The practical revised simplex method (Part 2)

Julian Hall

School of Mathematics

University of Edinburgh

January 25th 2007

The practical revised simplex method

- Practical implementation of the revised simplex method
 - Representation of B^{-1}

- Practical implementation of the revised simplex method
 - Representation of B^{-1}
 - Strategies for CHUZC

- Practical implementation of the revised simplex method
 - Representation of B^{-1}
 - Strategies for CHUZC
 - Partial pricing

- Practical implementation of the revised simplex method
 - Representation of B^{-1}
 - Strategies for CHUZC
 - Partial pricing
 - Hyper-sparsity

- Practical implementation of the revised simplex method
 - Representation of B^{-1}
 - Strategies for CHUZC
 - Partial pricing
 - Hyper-sparsity
- Parallel simplex

- Practical implementation of the revised simplex method
 - Representation of B^{-1}
 - Strategies for CHUZC
 - Partial pricing
 - Hyper-sparsity
- Parallel simplex
- Research frontiers

Representing B^{-1}

• The key to the efficiency of the revised simplex method is the representation of B^{-1} .

Representing B^{-1}

- The key to the efficiency of the revised simplex method is the representation of B^{-1} .
- Periodically: **INVERT** operation determines representation of B^{-1} using Gaussian elimination

Representing B^{-1}

- The key to the efficiency of the revised simplex method is the representation of B^{-1} .
- Periodically: **INVERT** operation determines representation of B^{-1} using Gaussian elimination
- Each iteration: **UPDATE** operation updates representation of B^{-1} according to

$$B^{-1} := \left(I - \frac{(\hat{\boldsymbol{a}}_q - \boldsymbol{e}_p)\boldsymbol{e}_p^T}{\hat{a}_{pq}}\right) B^{-1}$$

- The representation of B^{-1} is based on LU decomposition from Gaussian elimination.
- Classical factors B = LU fill in.
- Perform row and column interchanges to maintain sparsity: yields PBQ = LU.

- The representation of B^{-1} is based on LU decomposition from Gaussian elimination.
- Classical factors B = LU fill in.
- Perform row and column interchanges to maintain sparsity: yields PBQ = LU.
 - Markowitz pivot selection criterion (1957) preserves sparsity well.

- The representation of B^{-1} is based on LU decomposition from Gaussian elimination.
- Classical factors B = LU fill in.
- Perform row and column interchanges to maintain sparsity: yields PBQ = LU.
 - Markowitz pivot selection criterion (1957) preserves sparsity well.
 - Tomlin pivot selection criterion (1972) more efficient for most LP basis matrices.

- The representation of B^{-1} is based on LU decomposition from Gaussian elimination.
- Classical factors B = LU fill in.
- Perform row and column interchanges to maintain sparsity: yields PBQ = LU.
 - Markowitz pivot selection criterion (1957) preserves sparsity well.
 - Tomlin pivot selection criterion (1972) more efficient for most LP basis matrices.
- Permutations P and Q can be absorbed into the ordering of basic variables and constraints.

1 Find identity columns in B



1 Find identity columns in B



2a Find columns with only one nonzero



1 Find identity columns in B



2b Find rows with only one nonzero



2a Find columns with only one nonzero







2a Find columns with only one nonzero



2b Find rows with only one nonzero



3 Factorise the "bump" using GE







2a Find columns with only one nonzero



• Efficiency depends on the bump being small

2b Find rows with only one nonzero



 $3\,$ Factorise the "bump" using GE







2a Find columns with only one nonzero



2b Find rows with only one nonzero



3 Factorise the "bump" using GE



- Efficiency depends on the bump being small
- For many practical LP problems there is little or no bump

Using B = LU from Gaussian elimination

• Gaussian elimination yields

$$B = LU \text{ where } L = \begin{bmatrix} 1 & & & \\ & \ddots & & \\ & & 1 & \\ \boldsymbol{l}_1 & \dots & \boldsymbol{l}_{n-1} & 1 \end{bmatrix} \text{ and } U = \begin{bmatrix} u_{11} & \boldsymbol{u}_2 & \dots & \boldsymbol{u}_n \\ & u_{22} & & & \\ & & \ddots & & \\ & & & u_{nn} \end{bmatrix}$$

Using B = LU from Gaussian elimination

• Gaussian elimination yields

$$B = LU \text{ where } L = \begin{bmatrix} 1 & & & \\ & \ddots & & \\ & & 1 & \\ l_1 & \dots & l_{n-1} & 1 \end{bmatrix} \text{ and } U = \begin{bmatrix} u_{11} & u_2 & \dots & u_n \\ & u_{22} & & \\ & & \ddots & \\ & & & u_{nn} \end{bmatrix}$$

•
$$B \boldsymbol{x} = \boldsymbol{r}$$
 is solved by solving $L \boldsymbol{y} = \boldsymbol{r}$ and $U \boldsymbol{x} = \boldsymbol{y}$.

Using B = LU from Gaussian elimination

• Gaussian elimination yields

$$B = LU \text{ where } L = \begin{bmatrix} 1 & & & \\ & \ddots & & \\ & & 1 & \\ \boldsymbol{l}_1 & \dots & \boldsymbol{l}_{n-1} & 1 \end{bmatrix} \text{ and } U = \begin{bmatrix} u_{11} & \boldsymbol{u}_2 & \dots & \boldsymbol{u}_n \\ & u_{22} & & & \\ & & \ddots & & \\ & & & u_{nn} \end{bmatrix}$$

- $B\mathbf{x} = \mathbf{r}$ is solved by solving $L\mathbf{y} = \mathbf{r}$ and $U\mathbf{x} = \mathbf{y}$.
- In practice, r is transformed into x as follows.

$$egin{array}{lll} m{r} &:= &m{r} - r_km{l}_k & k = 1,\ldots,n-1; \ r_k &:= &m{r}_1 &:= &m{r}_1 \ u_{11} \end{array}, & m{r} &:= &m{r} - r_km{u}_k & k = n,\ldots,2; \end{array}$$

The practical revised simplex method

LU factors as a product of elementary matrices

B = LU may be written as

$$B = (\prod_{k=1}^{n-1} L_k) (\prod_{k=1}^n U_k)$$

where



LU factors as a product of elementary matrices

B = LU may be written as

$$B = (\prod_{k=1}^{n-1} L_k) (\prod_{k=1}^n U_k)$$

where



So

$$B^{-1} = \left(\prod_{k=1}^{n} U_k^{-1}\right) \left(\prod_{k=1}^{n-1} L_k^{-1}\right)$$

Product form and the eta file

• In general,
$$B^{-1} = \prod_{k=K_I}^{1} E_k^{-1}$$
 where

$$E_k^{-1} = \begin{bmatrix} 1 & -\eta_1^k & & \\ & \ddots & \vdots & & \\ & 1 & -\eta_{p_k-1}^k & & \\ & & 1 & & \\ & & -\eta_{p_k+1}^k & 1 & \\ & & \vdots & \ddots & \\ & & -\eta_m^k & & 1 \end{bmatrix} \begin{bmatrix} 1 & & & & \\ & \ddots & & & \\ & & \mu^k & & \\ & & & 1 & \\ & & & & 1 \end{bmatrix}.$$

Product form and the eta file

• In general,
$$B^{-1} = \prod_{k=K_I}^1 E_k^{-1}$$
 where



• The **pivot** is μ_k ; the index of the **pivotal row** is p_k .

Product form and the eta file

• In general,
$$B^{-1} = \prod_{k=K_I}^{1} E_k^{-1}$$
 where



- The **pivot** is μ_k ; the index of the **pivotal row** is p_k .
- The set $\{p_k, \mu_k, \eta_k\}_{k=1}^{K_I}$ is known as the **eta file**.

Eta file from Tomlin INVERT







Basis matrix is of dimension is 821; bump is of dimension 411



Basis matrix is of dimension is 356; bump is of dimension 324

The practical revised simplex method



Basis matrix is of dimension is 246077; bump is of dimension 32

The practical revised simplex method

• The solution of Bx = r is formed by transforming r into x as follows.

$$r_{p_k} := \mu_k r_{p_k}, \quad \boldsymbol{r} := \boldsymbol{r} - r_{p_k} \boldsymbol{\eta}_k \qquad k = 1, \dots, K_I.$$

• The solution of Bx = r is formed by transforming r into x as follows.

 $r_{p_k} := \mu_k r_{p_k}, \quad oldsymbol{r} := oldsymbol{r} - r_{p_k} oldsymbol{\eta}_k \qquad k = 1, \dots, K_I.$

This requires a sparse eta-vector to be added to r (up to) K_I times.

• The solution of Bx = r is formed by transforming r into x as follows.

$$r_{p_k} := \mu_k r_{p_k}, \quad \boldsymbol{r} := \boldsymbol{r} - r_{p_k} \boldsymbol{\eta}_k \qquad k = 1, \dots, K_I.$$

This requires a sparse eta-vector to be added to r (up to) K_I times.

• The solution of $B^T x = r$ is $(r^T B^{-1})^T$ and is formed by transforming r into x as follows.

$$r_{p_k} := \mu_k (r_{p_k} - \boldsymbol{r}^T \boldsymbol{\eta}_k) \qquad k = K_I, \dots, 1$$

• The solution of Bx = r is formed by transforming r into x as follows.

$$r_{p_k} := \mu_k r_{p_k}, \quad \boldsymbol{r} := \boldsymbol{r} - r_{p_k} \boldsymbol{\eta}_k \qquad k = 1, \dots, K_I.$$

This requires a sparse eta-vector to be added to r (up to) K_I times.

• The solution of $B^T x = r$ is $(r^T B^{-1})^T$ and is formed by transforming r into x as follows.

$$r_{p_k} := \mu_k (r_{p_k} - \boldsymbol{r}^T \boldsymbol{\eta}_k) \qquad k = K_I, \dots, 1$$

This requires the inner product between a sparse eta-vector and r (up to) K_I times.
UPDATE

Update of B^{-1} defined by

$$B^{-1}$$
 := $\left(I - rac{(\hat{\boldsymbol{a}}_q - \boldsymbol{e}_p)\boldsymbol{e}_p^T}{\hat{a}_{pq}}\right)B^{-1}$

UPDATE

Update of B^{-1} defined by



UPDATE

Update of B^{-1} defined by

for $p_k = p$, $\mu_k = 1/\hat{a}_{pq}$ and $\boldsymbol{\eta}_k = \hat{\boldsymbol{a}}_q - \hat{a}_{pq} \boldsymbol{e}_p$.

The practical revised simplex method

UPDATE (cont.)

• In general, K_U UPDATEs after INVERT

$$B^{-1} = \prod_{k=K_I+K_U}^{1} E_k^{-1}$$

UPDATE (cont.)

• In general, K_U UPDATEs after INVERT

$$B^{-1} = \prod_{k=K_I+K_U}^{1} E_k^{-1}$$

- Periodically it will more efficient, or necessary for numerical stability, to re-INVERT.
- Note that $K_U \ll K_I = O(m)$

FTRAN and BTRAN

Using the representation

$$B^{-1} = \prod_{k=K_I+K_U}^{1} E_k^{-1}$$

FTRAN and **BTRAN**

Using the representation

$$B^{-1} = \prod_{k=K_I+K_U}^{1} E_k^{-1}$$

• The pivotal column $\hat{a}_q = B^{-1}a_q$ is formed by transforming $r = a_q$ into \hat{a}_q by the **FTRAN** operation

 $r_{p_k}:=\mu_kr_{p_k}, \quad oldsymbol{r}:=oldsymbol{r}-r_{p_k}oldsymbol{\eta}_k \qquad k=1,\ldots,K_I+K_U.$

FTRAN and **BTRAN**

Using the representation

$$B^{-1} = \prod_{k=K_I+K_U}^{1} E_k^{-1}$$

• The pivotal column $\hat{a}_q = B^{-1} a_q$ is formed by transforming $r = a_q$ into \hat{a}_q by the FTRAN operation

 $r_{p_k}:=\mu_kr_{p_k}, \quad oldsymbol{r}:=oldsymbol{r}-r_{p_k}oldsymbol{\eta}_k \qquad k=1,\ldots,K_I+K_U.$

• The vector $\pi^T = e_p^T B^{-1}$ required to find the pivotal row is formed by transforming $r = e_p$ into π by the **BTRAN** operation

$$r_{p_k} := \mu_k (r_{p_k} - \boldsymbol{r}^T \boldsymbol{\eta}_k) \qquad k = K_I + K_U, \dots, 1.$$

The practical revised simplex method

Strategies for CHUZC

'Most negative reduced cost' rule is not the best.

Example:

minimize	f	=	$-2x_1$	_	x_2		
subject to			$100x_1$	+	x_2	\leq	1
			$x_1 \ge 0$	and	$\mid x_2$	> ().

Strategies for CHUZC

'Most negative reduced cost' rule is not the best.

Example:

Slack variable y_1 puts problem in standard form

The practical revised simplex method

minimize
$$f = -2x_1 - x_2$$

subject to $100x_1 + x_2 + y_1 = 1$
 $x_1 \ge 0, x_2 \ge 0$ and $y_1 \ge 0.$

- Starting from logical basis $y_1 = 1$:
- Reduced costs are

$$\begin{array}{ccc} -2 & \text{for} & x_1 \\ -1 & \text{for} & x_2 \end{array}$$

- Starting from logical basis $y_1 = 1$:
- Reduced costs are

-2	for	x_1
-1	for	x_2

- x_1 is best under Dantzig rule but a unit change in x_1 requires a change of -100 in y_1 .
- Problem solved in two iterations.

The practical revised simplex method



- Starting from logical basis $y_1 = 1$:
- Rates of change of objective are

$$-2/\sqrt{1+100^2} \approx -0.02$$
 for x_1
 $-1/\sqrt{1+1} \approx -0.71$ for x_2

- Starting from logical basis $y_1 = 1$:
- Rates of change of objective are

$$-2/\sqrt{1+100^2} \approx -0.02$$
 for x_1
 $-1/\sqrt{1+1} \approx -0.71$ for x_2

- x_2 is best under **steepest edge** rule.
- Problem solved in *one* iteration.

• 'Most negative reduced cost' rule (Dantzig)

- 'Most negative reduced cost' rule (Dantzig)
- 'Greatest reduction' rule (1951)
 - Too expensive in practice

- 'Most negative reduced cost' rule (Dantzig)
- 'Greatest reduction' rule (1951)
 - Too expensive in practice
- 'Steepest Edge' (1963)
 - Updates exact step lengths
 - Requires extra BTRAN and PRICE each iteration
 - Prohibitively expensive start-up cost unless B = I initially

- 'Most negative reduced cost' rule (Dantzig)
- 'Greatest reduction' rule (1951)
 - Too expensive in practice
- 'Steepest Edge' (1963)
 - Updates exact step lengths
 - Requires extra BTRAN and PRICE each iteration
 - Prohibitively expensive start-up cost unless B = I initially
- 'Devex' (1965)
 - Updates approximate step lengths
 - No additional calculation
 - No start-up cost

	Rows	Columns	Nonzeros	Columns per row	
Name	(m)	(n)	(au)	(n/m)	$ au/m^2$
fit2d	25	10500	129018	420	206
nw04	36	87482	724148	2430	559

• Some LP problems naturally have very many more columns than rows

	Rows	Columns	Nonzeros	Columns per row	
Name	(m)	(n)	(au)	(n/m)	$ au/m^2$
fit2d	25	10500	129018	420	206
nw04	36	87482	724148	2430	559

• PRICE completely dominates solution time

	Rows	Columns	Nonzeros	Columns per row	
Name	(m)	(n)	(au)	(n/m)	$ au/m^2$
fit2d	25	10500	129018	420	206
nw04	36	87482	724148	2430	559

- PRICE completely dominates solution time
- Simplex method only needs to choose *an* attractive column, not necessarily the best

	Rows	Columns	Nonzeros	Columns per row	
Name	(m)	(n)	(au)	(n/m)	$ au/m^2$
fit2d	25	10500	129018	420	206
nw04	36	87482	724148	2430	559

- PRICE completely dominates solution time
- Simplex method only needs to choose *an* attractive column, not necessarily the best
- Compute only a small subset of the reduced costs

	Rows	Columns	Nonzeros	Columns per row	
Name	(m)	(n)	(au)	(n/m)	$ au/m^2$
fit2d	25	10500	129018	420	206
nw04	36	87482	724148	2430	559

- PRICE completely dominates solution time
- Simplex method only needs to choose *an* attractive column, not necessarily the best
- Compute only a small subset of the reduced costs
- Number of iterations may increase but each iteration is vastly faster

- Hyper-sparsity exists when
 - the pivotal column $\hat{\boldsymbol{a}}_q = B^{-1} \boldsymbol{a}_q$ is sparse;

 (\mathbf{FTRAN})

- Hyper-sparsity exists when
 - the pivotal column $\hat{a}_q = B^{-1}a_q$ is sparse; the π vector $\pi^T = e_p^T B^{-1}$ is sparse;



- Hyper-sparsity exists when
 - the pivotal column $\hat{a}_q = B^{-1} a_q$ is sparse; the π vector $\pi^T = e_p^T B^{-1}$ is sparse;

 - the pivotal row $\hat{\boldsymbol{a}}_p^T = \boldsymbol{\pi}^T N$ is sparse.

(FTRAN) (BTRAN) (PRICE)

- Hyper-sparsity exists when
 - the pivotal column $\hat{a}_q = B^{-1} a_q$ is sparse;
 - the π vector $\pi^T = e_p^T B^{-1}$ is sparse;
 - the pivotal row $\hat{\boldsymbol{a}}_p^T = \boldsymbol{\pi}^T N$ is sparse.

(FTRAN) (BTRAN) (PRICE)

• Exploit hyper-sparsity both when forming and using these vectors.

See

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, 2005.

• FTRAN forms
$$\hat{\boldsymbol{a}}_q = B^{-1} \boldsymbol{a}_q$$
 using $\boldsymbol{r} = \boldsymbol{a}_q$ in

do 10,
$$k$$
 = 1, $K_I + K_U$
if $(r_{p_k} .eq. 0)$ go to 10
 $r_{p_k} := \mu_k r_{p_k}$
 $r := r - r_{p_k} \eta_k$

10 continue

• FTRAN forms $\hat{\boldsymbol{a}}_q = B^{-1} \boldsymbol{a}_q$ using $\boldsymbol{r} = \boldsymbol{a}_q$ in

do 10,
$$k$$
 = 1, $K_I + K_U$
if $(r_{p_k} \text{ .eq. 0})$ go to 10
 $r_{p_k} := \mu_k r_{p_k}$
 $\boldsymbol{r} := \boldsymbol{r} - r_{p_k} \boldsymbol{\eta}_k$

- 10 continue
- When \hat{a}_q is sparse, most of the work of FTRAN is the test for zero!

• FTRAN forms $\hat{\boldsymbol{a}}_q = B^{-1} \boldsymbol{a}_q$ using $\boldsymbol{r} = \boldsymbol{a}_q$ in

do 10,
$$k$$
 = 1, $K_I + K_U$
if $(r_{p_k} \text{ .eq. 0})$ go to 10
 $r_{p_k} := \mu_k r_{p_k}$
 $\boldsymbol{r} := \boldsymbol{r} - r_{p_k} \boldsymbol{\eta}_k$

- 10 continue
- When \hat{a}_q is sparse, most of the work of FTRAN is the test for zero!

Remedy:

Identify the INVERT etas to be applied by

• knowing etas associated with a nonzero in a particular row Hall and McKinnon (1999)

• FTRAN forms
$$\hat{oldsymbol{a}}_q = B^{-1}oldsymbol{a}_q$$
 using $oldsymbol{r} = oldsymbol{a}_q$ in

do 10,
$$k$$
 = 1, $K_I + K_U$
if $(r_{p_k} \text{ .eq. 0})$ go to 10
 $r_{p_k} := \mu_k r_{p_k}$
 $\boldsymbol{r} := \boldsymbol{r} - r_{p_k} \boldsymbol{\eta}_k$

10 continue

• When \hat{a}_q is sparse, most of the work of FTRAN is the test for zero!

Remedy:

Identify the INVERT etas to be applied by

- knowing etas associated with a nonzero in a particular row
- representing etas as a graph and perform a depth-first search Gilb

Hall and McKinnon (1999) Gilbert and Peierls (1988)

• BTRAN forms
$$oldsymbol{\pi}^T = oldsymbol{e}_p^T B^{-1}$$
 using $oldsymbol{r} = oldsymbol{e}_p^T$ in

do 10,
$$k$$
 = $K_I + K_U$, 1
 $r_{p_k} := \mu_k (r_{p_k} - \boldsymbol{r}^T \boldsymbol{\eta}_k)$

10 continue

• BTRAN forms $\boldsymbol{\pi}^T = \boldsymbol{e}_p^T B^{-1}$ using $\boldsymbol{r} = \boldsymbol{e}_p^T$ in

do 10, $k = K_I + K_U$, 1 $r_{p_k} := \mu_k (r_{p_k} - \boldsymbol{r}^T \boldsymbol{\eta}_k)$

- 10 continue
- When π is sparse, almost all operations in BTRAN are with zeros

• BTRAN forms $\boldsymbol{\pi}^T = \boldsymbol{e}_p^T B^{-1}$ using $\boldsymbol{r} = \boldsymbol{e}_p^T$ in

do 10, $k = K_I + K_U$, 1 $r_{p_k} := \mu_k (r_{p_k} - \boldsymbol{r}^T \boldsymbol{\eta}_k)$

- 10 continue
- When π is sparse, almost all operations in BTRAN are with zeros **Remedy**:
 - Store INVERT etas row-wise
 - Use special techniques developed for FTRAN

Hyper-sparsity in PRICE

- PRICE forms $\hat{\boldsymbol{a}}_p^T = \boldsymbol{\pi}^T N$
- When π is sparse, components of \hat{a}_p are inner products between two sparse vectors
Hyper-sparsity in PRICE

- PRICE forms $\hat{\boldsymbol{a}}_p^T = \boldsymbol{\pi}^T N$
- When π is sparse, components of \hat{a}_p are inner products between two sparse vectors
- Most operations are with zeros

Hyper-sparsity in PRICE

- PRICE forms $\hat{\boldsymbol{a}}_p^T = \boldsymbol{\pi}^T N$
- When π is sparse, components of \hat{a}_p are inner products between two sparse vectors
- Most operations are with zeros

Remedy:

- Store N row-wise
- Form \hat{a}_p^T by combining rows of N corresponding to nonzeros in π .

Parallelising the simplex method

Parallel architectures



Distributed Memory

Parallelising the simplex method

Parallel architectures





Shared Memory

- Standard simplex method parallelises beautifully!
 - Still slower than sequential revised simplex

- Standard simplex method parallelises beautifully!
 - Still slower than sequential revised simplex
- Revised simplex method viewed as sequential
 - Computational components performed in sequence
 - PRICE, CHUZC and CHUZR parallelise

- Standard simplex method parallelises beautifully!
 - Still slower than sequential revised simplex
- Revised simplex method viewed as sequential
 - Computational components performed in sequence
 - PRICE, CHUZC and CHUZR parallelise
 - BTRAN, FTRAN and INVERT are sequential

- Standard simplex method parallelises beautifully!
 - Still slower than sequential revised simplex
- Revised simplex method viewed as sequential
 - Computational components performed in sequence
 - PRICE, CHUZC and CHUZR parallelise
 - BTRAN, FTRAN and INVERT are sequential
- Hall and McKinnon (1995)
 - ParLP: *n*-processor asynchronous (Dantzig)
 - Speed-up factor: 2.5–4.8

- Standard simplex method parallelises beautifully!
 - Still slower than sequential revised simplex
- Revised simplex method viewed as sequential
 - Computational components performed in sequence
 - PRICE, CHUZC and CHUZR parallelise
 - BTRAN, FTRAN and INVERT are sequential
- Hall and McKinnon (1995)
 - ParLP: *n*-processor asynchronous (Dantzig)
 - Speed-up factor: 2.5–4.8
- Wunderling (1996)
 - SMoPlex: 2-processor synchronous (Steepest Edge)
 - Speed-up factor: 0.75–1.4

- Standard simplex method parallelises beautifully!
 - Still slower than sequential revised simplex
- Revised simplex method viewed as sequential
 - Computational components performed in sequence
 - PRICE, CHUZC and CHUZR parallelise
 - BTRAN, FTRAN and INVERT are sequential
- Hall and McKinnon (1995)
 - ParLP: *n*-processor asynchronous (Dantzig)
 - Speed-up factor: 2.5–4.8
- Wunderling (1996)
 - SMoPlex: 2-processor synchronous (Steepest Edge)
 - Speed-up factor: 0.75–1.4
- Hall and McKinnon (1996–97)
 - PARSMI: *n*-processor asynchronous (Devex)
 - Speed-up factor: 1.7–1.9

• Hypersparsity

- Column and row selection strategies to promote hypersparsity in the revised simplex method
 - * Very challenging

• Hypersparsity

- Column and row selection strategies to promote hypersparsity in the revised simplex method
 - * Very challenging
- Specialist variants of Gilbert-Peierls for LP

• Hypersparsity

- Column and row selection strategies to promote hypersparsity in the revised simplex method
 - * Very challenging
- Specialist variants of Gilbert-Peierls for LP

- SYNPLEX: stable synchronous version of PARSMI
 - * Implemented but needs improvements and/or parallel INVERT

• Hypersparsity

- Column and row selection strategies to promote hypersparsity in the revised simplex method
 - * Very challenging
- Specialist variants of Gilbert-Peierls for LP

- SYNPLEX: stable synchronous version of PARSMI
 - * Implemented but needs improvements and/or parallel INVERT
- Parallel Tomlin INVERT
 - * Serial simulation for hyper-sparse LPs gives very encouraging results.
 - * Parallel implementation due Summer 2007

• Hypersparsity

- Column and row selection strategies to promote hypersparsity in the revised simplex method
 - * Very challenging
- Specialist variants of Gilbert-Peierls for LP

- SYNPLEX: stable synchronous version of PARSMI
 - * Implemented but needs improvements and/or parallel INVERT
- Parallel Tomlin INVERT
 - * Serial simulation for hyper-sparse LPs gives very encouraging results.
 - * Parallel implementation due Summer 2007
- Parallel FTRAN and BTRAN

• Hypersparsity

- Column and row selection strategies to promote hypersparsity in the revised simplex method
 - * Very challenging
- Specialist variants of Gilbert-Peierls for LP

- SYNPLEX: stable synchronous version of PARSMI
 - * Implemented but needs improvements and/or parallel INVERT
- Parallel Tomlin INVERT
 - * Serial simulation for hyper-sparse LPs gives very encouraging results.
 - * Parallel implementation due Summer 2007
- Parallel FTRAN and BTRAN
 - * Some scope for standard LPs
 - * Very challenging for hyper-sparse LPs!

The simplex method has been around for 60 years

The simplex method has been around for 60 years

• Value of method means that the computational state-of-the-art is very advanced

The simplex method has been around for 60 years

- Value of method means that the computational state-of-the-art is very advanced
- There is still scope for further improvements

The simplex method has been around for 60 years

- Value of method means that the computational state-of-the-art is very advanced
- There is still scope for further improvements
- Making advances is demanding but attracts attention!

The simplex method has been around for 60 years

- Value of method means that the computational state-of-the-art is very advanced
- There is still scope for further improvements
- Making advances is demanding but attracts attention!

Slides available as

http://www.maths.ed.ac.uk/hall/RealSimplex/