Chapter 8

THE MOSEK INTERIOR POINT OPTIMIZER FOR LINEAR PROGRAMMING: AN IMPLEMENTATION OF THE HOMOGENEOUS ALGORITHM

Erling D. Andersen and Knud D. Andersen

Abstract The purpose of this work is to present the MOSEK optimizer intended for solution of large-scale sparse linear programs. The optimizer is based on the homogeneous interior-point algorithm which in contrast to the primal-dual algorithm detects a possible primal or dual infeasibility reliably. It employs advanced (parallelized) linear algebra, it handles dense columns in the constraint matrix efficiently, and it has a basis identification procedure.

This paper discusses in details the algorithm and linear algebra employed by the MOSEK interior point optimizer. In particular the homogeneous algorithm is emphasized. Furthermore, extensive computational results are reported. These results include comparative results for the XPRESS simplex and the MOSEK interior point optimizer. Finally, computational results are presented to demonstrate the possible speed-up, when using a parallelized version of the MOSEK interior point optimizer on a multiprocessor Silicon Graphics computer.

1. INTRODUCTION

During the last decade interior-point methods have gained acceptance within the optimization community. Indeed the methods are now considered an efficient alternative to the simplex method for solution of large-scale LP problems [13].

The renewed interest in interior-point methods started with Karmarkar's 1984 paper [27] and since then there has been a rapid development in both the theoretical foundation of the methods and in their practical implementation. We will not survey this development in detail, but refer the reader to the references [21, 7, 40, 42].

H Frenk et al. (eds.), High Performance Optimization, 197–232.

One of the main results of this development is the primal-dual infeasible-interior-point method, for brevity, called the primal-dual method. In fact this method forms the basis for the majority of the available interior-point based software. The feasible primal-dual method was first proposed by Kojima, Mizuno, and Yoshise [30] and later several researchers made important contributions to the development of the infeasible variant notably [32, 36, 29]. Also a series of papers published by Lustig, Marsten, and Shanno have strongly influenced the implementation of the primal-dual method, see [34] and the references therein. However, the primal-dual method as presented in [34] cannot detect a possible infeasible or unbounded status of the LP problem, which of course is a serious drawback in a general purpose optimizer. Moreover, if the LP solver is used to solve the LP relaxation within a branch and bound optimizer for integer programming, then reliable detection of infeasibility is important.

A possible remedy for this problem is suggested by Ye, Todd, and Mizuno [43] who present a homogeneous LP model. This model embeds the original LP problem in a slightly larger LP problem that always has a solution. The solution to the homogeneous model either proves that the original problem does not have an optimal solution or can easily be converted to an optimal solution to the original problem. The Ye, Todd, and Mizuno homogeneous model was later simplified by Xu, Hung, and Ye [41]. Furthermore, they present an implementation of the homogeneous algorithm and encouraging computational results.

In this paper we present our implementation of the homogeneous algorithm for LP. It is fairly similar to the one presented in [41], but there are differences in the choice of algorithmic parameters, the stopping criteria and notably in the starting point.

The outline of the paper is as follows. In Section 2 we present our notation. In Section 3 we introduce the homogeneous model. Our emphasis is on how the model detects a possible infeasible problem status. In the case the LP problem is infeasible we discuss how the solution to the homogeneous model can be useful in diagnosing the cause of the infeasibility. In Section 4 we present an implementation of the homogeneous algorithm. This is followed by a discussion of the linear algebra implemented in the MOSEK interior point optimizer in Section 5. In Section 7 we briefly discuss the identification of an optimal basis starting from the optimal interior-point solution. In Section 8, preprocessing of the LP is discussed. Finally, in Section 9, we report computational results for several large-scale test problems, and compare the computational efficiency of the XPRESS simplex and the MOSEK interior point optimizers.