
GPU parallelization of ALNS for the DCVRP

Lukas Bach^{*1}, Geir Hasle¹, and Christian Schulz¹

¹SINTEF – P.O. Box 124 Blindern, 0314 Oslo, Norway

Abstract

In recent years the computational capacity of the Graphics Processing Unit (GPU) in ordinary desktop computers has increased significantly compared to the Central Processing Unit (CPU). It is interesting to explore how this alternative source of computing power can be utilized. Most investigations of GPU-based solution methods in discrete optimization are based on swarm intelligence or evolutionary algorithms.

One of the best single-solution metaheuristics for discrete optimization is Adaptive Large Neighborhood Search (ALNS). GPU parallelization of ALNS has not been reported in the literature. We investigate ALNS on the GPU by developing a data parallel ALNS algorithm for the Distance-constrained Capacitated Vehicle Routing Problem (DCVRP). To achieve good utilization of the GPU, it was necessary to adapt the set of destroy and repair operators of a state-of-the-art CPU implementation.

The data parallel ALNS is implemented in NVIDIA CUDA. The Performance of a state-of-the-art CPU implementation and our GPU version is compared experimentally on an Intel Core i7-4770K with 3.5 GHz and an NVIDIA GeForce GTX TITAN. We use three standard DCVRP benchmarks: the Christofides, Mingozzi, and Toth instances, the Kelly instances, and the Li, Golden, and Wasil instances – in total a set of 46 instances with the number of customers ranging from 50 to 1200. On average, our GPU implementation of ALNS yields competitive solution quality with less runtime than the CPU implementation. However, on larger instances it is easier to utilize the parallelism of the GPU and achieve both improved solution quality and considerably improved runtime.

^{*}Speaker