

Parallel basis matrix triangularisation for hyper-sparse LP problems

Julian Hall

School of Mathematics

University of Edinburgh

July 4th 2007



Overview

- Nature of the challenge of matrix inversion for the revised simplex method
- Dominant costs of serial matrix inversion
- Parallel matrix inversion strategies
- Parallel triangularisation scheme
- Results
 - serial simulation
 - parallel implementation


Solving LP problems

$$\begin{array}{ll} \text{minimize} & f = \mathbf{c}^T \mathbf{x} \\ \text{subject to} & A\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0} \\ \text{where} & \mathbf{x} \in \mathbb{R}^n \quad \text{and} \quad \mathbf{b} \in \mathbb{R}^m \end{array}$$

- At any vertex the variables may be partitioned into index sets
 - $\mathcal{B}$ of m basic variables $\mathbf{x}_B \geq \mathbf{0}$
 - $\mathcal{N}$ of $n - m$ nonbasic variables $\mathbf{x}_N = \mathbf{0}$
- Columns of A are partitioned correspondingly into
 - the **basis matrix** B
 - the matrix N

Simplex method and hyper-sparsity


- Major computational cost each iteration is forming $B^{-1}\mathbf{r}_F$, $\mathbf{r}_B^T B^{-1}$ and $\mathbf{r}_\pi^T N$
- Vectors $\mathbf{r}_F$, $\mathbf{r}_B$ are always sparse
- Vector $\mathbf{r}_\pi$ may be sparse
- If the results of these operations are (usually) sparse then LP is said to be **hyper-sparse**
- When exploiting hyper-sparsity
 - For hyper-sparse LP problems (○)
Simplex method typically better than interior point methods
 - For other LPs (●)
Interior point methods frequently better than simplex method



Hall and McKinnon: COAP **32** (2005) 259–283



Basis matrix inversion: The nature of the challenge



- Consider optimal block triangular form (Tarjan) of B
 - All diagonal blocks are 1×1 (trivial) or irreducible
 - For LU factors of B , elimination is restricted to diagonal blocks
- Total cost of inverting B depends on
 - Cost of finding irreducible blocks
 - Cost of factorising irreducible blocks
- Relative number/size of 1×1 and irreducible Tarjan blocks is closely related to hyper-sparsity of the LP

Tomlin INVERT (1972)

- Matrix inversion procedure designed for the revised simplex method
- Operates in two phases

Triangularisation:


- Active row/column **singleton** entries are identified as pivots until all active rows/columns have at least *two* active nonzeros
- Corresponds to Markowitz pivot selection so long as count is zero
- Yields approximate Tarjan form at *much* lower cost

Factorisation:

- Gaussian elimination applied the residual **bump**



Relative time for triangularisation and factorization



- Triangularisation can be hugely dominant: particularly for large hyper-sparse LPs
- Bump factorization can be hugely dominant: particularly for smaller non hyper-sparse LPs

Parallel basis matrix inversion: Overview

- Both phases of Tomlin matrix inversion can dominate the computational cost
- Depends on the nature of the LP
- Parallel bump factorisation (important for smaller, non hyper-sparse LPs)
 - Of lesser interest
 - Consider using existing codes such as MUMPS or SuperLU
- Parallel triangularisation (important for larger, hyper-sparse LPs)
 - Of greater interest
 - The subject of the rest of this talk