

Abstract

In this paper we face a very fundamental problem in Operations Research: to find good dual bounds to generic mixed integer mathematical programs (MIPs) as quickly as possible. In particular, we focus on the scenario where large scale data needs to be considered, multi-core CPU architectures are available, and massive parallelism can be exploited by means of decomposition methods. We consider column generation techniques to solve extended formulations obtained by means of Dantzig-Wolfe decomposition for MIPs. We propose a concurrent algorithm, that relaxes the synchronized behavior of classical column generation. Our approach relies on simple data structures and efficient synchronization, still providing the same global convergence properties of classical sequential column generation methods. We present and discuss the results of an extensive experimental campaign, comparing our concurrent algorithm to both a naive parallelization of column generation and the cutting planes algorithm implemented in state-of-the-art commercial optimization packages, considering large scale datasets of a hard packing problem from the literature as representative benchmark. Our approach turns out to be on average one order of magnitude faster than competitors, attaining almost linear speedups as the number of available CPU cores increases.