Parallel distributed-memory simplex for large-scale stochastic LP problems

J. A. J. Hall M. Lubin

19 June 2014

Abstract

Although parallel efficiency using the revised simplex method has not been achieved for general large sparse LP problems, this talk will show how the particular structure of stochastic LP problems gives scope for efficient data parallelism. Issues relating to algorithmic design and data distribution will be discussed. Results obtained on a large cluster and supercomputer using a distributed-memory implementation are presented for stochastic LP problems with up to 10^8 variables and constraints.

Keywords: Simplex method; Parallel computing; Stochastic optimization; Block-angular

ACM Classification: 90C05; 90C15; 68W10

1 Introduction

In this talk, we present a parallel solution procedure based on the revised simplex method for linear programming (LP) problems with a special structure of the form

$$\begin{array}{lccccccccccc}
\text{minimize} & \boldsymbol{c}_0^T\boldsymbol{x}_0 & + & \boldsymbol{c}_1^T\boldsymbol{x}_1 & + & \boldsymbol{c}_2^T\boldsymbol{x}_2 & + & \dots & + & \boldsymbol{c}_N^T\boldsymbol{x}_N & & \\
\text{subject to} & A\boldsymbol{x}_0 & & & & & & & & & = & \boldsymbol{b}_0, \\
 & T_1\boldsymbol{x}_0 & + & W_1\boldsymbol{x}_1 & & & & & & & = & \boldsymbol{b}_1, \\
 & T_2\boldsymbol{x}_0 & & & + & W_2\boldsymbol{x}_2 & & & & & = & \boldsymbol{b}_2, \\
 & \vdots & & & & & & \ddots & & & & \vdots \\
 & T_N\boldsymbol{x}_0 & & & & & & & + & W_N\boldsymbol{x}_N & = & \boldsymbol{b}_N, \\
 & \multicolumn{11}{l}{\boldsymbol{x}_0 \ge \boldsymbol{0}, \quad \boldsymbol{x}_1 \ge \boldsymbol{0}, \quad \boldsymbol{x}_2 \ge \boldsymbol{0}, \quad \dots, \quad \boldsymbol{x}_N \ge \boldsymbol{0}.}
\end{array}$$

Very large instances of such problems have been considered to be too big to solve with the simplex method; instead, decomposition approaches based on Benders decomposition or, more recently, interior-point methods are generally used. However, these approaches do not provide optimal basic solutions which allow for the efficient hot-starts required, for example, in a branch-and-bound context, and can provide important sensitivity information. Our approach exploits the dual block-angular structure of these problems inside the linear algebra of the revised simplex method in a manner suitable for high-performance distributed-memory clusters or supercomputers. While the focus is on stochastic LPs, the work is applicable to all problems with a dual block-angular structure. Our implementation is competitive in serial with highly efficient sparsity-exploiting simplex codes and achieves parallel efficiency when using up to 128 cores and runs up to 100 times faster than the leading open-source serial solver. Additionally, very large problems with hundreds of millions of variables have been successfully solved to optimality.

2 Data parallel linear algebra

The structure of the basis matrix B and matrix N corresponding to the nonbasic variables permits distribution of the data and computation relating to solution of systems of equations involving B and products involving N. Minimal duplicated computation leads to relatively little data transfer being required. The numerical linear algebra is was developed from the COIN-OR utilities and is efficient with respect to the hyper-sparsity present in the problems. Thus the implementation is comparable with world-class open source revised simplex solvers.

3 Test problems, results and conclusions

The principal source of test problems are deterministic LP problems of scenarios resulting from sampling a minimum expected cost stochastic model of wind power generation in the state of Illinois. Increasing the number of scenarios yields ever larger deterministic LP problems, allowing a range of experiments to be performed on two distributed-memory machines: a 320-node cluster of dual quad-core Xeon processors with an InfiniBand QDR interconnect and a Blue Gene/P (BG/P) supercomputer with 40,960 nodes of quad-core 850 MHz PowerPC processors. The nature of the scenario sampling allows a bootstrapping approach to be used to deduce an advanced initial basis, considerably reducing the solution time which would be required otherwise. The largest instance had 463,113,276 variables and 486,899,712 constraints and was solved to optimality: possibly the largest LP ever solved using the simplex method.

This is the largest-scale parallel sparsity-exploiting revised simplex implementation that has been developed to date and the first truly distributed solver. It is built on novel analysis of the linear algebra for dual block-angular LP problems when solved by using the revised simplex method and a novel parallel scheme for applying product-form updates.

This work was awarded the COIN-OR INFORMS 2013 Cup.