An asynchronous parallel bundle method with an application in train timetabling

Frank Fischer, Christoph Helmberg

October 24, 2016

Abstract

The train timetabling problem asks for conflict free schedules for a set of trains in an railway network. One often used model for aperiodic instances is based on time expanded networks. The schedule of each single train corresponds to a path in its network and certain restrictions like station capacities or headway times are modelled by coupling constraints. The standard approach for solving large scale instances is Lagrangian relaxation of the coupling constraints. The resulting dual problem is then solved using subgradient or bundle methods to determine optimal Lagrangian multipliers.

The classical bundle method will solve all subproblems in each iteration, possibly in parallel, to compute the next iterate. If the number of subproblems is large, many of these subproblems might not provide useful information for the next candidate. In this talk we propose a different approach that optimizes, asynchronously in parallel, subspaces of multipliers that are selected dynamically. The algorithm starts several parallel processes and in a kind of parallel coordinated descent each process selects a subspace with large predicted decrease. Then each process optimizes the associated multipliers independently until a certain improvement level has been achieved and writes its solution back to the global data. Because this improvement may lead to increased violation of other constraints, the algorithm automatically detects and tracks these dependencies and respects them in future subspace selections ensuring global convergence.

The algorithm is intended for problems in which the subproblems have a certain loosely coupled structure in the sense that each coupling constraint connects only few subproblems. Then one can hope that the number of subproblems in each subspace problem is only a small fraction of all subproblems. The train timetabling problem has this structure and we present some computational experiments comparing the new algorithm with the classical approach on some large scale instances.