Abstract

Linear programming (LP) is nowadays probably the most frequently used optimization technique in science and industry. During the last fifteen years the dual simplex method has become a strong contender in solving large scale LP problems. Furthermore, it constitutes an indispensable subroutine within branch-and-cut methods deployed to solve mixed-integer linear programming problems.

Despite of its success and relevance for future research only very publications in research literature discuss mathematical or computational techniques proposed for the dual simplex algorithm from the perspective of implementation. The lack of descriptions of important implementation details has led to a great performance gap between open-source research codes and commercial LP-systems.

In this thesis we present the mathematical algorithms, computational techniques and implementation details, which are the key factors for our dual simplex code to close this gap. Special attention is given to three main issues: 1. Dual phase I methods. 2. Exploitation of hypersparsity. 3. Implementation of dual pricing and ratio test.

In our study of dual phase I methods we show that the task of minimizing the sum of dual infeasibilities can be explicitly modeled as a subproblem and directly be solved by the dual phase II. This subproblem approach is much easier to implement and as powerful as previously proposed algorithmic methods. Furthermore, we propose a new method, which combines Pan's dual phase I algorithm and the subproblem approach and turns out to be superior to other methods in our computational tests. We also discuss the impact of LP preprocessing on dual feasibility.

Our subroutines to solve the required systems of linear equations are based on an LU-factorization of the basis. We give a mathematical description of this technique, which allows to use only one instead of two permutation matrices in the FTran and BTran operations. Based on this framework, we present the first detailed description of how to exploit hypersparsity in the dual simplex algorithm.

We describe several techniques to solve numerically difficult LP problems and reduce the number of degenerate iterations. Here, our main contribution is the conceptual integration of Harris' ratio test, bound flipping and cost shifting techniques and the description of a sophisticated and efficient implementation. Furthermore, we address important issues of the implementation of dual steepest edge pricing and show how to maintain a vector of primal infeasibilities.

Finally we show, that our dual simplex code outperforms the best existing open-source and research codes and is competitive to the leading commercial LP-systems on our most difficult test problems.