propose a set of algorithms to calculate and setup dynamically the CUDA runtime variables. All CDP-based algorithms mentioned in this work can be implemented without using CDP, invoking CUDA kernels from the host. We present two

implementations which have the semantics of the CDP-based strategy. We show that a smaller interference of the host combined with a block-based child search seems worthwhile for irregular tree search algorithms. On the other hand, in situations where the load processed by the child kernels is small, the iterative non-CDP version is a better option. We also identify some difficulties, limitations, and bottlenecks concerning the CDP programming model which may be useful for helping potential users.

Banca:

- Prof. Dr. Francisco Heron de Carvalho Junior (MDCC/UFC - Orientador)
- Prof. Dr. Igor Machado Coelho (UERJ)
- Prof. Dr. Albert Einstein Fernandes Muritiba (UFC)
- Prof. Dr. Rafael Castro de Andrade (UFC)