High performance numerical linear algebra for the revised simplex method

Julian Hall

School of Mathematics University of Edinburgh

jajhall@ed.ac.uk

Workshop on Linear Algebra for PDEs and Optimization, Edinburgh

4 September 2017



- Primal revised simplex method
 - Solving linear systems (FTRAN and BTRAN) with focus on exploiting hyper-sparsity
 - Exploiting $\bar{B} = B + (\boldsymbol{a}_q B\boldsymbol{e}_p)\boldsymbol{e}_p^T$ (UPDATE-BASIS)
 - Traditional techniques
 - Novel techniques
- Exploiting parallelism
 - Background
 - Data parallelism by exploiting LP structure
 - Task parallelism for general LP problems
- Conclusions

The simplex algorithm: Choosing a column



Primal algorithm: Assume $\widehat{\boldsymbol{b}} \geq \boldsymbol{0}$ Seek $\widehat{\boldsymbol{c}}_{\scriptscriptstyle N} \geq \boldsymbol{0}$

 $\begin{array}{ll} \text{Scan } \widehat{c_j} < 0, \ i \in \mathcal{N}, \ \text{for a good candidate } q \ \text{to leave } \mathcal{N} & \quad \text{CHUZC} \\ \text{Scan } \widehat{b_i} / \widehat{a_{iq}} > 0, \ i \in \mathcal{B}, \ \text{for a good candidate } p \ \text{to leave } \mathcal{B} & \quad \text{CHUZR} \end{array}$

The simplex algorithm: Choosing a row



Primal algorithm: Assume $\widehat{\boldsymbol{b}} \geq \boldsymbol{0}$ Seek $\widehat{\boldsymbol{c}}_{\scriptscriptstyle N} \geq \boldsymbol{0}$

 $\begin{array}{ll} \text{Scan } \widehat{c_j} < 0, \ i \in \mathcal{N}, \ \text{for a good candidate } q \ \text{to leave } \mathcal{N} & \quad \text{CHUZC} \\ \text{Scan } \widehat{b_i} / \widehat{a_{iq}} > 0, \ i \in \mathcal{B}, \ \text{for a good candidate } p \ \text{to leave } \mathcal{B} & \quad \text{CHUZR} \end{array}$

The simplex algorithm: Update



Update: Exchange p and q between ${\mathcal B}$ and ${\mathcal N}$

Update $\widehat{m{b}}:=\widehat{m{b}}-lpha_{m{P}}\widehat{m{a}}_{m{q}}$	$lpha_{m{P}}=\widehat{m{b}}_{m{p}}/\widehat{m{a}}_{m{pq}}$	UPDATE-PRIMAL
Update $\widehat{\boldsymbol{c}}_{\scriptscriptstyle N}^{ \mathcal{T}} := \widehat{\boldsymbol{c}}_{\scriptscriptstyle N}^{ \mathcal{T}} + \alpha_D \widehat{\boldsymbol{a}}_{\scriptscriptstyle P}^{ \mathcal{T}}$	$lpha_{D} = -\widehat{c}_{q}/\widehat{a}_{pq}$	UPDATE-DUAL

Data required

• Pivotal row
$$\widehat{\boldsymbol{a}}_p^T = \boldsymbol{e}_p^T B^{-1} N$$

• Pivotal column
$$\widehat{\boldsymbol{a}}_q = B^{-1} \boldsymbol{a}_q$$

Why does it work?

Objective improves by
$$\left| \frac{\widehat{b}_p imes \widehat{c}_q}{\widehat{a}_{pq}} \right|$$
 each iteration

Standard simplex method (SSM): Major computational component



Update of tableau:
$$\widehat{N} := \widehat{N} - \frac{1}{\widehat{a}_{pq}} \widehat{a}_q \widehat{a}_p^T$$

where $\widehat{N} = B^{-1} N$

- Hopelessly inefficient for sparse LP problems
- Prohibitively expensive for large LP problems

Revised simplex method (RSM): Major computational components

Pivotal row via $B^T \pi_p = \boldsymbol{e}_p$ BTRANand $\widehat{\boldsymbol{a}}_p^T = \pi_p^T N$ PRICEPivotal column via $B \, \widehat{\boldsymbol{a}}_q = \boldsymbol{a}_q$ FTRANRepresent B^{-1} INVERT

Recall: major computational components

- BTRAN: Form $\pi_p = B^{-T} \boldsymbol{e}_p$
- **PRICE**: Form $\widehat{a}_p^T = \pi_p^T N$
- FTRAN: Form $\widehat{a}_q = B^{-1} a_q$

Phenomenon of hyper-sparsity

- Vectors \boldsymbol{e}_p and \boldsymbol{a}_q are sparse
- Results π_p , \hat{a}_p^T and \hat{a}_q may be sparse—because B^{-1} is sparse
 - In BTRAN, π_p is a row of B^{-1}
 - In PRICE, \widehat{a}_{p}^{T} is a linear combination of a few rows of N
 - In FTRAN, \hat{a}_q is a linear combination of a few columns of B^{-1}

Random sparse matrix Bm = 356 and density 2.5%



 B^{-1} has density of 99%



Optimal *B* for LP problem stair m = 356 and density 2.5%



 B^{-1} has density of 58% $B^{-1}\boldsymbol{b}$ for sparse \boldsymbol{b} is typically dense



Optimal *B* for LP problem pds-02 m = 2953 and density 0.07%



 B^{-1} has density of 0.52% $B^{-1}\boldsymbol{b}$ for sparse \boldsymbol{b} is typically sparse



Random matrix m = 2953 and density 0.07%



 B^{-1} has density of 100%



- Use solution of $L \mathbf{x} = \mathbf{b}$
 - To illustrate the phenomenon of hyper-sparsity
 - To demonstrate how to exploit hyper-sparsity
- Apply principles to other triangular solves in the simplex method

Recall: Solve Lx = b using

function ftranL(L, b, x)

$$r = b$$

for all $j \in \{1, ..., m\}$ do
for all $i : L_{ij} \neq 0$ do
 $r_i = r_i - L_{ij}r_j$
 $x = r$

When **b** is **sparse**

• Inefficient until r fills in

Better: Check r_j for zero

$$\begin{array}{l} \textbf{function ftranL}(L, \ \boldsymbol{b}, \ \boldsymbol{x}) \\ \boldsymbol{r} = \boldsymbol{b} \\ \textbf{for all } j \in \{1, \dots, m\} \ \textbf{do} \\ \textbf{if } r_j \neq 0 \ \textbf{then} \\ \textbf{for all } i : L_{ij} \neq 0 \ \textbf{do} \\ r_i = r_i - L_{ij}r_j \\ \boldsymbol{x} = \boldsymbol{r} \end{array}$$

When x is sparse

- Few values of r_j are nonzero
- Check for zero dominates
- Requires more efficient identification of set X of indices j such that r_j ≠ 0

Gilbert and Peierls (1988) H and McKinnon (1998–2005) COAP best paper prize (2005)

• Each iteration:

- Solve $B^T \pi_p = \boldsymbol{e}_p$
- Solve $B \hat{a}_q = a_q$
- Column p of B replaced by \boldsymbol{a}_q to give $\bar{B} = B + (\boldsymbol{a}_q B\boldsymbol{e}_p)\boldsymbol{e}_p^T$
- Sparsity-exploiting decomposition: PBQ = LU
- Challenge: Solve systems involving \bar{B} at minimal cost

Simplex method: Product form update (PFI) for B

• Given

$$ar{B} = B + (oldsymbol{a}_q - Boldsymbol{e}_p)oldsymbol{e}_p^T$$

• Take *B* out as a factor on the left

$$\bar{B} = B[I + (B^{-1}\boldsymbol{a}_q - \boldsymbol{e}_p)\boldsymbol{e}_p^T] = BE$$

where $E = I + (\widehat{\boldsymbol{a}}_q - \boldsymbol{e}_p)\boldsymbol{e}_p^T = \begin{bmatrix} 1 & \eta_1 & \\ \ddots & \vdots & \\ & \mu & \\ & \vdots & \ddots & \\ & \eta_m & & 1 \end{bmatrix}$

 $\mu=\widehat{a}_{pq}$ is the **pivot**; remaining entries in \widehat{a}_q form the **eta vector** η

• Can solve $\bar{B}\mathbf{x} = \mathbf{r}$ as $B\mathbf{x} = \mathbf{r}$ then $\mathbf{x} := E^{-1}\mathbf{x}$ as

$$x_{p} := x_{p}/\mu$$
 then $\mathbf{x} := \mathbf{x} - x_{p} \boldsymbol{\eta}$

Dantzig and Orchard-Hays (1954)

Simplex method: Forrest-Tomlin update (FT) for B

• Given

$$ar{B} = B + (oldsymbol{a}_q - Boldsymbol{e}_p)oldsymbol{e}_p^T$$
 where (wlog) $B = LU$

• Multiply \bar{B} by L^{-1} to give

$$L^{-1}\bar{B} = U + (L^{-1}\boldsymbol{a}_q - U\boldsymbol{e}_p)\boldsymbol{e}_p^T = U + (\tilde{\boldsymbol{a}}_q - \boldsymbol{u}_p)\boldsymbol{e}_p^T = U' \quad (a)$$

• Eliminate entries in row p to give $R^{-1}U' = \bar{U}$ (b)



- Yields $\bar{B} = LR\bar{U}$
- Compute \widetilde{a}_q when forming \widehat{a}_q
- Represent R like E
- FT more efficient than PFI with respect to sparsity

Forrest and Tomlin (1972)

Simplex method: Multiple updates

- Suppose $B_0 = L_0 U_0$ and k updates are performed to obtain B_k
 - Pivot in rows $\{p_i\}_{i=1}^k$
 - Introduce columns $\{a_{q_i}\}_{i=1}^k$
- PFI generalizes as

$$B_k = B_0 E_1 E_2 \dots E_k$$

• FT generalizes as

$$B_k = L_0 R_1 R_2 \dots R_k U_k$$

Eventually more computationally efficient or numerically prudent to reinvert some B_k

Novel update techniques for the revised simplex method: 1

Alternative product form update (APF)

- **Recall:** Column *p* of *B* is replaced by \boldsymbol{a}_q to give $\bar{B} = B + (\boldsymbol{a}_q B\boldsymbol{e}_p)\boldsymbol{e}_p^T$
 - Traditional PFI takes B out as a factor on the left so $\bar{B} = BE$
- Idea: Why not take it out on the right!

$$ar{B} = [I + (oldsymbol{a}_q - Boldsymbol{e}_p)oldsymbol{e}_p^T B^{-1}]B = TB$$

where $T = I + (oldsymbol{a}_q - oldsymbol{a}_{p'})oldsymbol{\widehat{e}}_p^T$

- T is formed of known data and readily invertible (like E for PFI) Naturally compute \hat{e}_p when solving $B^T \pi_p = e_p$
- But: Is this useful?

Middle product form update (MPF)

- **Recall:** Column *p* of *B* is replaced by \boldsymbol{a}_q to give $\bar{B} = B + (\boldsymbol{a}_q B\boldsymbol{e}_p)\boldsymbol{e}_p^T$
- Idea: Substitute B = LU and take factors L on the left and U on the right!

$$\begin{split} \bar{B} &= LU + (\boldsymbol{a}_q - B\boldsymbol{e}_p)\boldsymbol{e}_p^T \\ &= LU + LL^{-1}(\boldsymbol{a}_q - B\boldsymbol{e}_p)\boldsymbol{e}_p^T U^{-1}U \\ &= L[I + (\widetilde{\boldsymbol{a}}_q - U\boldsymbol{e}_p)\widetilde{\boldsymbol{e}}_p^T]U \\ &= LMU \quad \text{where} \quad M = I + (\widetilde{\boldsymbol{a}}_q - \boldsymbol{u}_p)\widetilde{\boldsymbol{e}}_p^T \end{split}$$

- *M* is formed of known data and readily invertible (like *E* for PFI) Naturally compute \tilde{a}_q when solving $B \hat{a}_q = a_q$ and \tilde{e}_p when solving $B^T \pi_p = e_p$
- But: Is this useful?

Novel update techniques for the revised simplex method: 3

Forrest-Tomlin update

- **Recall:** Column p of B is replaced by \boldsymbol{a}_q to give $\bar{B} = B + (\boldsymbol{a}_q B\boldsymbol{e}_p)\boldsymbol{e}_p^T$
- Idea: Substitute B = LU and take factor L on the left

$$ar{B} = L[U + (L^{-1}oldsymbol{a}_q - Uoldsymbol{e}_p)oldsymbol{e}_p^T] = LU'$$

where $U' = U + (\tilde{\boldsymbol{a}}_q - \boldsymbol{u}_p) \boldsymbol{e}_p^T$ is "spiked" upper triangular and then eliminate

Collective Forrest-Tomlin (CFT) update

- Update Forrest-Tomlin representation of B after multiple basis changes
- Don't have data to perform a sequence of standard FT updates
- Have to perform elimination corresponding to multiple spikes
- But: Is this useful?

Huangfu and H (2013)

Novel update techniques for the revised simplex method: Multiple updates

- Suppose $B_0 = L_0 U_0$ and k updates are performed to obtain B_k
- **Recall:** PFI generalizes as $B_k = B_0 E_1 E_2 \dots E_k$
- **Recall:** FT generalizes as $B_k = L_0 R_1 R_2 \dots R_k U_k$
- APF generalizes as

$$B_k = T_k \dots T_2 T_1 B$$

• MPF generalizes as

$$B_k = L_0 M_1 M_2 \dots M_k U_0$$

Novel update techniques for the revised simplex method: Results

- Test environment
 - 30 representative LP problems from standard test sets
 - Same sequence of basis changes for all update techniques
- Geometric mean time for operations to form and use update data

Update	PFI	FT	APF	MPF	CFT
Mean	1.00	0.30	0.95	0.45	0.29

Conclusions:

- FT much better than PFI
- APF little better than PFI
- MPF closer to FT than PFI
- CFT as good as FT

Parallelising the revised simplex method

Parallelising the simplex method

History

- SSM: Good parallel efficiency of $\widehat{N} := \widehat{N} \frac{1}{\widehat{a}_{pq}} \widehat{a}_{p}^{T}$ was achieved Many! (1988–date)
- Only parallel revised simplex is worthwhile: goal of H and McKinnon in 1992!

Parallel primal revised simplex method

• Overlap computational components for different iterations

Wunderling (1996), H and McKinnon (1995-2005)

• Modest speed-up was achieved on general sparse LP problems

Parallel dual revised simplex method

- Only immediate parallelism is in forming $\pi_p^T N$
- When $n \gg m$ significant speed-up was achieved

Parallelising the simplex method: Matrix-vector product $\pi_{\rho}^{T}N$ (PRICE)

• In theory:

- Partition $N = \begin{bmatrix} N_1 & N_2 & \dots & N_P \end{bmatrix}$, where P is the number of processes
- Compute $\pi_p^T N_j$ on process j for j = 1, 2, ... P
- Execution time is reduced by a factor P: linear speed-up

In practice

- On a distributed memory machine with N_j on processor j Linear speed-up achieved
- On a shared memory machine
 - If *N_j* fits into cache for process *j* Linear speed-up achieved
 - Otherwise, computation is limited by speed of reading data from memory Speed-up limited to number of memory channels

Data parallelism by exploiting LP structure

PIPS-S (2011-date)

Overview

- Written in C++ to solve stochastic MIP relaxations in parallel
- Dual simplex
- Based on NLA routines in clp
- Product form update

Concept

- Exploit data parallelism due to block structure of LPs
- Distribute problem over processes

Paper: Lubin, H, Petra and Anitescu (2013)

- COIN-OR INFORMS 2013 Cup
- COAP best paper prize (2013)

PIPS-S: Stochastic MIP problems

Two-stage stochastic LPs have column-linked block angular (BALP) structure

- Variables $x_0 \in \mathbb{R}^{n_0}$ are first stage decisions
- Variables $\mathbf{x}_i \in \mathbb{R}^{n_i}$ for i = 1, ..., N are second stage decisions Each corresponds to a scenario which occurs with modelled probability
- The objective is the expected cost of the decisions
- In stochastic MIP problems, some/all decisions are discrete

- Power systems optimization project at Argonne
- Integer second-stage decisions
- Stochasticity from wind generation
- Solution via branch-and-bound
 - Solve root using parallel IPM solver PIPS

Lubin, Petra et al. (2011)

• Solve nodes using parallel dual simplex solver PIPS-S





Convenient to permute the LP thus:

PIPS-S: Exploiting problem structure

- Inversion of the basis matrix B is key to revised simplex efficiency
- For column-linked BALP problems

$$B = \begin{bmatrix} W_1^B & & T_1^B \\ & \ddots & & \vdots \\ & & W_N^B & T_N^B \\ & & & & A^B \end{bmatrix}$$

• W_i^B are columns corresponding to n_i^B basic variables in scenario i

•
$$\begin{bmatrix} T_1^B \\ \vdots \\ T_N^B \\ A^B \end{bmatrix}$$
 are columns corresponding to n_0^B basic first stage decisions

PIPS-S: Exploiting problem structure

- Inversion of the basis matrix B is key to revised simplex efficiency
- For column-linked BALP problems

$$B = \begin{bmatrix} W_{1}^{B} & & T_{1}^{B} \\ & \ddots & & \vdots \\ & & W_{N}^{B} & T_{N}^{B} \\ & & & A^{B} \end{bmatrix}$$



- B is nonsingular so
 - $W_i^{\scriptscriptstyle B}$ are "tall": full column rank
 - $\begin{bmatrix} W_i^B & T_i^B \end{bmatrix}$ are "wide": full row rank
 - $\tilde{A}^{\scriptscriptstyle B}$ is "wide": full row rank
- Scope for parallel inversion is immediate and well known

• Eliminate sub-diagonal entries in each W_i^B (independently)



• Apply elimination operations to each T_i^B (independently)

• Accumulate non-pivoted rows from the $W_i^{\scriptscriptstyle B}$ with $A^{\scriptscriptstyle B}$ and complete elimination



Scope for parallelism

- Parallel Gaussian elimination yields block LU decomposition of B
- Scope for parallelism in block forward and block backward substitution
- Scope for parallelism in PRICE

Implementation

- Distribute problem data over processes
- Perform data-parallel BTRAN, FTRAN and PRICE over processes
- Used MPI

PIPS-S: Results

n Fusion cluster: Performance relative to c1p							
	Dimension	Cores	Storm	SSN	UC12	UC24	•
	$m+n=O(10^6)$	1 32	0.34 8.5	0.22 6.5	0.17 2.4	0.08 0.7	-
	$\overline{m+n=O(10^7)}$	256	299	45	67	68	

On Blue Gene

- Instance of UC12
- $m + n = O(10^8)$
- Requires 1 TB of RAM
- Runs from an advanced basis

Cores	Iterations	Time (h)	lter/sec
1024	Exceeded	execution	time limit
2048	82,638	6.14	3.74
4096	75,732	5.03	4.18
8192	86,439	4.67	5.14

Task parallelism for general LP problems

h₂gmp (2011-date)

Overview

- Written in C++ to study parallel simplex
- Dual simplex with steepest edge and BFRT
- Forrest-Tomlin update
 - complex and inherently serial
 - efficient and numerically stable

Concept

- Exploit limited task and data parallelism in standard dual RSM iterations (sip)
- Exploit greater task and data parallelism via minor iterations of dual SSM (pami)
- Test-bed for research
- Work-horse for consultancy

Huangfu, H and Galabova (2011-date)

h_2gmp : Multiple iteration parallelism with pami option

- Perform standard dual simplex minor iterations for rows in set $\mathcal{P}~(|\mathcal{P}|\ll m)$
- Suggested by Rosander (1975) but never implemented efficiently in serial



- Task-parallel multiple BTRAN to form $m{\pi}_{\mathcal{P}}=B^{-1}m{e}_{\mathcal{P}}$
- Data-parallel PRICE to form \widehat{a}_{p}^{T} (as required)
- Task-parallel multiple FTRAN for primal, dual and weight updates

Huangfu and H (2011–2014) COAP best paper prize (2015)

Novel update procedures: are they useful?

Update 1: Alternative product form update (APF)

- **Recall:** $\bar{B} = TB$ with T easily invertible
- Used to get $\pi_{\mathcal{P}} = \bar{B}^{-T} \boldsymbol{e}_{\mathcal{P}}$ from $B^{-T} \boldsymbol{e}_{\mathcal{P}}$ in pami
- Used to compute multiple $B^{-1}\boldsymbol{a}_F$ efficiently after multiple BFRT in pami

Update 2: Middle product form update (MPF)

- **Recall:** $\bar{B} = LMU$ with *M* easily invertible
- Not used by pami!
- Used by **Google** in glop

Update 3: Multiple Forrest-Tomlin update

Used to perform multiple Forrest-Tomlin updates after minor iterations in pami

h_2gmp : Performance and reliability

Extended testing using 159 test problems

- 98 Netlib
- 16 Kennington
- 4 Industrial
- 41 Mittelmann

Exclude 7 which are "hard"

Performance

Benchmark against clp (v1.16) and cplex (v12.5)

- Dual simplex
- No presolve
- No crash

Ignore results for 82 LPs with minimum solution time below 0.1s

h₂gmp: Performance



Conclusions

High performance NLA is essential for the revised simplex method

- Hyper-sparsity: advance in serial simplex
- Stochastic LP: large scale data-parallelism for special problems
- Novel updates: underpin open-source parallel simplex solver

J. A. J. Hall and K. I. M. McKinnon.

Hyper-sparsity in the revised simplex method and how to exploit it. Computational Optimization and Applications, 32(3):259–283, December 2005.

Q. Huangfu and J. A. J. Hall.

Parallelizing the dual revised simplex method.

Technical Report ERGO-14-011, School of Mathematics, University of Edinburgh, 2014. Accepted for publication in Mathematical Programming Computation.



Q. Huangfu and J. A. J. Hall.

Novel update techniques for the revised simplex method. Computational Optimization and Applications, 60(4):587–608, 2015.



M. Lubin, J. A. J. Hall, C. G. Petra, and M. Anitescu.

Parallel distributed-memory simplex for large-scale stochastic LP problems. Computational Optimization and Applications, 55(3):571–596, 2013.

Slides: http://www.maths.ed.ac.uk/hall/NLAPDEO17/