

High performance computational techniques for the simplex method

Julian Hall

School of Mathematics
University of Edinburgh
jajhall@ed.ac.uk

CO@Work 2020

17 September 2024



THE UNIVERSITY
of EDINBURGH



- Computational view of simplex algorithms
- Serial techniques
 - Hyper-sparsity
 - Cost perturbation
- Parallel techniques for general LP problems

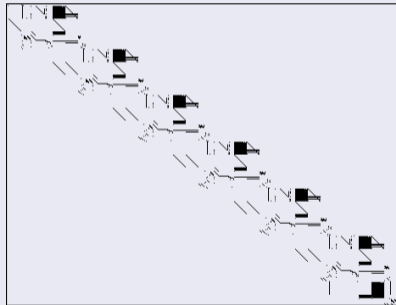
Solving LP problems: Background

$$\text{minimize } \mathbf{c}^T \mathbf{x} \quad \text{such that} \quad \mathbf{Ax} = \mathbf{b} \quad \text{and} \quad \mathbf{x} \geq \mathbf{0}$$

Background


- Fundamental model in optimal decision-making
- Solution techniques
 - Simplex method (1947)
 - Interior point methods (1984)
 - First order methods (2021)
- Large problems have
 - 10^3 – 10^8 variables
 - 10^3 – 10^8 constraints
- Matrix A is usually **sparse** and may be **structured**

Example



STAIR: 356 rows, 467 columns and 3856 nonzeros

Solving LP problems: Background



minimize $\mathbf{c}^T \mathbf{x}$ such that $A\mathbf{x} = \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$

- A **vertex** of the **feasible region** $K \subset \mathbb{R}^n$ has
 - m **basic** components, $i \in \mathcal{B}$
 - $n - m$ zero **nonbasic** components, $j \in \mathcal{N}$
- A and \mathbf{x} are partitioned according to $\mathcal{B} \cup \mathcal{N}$

$$B\mathbf{x}_B + N\mathbf{x}_N = \mathbf{b} \Rightarrow \mathbf{x}_B = B^{-1}(\mathbf{b} - N\mathbf{x}_N) = \hat{\mathbf{b}} - \hat{N}\mathbf{x}_N$$

since the **basis matrix** B is nonsingular

- Reduced objective is then $f = \hat{f} + \hat{\mathbf{c}}_N^T \mathbf{x}_N$, where $\hat{f} = \mathbf{c}_B^T \hat{\mathbf{b}}$ and $\hat{\mathbf{c}}_N^T = \mathbf{c}_N^T - \mathbf{c}_B^T B^{-1}N$
- For $\mathbf{x}_N = \mathbf{0}$, partition yields an optimal solution if there is **Primal feasibility** $\hat{\mathbf{b}} \geq \mathbf{0}$; **Dual feasibility** $\hat{\mathbf{c}}_N \geq \mathbf{0}$

- **Reduced LP** corresponding to partition $\mathcal{B} \cup \mathcal{N}$ of $\{1, \dots, n\}$ with B nonsingular is

$$\text{minimize } \hat{f} + \hat{\mathbf{c}}_N^T \mathbf{x}_N \quad \text{such that } \mathbf{x}_B + \hat{N} \mathbf{x}_N = \hat{\mathbf{b}} \quad \text{and } \mathbf{x} \geq \mathbf{0}$$

- Convenient to represent this in the **simplex tableau**

	\mathcal{N}	RHS
\mathcal{B}	\hat{N}	$\hat{\mathbf{b}}$
	$\hat{\mathbf{c}}_N^T$	

Primal simplex algorithm

Primal simplex algorithm: Choose a column

Assume $\hat{\mathbf{b}} \geq \mathbf{0}$ Seek $\hat{\mathbf{c}}_N \geq \mathbf{0}$

Scan $\hat{c}_j < 0$ for q to leave \mathcal{N}



	\mathcal{N}	RHS
\mathcal{B}		
	\hat{c}_q $\hat{\mathbf{c}}_N^T$	

Primal simplex algorithm: Choose a row

Assume $\hat{\mathbf{b}} \geq \mathbf{0}$ Seek $\hat{\mathbf{c}}_N \geq \mathbf{0}$

Scan $\hat{c}_j < 0$ for q to leave \mathcal{N}

Scan $\hat{b}_i / \hat{a}_{iq} < 0$ for p to leave \mathcal{B}

	\mathcal{N}	RHS
\mathcal{B}		

Primal simplex algorithm: Update cost and RHS

Assume $\hat{\mathbf{b}} \geq \mathbf{0}$ Seek $\hat{\mathbf{c}}_N \geq \mathbf{0}$

Scan $\hat{c}_j < 0$ for q to leave \mathcal{N}

Scan $\hat{b}_i / \hat{a}_{iq} < 0$ for p to leave \mathcal{B}

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

Update $\hat{\mathbf{b}} := \hat{\mathbf{b}} - \alpha_P \hat{\mathbf{a}}_q$ $\alpha_P = \hat{b}_p / \hat{a}_{pq}$

Update $\hat{\mathbf{c}}_N^T := \hat{\mathbf{c}}_N^T + \alpha_D \hat{\mathbf{a}}_p^T$ $\alpha_D = -\hat{c}_q / \hat{a}_{pq}$

	\mathcal{N}		RHS
\mathcal{B}	$\hat{\mathbf{a}}_q$		$\hat{\mathbf{b}}$
	\hat{a}_{pq}	$\hat{\mathbf{a}}_p^T$	\hat{b}_p
	\hat{c}_q	$\hat{\mathbf{c}}_N^T$	

Primal simplex algorithm: Data required

Assume $\hat{\mathbf{b}} \geq \mathbf{0}$ Seek $\hat{\mathbf{c}}_N \geq \mathbf{0}$

Scan $\hat{c}_j < 0$ for q to leave \mathcal{N}

Scan $\hat{b}_i / \hat{a}_{iq} < 0$ for p to leave \mathcal{B}

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

Update $\hat{\mathbf{b}} := \hat{\mathbf{b}} - \alpha_P \hat{\mathbf{a}}_q$ $\alpha_P = \hat{b}_p / \hat{a}_{pq}$

Update $\hat{\mathbf{c}}_N^T := \hat{\mathbf{c}}_N^T + \alpha_D \hat{\mathbf{a}}_p^T$ $\alpha_D = -\hat{c}_q / \hat{a}_{pq}$

Data required

Pivotal row $\hat{\mathbf{a}}_p^T = \mathbf{e}_p^T B^{-1} N$

Pivotal column $\hat{\mathbf{a}}_q = B^{-1} \mathbf{a}_q$

	\mathcal{N}		RHS
\mathcal{B}	$\hat{\mathbf{a}}_q$		$\hat{\mathbf{b}}$
	\hat{a}_{pq}	$\hat{\mathbf{a}}_p^T$	\hat{b}_p
	\hat{c}_q	$\hat{\mathbf{c}}_N^T$	

Primal simplex algorithm: Revised simplex method computation

Assume $\hat{\mathbf{b}} \geq \mathbf{0}$ Seek $\hat{\mathbf{c}}_N \geq \mathbf{0}$

Scan $\hat{c}_j < 0$ for q to leave \mathcal{N}

Scan $\hat{b}_i / \hat{a}_{iq} < 0$ for p to leave \mathcal{B}

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

Update $\hat{\mathbf{b}} := \hat{\mathbf{b}} - \alpha_P \hat{\mathbf{a}}_q$ $\alpha_P = \hat{b}_p / \hat{a}_{pq}$

Update $\hat{\mathbf{c}}_N^T := \hat{\mathbf{c}}_N^T + \alpha_D \hat{\mathbf{a}}_p^T$ $\alpha_D = -\hat{c}_q / \hat{a}_{pq}$

Data required

Pivotal row $\hat{\mathbf{a}}_p^T = \mathbf{e}_p^T B^{-1} N$ via $B^T \boldsymbol{\pi}_p = \mathbf{e}_p$; $\hat{\mathbf{a}}_p^T = \boldsymbol{\pi}_p^T N$

Pivotal column $\hat{\mathbf{a}}_q = B^{-1} \mathbf{a}_q$ via $B \hat{\mathbf{a}}_q = \mathbf{a}_q$

	\mathcal{N}		RHS
\mathcal{B}	$\hat{\mathbf{a}}_q$		$\hat{\mathbf{b}}$
	\hat{a}_{pq}	$\hat{\mathbf{a}}_p^T$	\hat{b}_p
	\hat{c}_q	$\hat{\mathbf{c}}_N^T$	

Dual simplex algorithm: Revised simplex method computation

Assume $\hat{\mathbf{c}}_N \geq \mathbf{0}$ Seek $\hat{\mathbf{b}} \geq \mathbf{0}$

Scan $\hat{b}_i < 0$ for p to leave \mathcal{B}

Scan $\hat{c}_j/\hat{a}_{pj} < 0$ for q to leave \mathcal{N}

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

Update $\hat{\mathbf{b}} := \hat{\mathbf{b}} - \alpha_P \hat{\mathbf{a}}_q$ $\alpha_P = \hat{b}_p/\hat{a}_{pq}$

Update $\hat{\mathbf{c}}_N^T := \hat{\mathbf{c}}_N^T + \alpha_D \hat{\mathbf{a}}_p^T$ $\alpha_D = -\hat{c}_q/\hat{a}_{pq}$

Data required

Pivotal row $\hat{\mathbf{a}}_p^T = \mathbf{e}_p^T B^{-1} N$ via $B^T \boldsymbol{\pi}_p = \mathbf{e}_p$; $\hat{\mathbf{a}}_p^T = \boldsymbol{\pi}_p^T N$

Pivotal column $\hat{\mathbf{a}}_q = B^{-1} \mathbf{a}_q$ via $B \hat{\mathbf{a}}_q = \mathbf{a}_q$

	\mathcal{N}		RHS
\mathcal{B}	$\hat{\mathbf{a}}_q$		$\hat{\mathbf{b}}$
	\hat{a}_{pq}	$\hat{\mathbf{a}}_p^T$	\hat{b}_p
	\hat{c}_q	$\hat{\mathbf{c}}_N^T$	

Solving LP problems: Primal or dual simplex?

Primal simplex algorithm

- Traditional variant
- Solution generally not primal feasible when (primal) LP is tightened

Dual simplex algorithm

- Preferred variant
- Easier to get dual feasibility
- More progress in many iterations
- Solution dual feasible when primal LP is tightened

In practice, both are required for efficiency and robustness

Simplex method: Computation

Standard simplex method (SSM): Major computational component

	\mathcal{N}	RHS
\mathcal{B}	\hat{N}	$\hat{\mathbf{b}}$
	$\hat{\mathbf{c}}_N^T$	

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

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

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

Revised simplex method (RSM): Major computational components

Pivotal row via $B^T \boldsymbol{\pi}_p = \mathbf{e}_p$ **BTRAN** and $\hat{\mathbf{a}}_p^T = \boldsymbol{\pi}_p^T N$ **PRICE**

Pivotal column via $B \hat{\mathbf{a}}_q = \mathbf{a}_q$ **FTRAN** Represent B^{-1} **INVERT**

Update B^{-1} exploiting $\bar{B} = B + (\mathbf{a}_q - B\mathbf{e}_p)\mathbf{e}_p^T$ **UPDATE**

Mittelmann LP test set (2020)

Industry standard set of 40 LP problems

	Rows	Cols	Nonzeros	$\frac{\text{Rows}}{\text{Cols}}$	$\frac{\text{Nonzeros}}{\text{Rows} \times \text{Cols}}$	$\frac{\text{Nonzeros}}{\max(\text{Rows}, \text{Cols})}$
Min	960	1560	38304	1/255	0.0005%	2.2
Geomean	54256	72442	910993	0.75	0.02%	6.5
Max	986069	1259121	11279748	85	16%	218.0

Mittelmann measure for solvers


- Unsolved problems given “timeout” solution time
- Shift all solution times up by 10s
- Compute geometric mean of logs of shifted times
- **Solution time measure** is exponent of geometric mean shifted down by 10s
- **Mittelmann measure** for a solver is its solution time measure relative to the best

Hyper-sparsity: Solve $B\mathbf{x} = \mathbf{r}$ for sparse \mathbf{r}


- In simplex, RHS of linear systems is sparse: column from A or unit vector
- When \mathbf{r} is sparse, solution $\mathbf{x} = B^{-1}\mathbf{r}$ combines a few columns of B^{-1}
- Although B^{-1} is never formed explicitly, studying it is instructive

Inverse of a sparse matrix and solution of $B\mathbf{x} = \mathbf{r}$

Optimal B for LP problem STAIR




$B^{-1}\mathbf{r}$ is typically dense




Hyper-sparsity: Solve $Bx = r$ for sparse r

Inverse of a sparse matrix and solution of $Bx = r$

Optimal B for LP problem PDS-02



$B^{-1}r$ is typically **sparse**



- If B^{-1} is sparse, then the LP is said to be **hyper-sparse**
- Huge performance gains from exploiting this in the simplex method

H and McKinnon (2005)

Hyper-sparsity: Effectiveness

Testing environment

- Mittelman test set of 40 LPs
- HiGHS dual simplex solver with/without exploiting hyper-sparsity
- Time limit of 10,000 seconds

Results

- When exploiting hyper-sparsity: solves 37 problems
- When not exploiting hyper-sparsity: solves 34 problems

	Min	Geomean	Max
Iteration count increase	0.75	1.08	3.17
Solution time increase	0.83	2.31	67.13
Iteration speed decrease	0.92	2.14	66.43
Mittelman measure	2.57		

Dual simplex: Cost perturbation

Dual degeneracy

- If some nonbasic dual values $\mathbf{c}_N^T - \mathbf{c}_B^T B^{-1} N$ are zero, the vertex is **dual degenerate**
- At a dual degenerate vertex, an iteration of the dual simplex algorithm may not lead to a strict increase in the dual objective
- Stalling or cycling may occur

Cost perturbation

- Add a small random value to some/all of the cost coefficients \mathbf{c}
- Nonbasic dual values then (at worst) take small positive values
- An iteration of the dual simplex algorithm yields (at least) a small positive increase in the dual objective
- When optimal, remove perturbations
- May require primal simplex iterations to regain optimality

Results using Mittelmann test set

- With cost perturbation: HiGHS solves 37/40 problems
- Without cost perturbation: solves 27 problems

	Min	Geomean	Max
Iteration count increase	0.80	1.36	7.21
Solution time increase	0.57	1.46	13.31
Iteration speed decrease	0.49	1.07	4.02
Mittelmann measure		3.80	

Parallel simplex for general LP problems

Past work

- High value problem: many attempts, but almost nothing of practical value
 - Parallel tableau simplex: “easy” but useless
 - Parallel PRICE $\pi_p^T N$: “easy” but Amdahl reigns unless $m \ll n$
- Crazy asynchronous schemes: H and McKinnon (mid-90s)

State-of-the-art

- High performance serial dual simplex solver with standard algorithmic enhancements (`hsol`)
- 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`)

Huangfu and H

Multiple iteration parallelism

- 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

	\mathcal{N}	RHS
\mathcal{B}	$\widehat{\mathbf{a}}_{\mathcal{P}}^T$	$\widehat{\mathbf{b}}$ $\widehat{b}_{\mathcal{P}}$
	$\widehat{\mathbf{c}}_N^T$	

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

Huangfu and H (2011–2014)

Serial overhead of pami

- HiGHS pami solver in serial: solves 34/40 problems

	Min	Geomean	Max
Iteration count increase	0.43	1.02	2.98
Solution time increase	0.31	1.62	5.36
Iteration speed decrease	0.69	1.59	5.11
Mittelmann measure	2.08		

Parallel speed-up of pami with 8 threads

	Min	Geomean	Max
Iteration count decrease	1.00	1.00	1.00
Solution time decrease	1.15	1.88	2.39
Iteration speed increase	1.15	1.88	2.39

Performance enhancement using parallel pami with 8 threads

	Min	Geomean	Max
Iteration count decrease	0.34	0.98	2.34
Solution time decrease	0.34	1.16	6.44
Iteration speed increase	0.38	1.18	2.75
Mittelman measure	1.21		

Observations

- There is significant scope to improve pami performance further
- Use pami tactically: switch it off if it is ineffective

Commercial impact

- Huangfu applied the parallel dual simplex techniques within the Xpress solver
- For much of 2013–2018 the Xpress simplex solver was the best in the world

PDLP: A first order method for solving LPs

- PDLP is a new method for solving

$$\text{minimize } f = \mathbf{c}^T \mathbf{x} \quad \text{such that } \mathbf{A}\mathbf{x} = \mathbf{b}; \mathbf{x} \geq \mathbf{0}$$

- Finds a saddle-point of $\min_{\mathbf{x} \geq \mathbf{0}} \max_{\mathbf{y} \geq \mathbf{0}} L(\mathbf{x}, \mathbf{y}) := \mathbf{c}^T \mathbf{x} - \mathbf{y}^T \mathbf{A}\mathbf{x} + \mathbf{b}^T \mathbf{y}$
- Uses primal-dual hybrid gradient Chambolle-Pock (2011) update

$$\mathbf{x}^{t+1} = [\mathbf{x}^t - \tau(\mathbf{c} - \mathbf{A}^T \mathbf{y}_t)]^+; \quad \mathbf{y}^{t+1} = [\mathbf{y}^t + \mu \mathbf{A}(\mathbf{A}\mathbf{x}^{t+1} - \mathbf{b})]^+$$

- Google's implementation in C++ on CPU is in the Mittelmann benchmarks
(Applegate, Hinder, Lu, and Lubin – 2021)
- Better results for cuPDLP-C implementation in C+CUDA on GPU
(Lu, Yang, Hu, Huangfu, Liu, Liu, Ye, Zhang – Dec 2023)
 - Available under MIT license on GitHub using HiGHS for file reading and presolve
 - CPU/GPU version COPT v7.0
 - CPU version in HiGHS v1.7.0

(March 2024)

HiGHS: Open-source software for large-scale sparse linear optimization

HiGHS: **H**all, **i**vet **G**alabova, **H**uangfu, **S**chork

$$\text{minimize } f = \frac{1}{2} \mathbf{x}^T \mathbf{Q} \mathbf{x} + \mathbf{c}^T \mathbf{x} \quad \text{such that } \mathbf{A} \mathbf{x} = \mathbf{b}; \mathbf{x} \geq \mathbf{0}, \quad x_i \in \mathbb{Z}, \forall i \in \mathcal{I}$$

Features

- Simplex, interior point and first order solvers for LP
- Branch-and-cut solver for MIP
- Active set solver for QP
- Written in C++
- Interfaces to other languages and systems



Availability

- Open-source (MIT license)
- No third-party code
- <https://HiGHS.dev/>

The world's best open-source linear optimization software

Mittelmann (2022–date)

Practical LP solution

Must exploit sparsity (and maybe structure)

- Simplex method
- Interior point methods
- First order methods

High performance simplex

- Best when solving families of related problems (MIP; SLP)
- Many (more) algorithmic and computational tricks in serial
- Parallel simplex has some impact on performance



D. Applegate, M. Díaz, O. Hinder, H. Lu, M. Lubin, B. O'Donoghue, and W. Schudy.
Practical large-scale linear programming using primal-dual hybrid gradient.
[Advances in Neural Information Processing Systems](#), 34:20243–20257, 2021.



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.
[Mathematical Programming Computation](#), 10(1):119–142, 2018.