Computational linear optimization

Julian Hall

School of Mathematics

University of Edinburgh

March 15th 2007

Linear programming problems

$$\begin{aligned} \text{minimize} \quad & f = \boldsymbol{c}^T \boldsymbol{x} \\ \text{subject to} \quad & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

Linear programming problems

$$\begin{array}{rl} \text{minimize} & f = \boldsymbol{c}^T \boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Linear programming (LP) is the fundamental model in optimal decision-making

Linear programming problems

$$\begin{array}{rl} \text{minimize} & f = \boldsymbol{c}^T \boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Linear programming (LP) is the fundamental model in optimal decision-making
- Solution techniques
  - Simplex method (1947)
  - Interior point method (1984)

Linear programming problems

$$\begin{array}{rl} \text{minimize} & f = \boldsymbol{c}^T \boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Linear programming (LP) is the fundamental model in optimal decision-making
- Solution techniques
  - Simplex method (1947)
  - Interior point method (1984)
- Large problems have
  - $10^3 - 10^7$ variables
  - $10^3 - 10^7$ constraints

Linear programming problems

$$\begin{array}{ll} \text{minimize} & f = \boldsymbol{c}^T \boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Linear programming (LP) is the fundamental model in optimal decision-making
- Solution techniques
  - Simplex method (1947)
  - Interior point method (1984)
- Large problems have
  - $10^3 - 10^7$ variables
  - $10^3 - 10^7$ constraints
- Matrix A is sparse

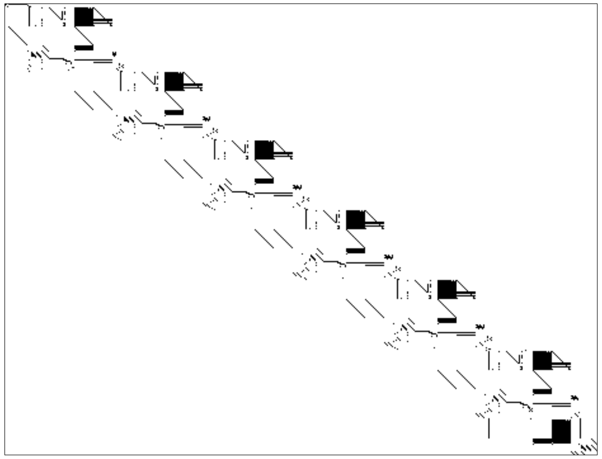

STAIR: 356 rows, 467 columns and 3856 nonzeros

Linear programming problems

$$\begin{array}{rl} \text{minimize} & f = \boldsymbol{c}^T\boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Linear programming (LP) is the fundamental model in optimal decision-making
- Solution techniques
  - Simplex method (1947)
  - Interior point method (1984)
- Large problems have
  - $10^3 - 10^7$ variables
  - $10^3 - 10^7$ constraints
- Matrix A is sparse

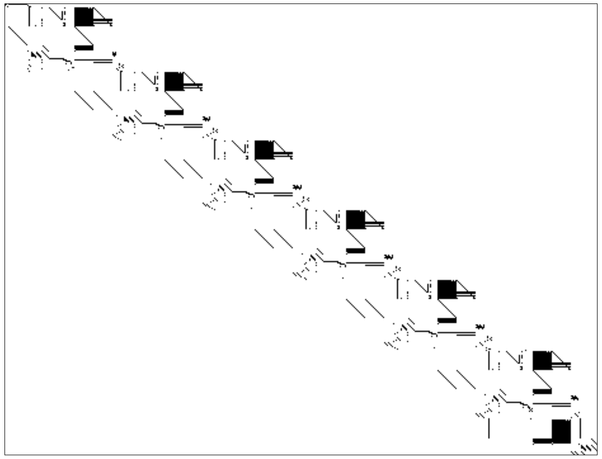

STAIR: 356 rows, 467 columns and 3856 nonzeros

- Structure as much as size determines the computational challenge

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$\begin{array}{rcll} A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \mathbf{0} & \text{primal feasibility} \\ A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \mathbf{0} & \text{dual feasibility} \\ \boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & \text{complementarity condition} \end{array}$$

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$\begin{array}{rcll} A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & \text{primal feasibility} \\ A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & \text{dual feasibility} \\ \boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & \text{complementarity condition} \end{array}$$

Simplex method

- Uses a combinatorial approach

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$
\begin{array}{rlll}
A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & \text{primal feasibility} \\
A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & \text{dual feasibility} \\
\boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & \text{complementarity condition}
\end{array}
$$

Simplex method

- Uses a combinatorial approach
- Modifies a partition $\mathcal{B} \cup \mathcal{N}$ of variables until $\boldsymbol{s} \geq \boldsymbol{0}$

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$\begin{array}{rcll} A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & \text{primal feasibility} \\ A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & \text{dual feasibility} \\ \boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & \text{complementarity condition} \end{array}$$

Simplex method

- Uses a combinatorial approach
- Modifies a partition $\mathcal{B} \cup \mathcal{N}$ of variables until $\boldsymbol{s} \geq \boldsymbol{0}$
- Moves along edges of the feasible region

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$\begin{array}{rclll} A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & & \text{primal feasibility} \\ A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & & \text{dual feasibility} \\ \boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 & \forall j & \text{complementarity condition} \end{array}$$

Simplex method

- Uses a combinatorial approach
- Modifies a partition $\mathcal{B} \cup \mathcal{N}$ of variables until $\boldsymbol{s} \geq \boldsymbol{0}$
- Moves along edges of the feasible region
- Terminates at an optimal vertex

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$
\begin{array}{rcll}
A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & \text{primal feasibility} \\
A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & \text{dual feasibility} \\
\boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & \text{complementarity condition}
\end{array}
$$

Simplex method

- Uses a combinatorial approach
- Modifies a partition $\mathcal{B} \cup \mathcal{N}$ of variables until $\boldsymbol{s} \geq \boldsymbol{0}$
- Moves along edges of the feasible region
- Terminates at an optimal vertex

Interior point methods (IPM)

- Use an iterative approach

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$\begin{array}{rcll} A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & \text{primal feasibility} \\ A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & \text{dual feasibility} \\ \boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & \text{complementarity condition} \end{array}$$

Simplex method

- Uses a combinatorial approach
- Modifies a partition $\mathcal{B} \cup \mathcal{N}$ of variables until $\boldsymbol{s} \geq \boldsymbol{0}$
- Moves along edges of the feasible region
- Terminates at an optimal vertex

Interior point methods (IPM)

- Use an iterative approach
- Reduce $x_j s_j = \mu$ to zero $\forall j$

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$\begin{array}{rcll} A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & \text{primal feasibility} \\ A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & \text{dual feasibility} \\ \boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & \text{complementarity condition} \end{array}$$

Simplex method

- Uses a combinatorial approach
- Modifies a partition $\mathcal{B} \cup \mathcal{N}$ of variables until $\boldsymbol{s} \geq \boldsymbol{0}$
- Moves along edges of the feasible region
- Terminates at an optimal vertex

Interior point methods (IPM)

- Use an iterative approach
- Reduce $x_j s_j = \mu$ to zero $\forall j$
- Move through the interior of the feasible region

Optimality conditions for LP

Lagrange multipliers $\boldsymbol{y}$ and $\boldsymbol{s}$ exist such that the following conditions hold

$$
\begin{array}{rclll}
A\boldsymbol{x} = \boldsymbol{b} & & \boldsymbol{x} \geq \boldsymbol{0} & & \text{primal feasibility} \\
A^T\boldsymbol{y} + \boldsymbol{s} = \boldsymbol{c} & & \boldsymbol{s} \geq \boldsymbol{0} & & \text{dual feasibility} \\
\boldsymbol{s}^T\boldsymbol{x} = 0 & \Rightarrow & x_j s_j = 0 \quad \forall j & & \text{complementarity condition}
\end{array}
$$

Simplex method

- Uses a combinatorial approach
- Modifies a partition $\mathcal{B} \cup \mathcal{N}$ of variables until $\boldsymbol{s} \geq \boldsymbol{0}$
- Moves along edges of the feasible region
- Terminates at an optimal vertex

Interior point methods (IPM)

- Use an iterative approach
- Reduce $x_j s_j = \mu$ to zero $\forall j$
- Move through the interior of the feasible region
- Converge to within ϵ of an optimal vertex

Simplex vs interior point methods

Simplex method

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**

Interior point methods

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1}\boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1}\boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**
- Complexity
  - Theory: $O(2^n)$ iterations

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1}\boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**

- Complexity
  - Theory: $O(2^n)$ iterations
  - Practice: $O(n)$ iterations

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1} \boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**
- Complexity
  - Theory: $O(2^n)$ iterations
  - Practice: $O(n)$ iterations

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1}\boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**
- Complexity
  - Theory: $O(\sqrt{n}\log n)$ iterations

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**
- Complexity
  - Theory: $O(2^n)$ iterations
  - Practice: $O(n)$ iterations

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1} \boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**
- Complexity
  - Theory: $O(\sqrt{n}\log n)$ iterations
  - Practice: $O(\log n)$ iterations

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**
- Complexity
  - Theory: $O(2^n)$ iterations
  - Practice: $O(n)$ iterations

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1} \boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**
- Complexity
  - Theory: $O(\sqrt{n}\log n)$ iterations
  - Practice: $O(\log n)$ iterations

Questions

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**
- Complexity
  - Theory: $O(2^n)$ iterations
  - Practice: $O(n)$ iterations

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1} \boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**
- Complexity
  - Theory: $O(\sqrt{n}\log n)$ iterations
  - Practice: $O(\log n)$ iterations

Questions

- Why is the simplex method still competitive?

Simplex vs interior point methods

Simplex method

- Each iteration
  - Form $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$
  - Each $[B : N]$ is a partition of A
  - $\boldsymbol{r}_F$, $\boldsymbol{r}_B$ (and $\boldsymbol{r}_\pi$) are **sparse**
- Complexity
  - Theory: $O(2^n)$ iterations
  - Practice: $O(n)$ iterations

Interior point methods

- Each iteration
  - Form $(A\Theta A^T)^{-1}\boldsymbol{r}$ or $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}^{-1}\boldsymbol{r}$
  - Diagonal matrix Θ is iteration-dependent
  - $\boldsymbol{r}$ is **full**
- Complexity
  - Theory: $O(\sqrt{n}\log n)$ iterations
  - Practice: $O(\log n)$ iterations

Questions

- Why is the simplex method still competitive?
- What are the computational challenges?

Simplex: achievements and challenges

- Practical problems whose structure yields sparse B^{-1} are **hyper-sparse**

Simplex: achievements and challenges

- Practical problems whose structure yields sparse B^{-1} are **hyper-sparse**
  - $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$ are sparse vectors

Simplex: achievements and challenges

- Practical problems whose structure yields sparse B^{-1} are **hyper-sparse**
  - $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$ are sparse vectors
  - Exploiting hyper-sparsity—Hall and McKinnon won COAP prize for 2005
  - Simplex now out-performs IPM for hyper-sparse LPs

Simplex: achievements and challenges

- Practical problems whose structure yields sparse B^{-1} are **hyper-sparse**
  - $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$ are sparse vectors
  - Exploiting hyper-sparsity—Hall and McKinnon won COAP prize for 2005
  - Simplex now out-performs IPM for hyper-sparse LPs
  - **Challenge:** better understanding and promotion of hyper-sparsity

Simplex: achievements and challenges

- Practical problems whose structure yields sparse B^{-1} are **hyper-sparse**
  - $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$ are sparse vectors
  - Exploiting hyper-sparsity—Hall and McKinnon won COAP prize for 2005
  - Simplex now out-performs IPM for hyper-sparse LPs
  - **Challenge:** better understanding and promotion of hyper-sparsity
- Parallel simplex solver
  - No useful codes

Simplex: achievements and challenges

- Practical problems whose structure yields sparse B^{-1} are **hyper-sparse**
  - $B^{-1}\boldsymbol{r}_F$, $\boldsymbol{r}_B^T B^{-1}$ and $N\boldsymbol{r}_\pi$ are sparse vectors
  - Exploiting hyper-sparsity—Hall and McKinnon won COAP prize for 2005
  - Simplex now out-performs IPM for hyper-sparse LPs
  - **Challenge:** better understanding and promotion of hyper-sparsity
- Parallel simplex solver
  - No useful codes
  - **Challenge:** develop one!

Interior point methods: achievements and challenges

- Interior point methods better than simplex for many problems

Interior point methods: achievements and challenges

- Interior point methods better than simplex for many problems

  Challenge: Maintain sparsity when factoring $A\Theta A^T$

Interior point methods: achievements and challenges

- Interior point methods better than simplex for many problems
  Challenge: Maintain sparsity when factoring $A\Theta A^T$
- **But:** $A\Theta A^T$ can fill in badly

Interior point methods: achievements and challenges

- Interior point methods better than simplex for many problems

  Challenge: Maintain sparsity when factoring $A\Theta A^T$
- **But:** $A\Theta A^T$ can fill in badly
  - Network LP problems are *very* sparse
  - Constraint (node-arc incidence) matrix for the graph

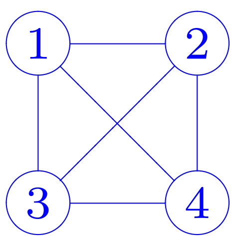

is $A = \begin{bmatrix} \bullet & & \bullet & \bullet & & \\ \bullet & \bullet & & & & \bullet \\ & \bullet & & \bullet & \bullet & \\ & & \bullet & & \bullet & \bullet \end{bmatrix}$

Interior point methods: achievements and challenges

- Interior point methods better than simplex for many problems

  Challenge: Maintain sparsity when factoring $A\Theta A^T$
- **But:** $A\Theta A^T$ can fill in badly
  - Network LP problems are *very* sparse
  - Constraint (node-arc incidence) matrix for the graph

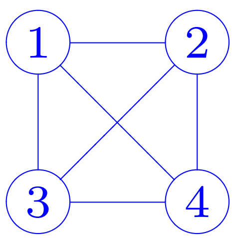

is $A = \begin{bmatrix} \bullet & & \bullet & \bullet & & \\ \bullet & \bullet & & & & \bullet \\ & \bullet & & \bullet & \bullet & \\ & & \bullet & & \bullet & \bullet \end{bmatrix}$ but $A\Theta A^T = \begin{bmatrix} \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \end{bmatrix}$

Interior point methods: achievements and challenges

- Interior point methods better than simplex for many problems

  Challenge: Maintain sparsity when factoring $A\Theta A^T$

- **But:** $A\Theta A^T$ can fill in badly
  - Network LP problems are *very* sparse
  - Constraint (node-arc incidence) matrix for the graph

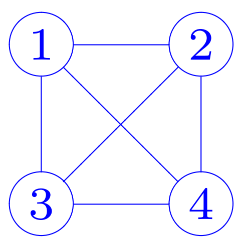

$$\text{is} \quad A = \begin{bmatrix} \bullet & & \bullet & \bullet & & \\ \bullet & \bullet & & & & \bullet \\ & \bullet & & \bullet & \bullet & \\ & & \bullet & & \bullet & \bullet \end{bmatrix} \quad \text{but} \quad A\Theta A^T = \begin{bmatrix} \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \end{bmatrix}$$

  - Network(-like) LPs are hyper-sparse and simplex out-performs IPM

Interior point methods: achievements and challenges

- Interior point methods better than simplex for many problems

  Challenge: Maintain sparsity when factoring $A\Theta A^T$
- **But:** $A\Theta A^T$ can fill in badly
  - Network LP problems are *very* sparse
  - Constraint (node-arc incidence) matrix for the graph

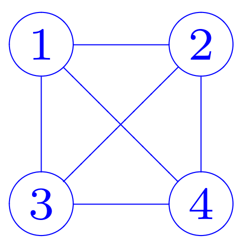

$$\text{is} \quad A = \begin{bmatrix} \bullet & & \bullet & \bullet & & \\ \bullet & \bullet & & & & \bullet \\ & \bullet & & \bullet & \bullet & \\ & & \bullet & & \bullet & \bullet \end{bmatrix} \quad \text{but} \quad A\Theta A^T = \begin{bmatrix} \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \end{bmatrix}$$

  - Network(-like) LPs are hyper-sparse and simplex out-performs IPM
  - **Challenge:** Factor $\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix}$

Quadratic programming problems

$$\begin{aligned} \text{minimize} \quad & f = \boldsymbol{c}^T\boldsymbol{x} + \tfrac{1}{2}\boldsymbol{x}^T Q\boldsymbol{x} \\ \text{subject to} \quad & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{aligned}$$

Quadratic programming problems

$$\begin{array}{rl} \text{minimize} & f = \boldsymbol{c}^T\boldsymbol{x} + \frac{1}{2}\boldsymbol{x}^T Q\boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Solved by natural extension to interior point methods for LP

Quadratic programming problems

$$\begin{array}{rl} \text{minimize} & f = \boldsymbol{c}^T\boldsymbol{x} + \frac{1}{2}\boldsymbol{x}^T Q \boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Solved by natural extension to interior point methods for LP
- Each iteration
  - Form

$$(A[\Theta^{-1} + Q]^{-1}A^T)^{-1}\boldsymbol{r} \qquad \text{or} \qquad \begin{bmatrix} -\Theta^{-1} - Q & A^T \\ A & 0 \end{bmatrix}^{-1}\boldsymbol{r}$$

  - Matrix $\Theta^{-1} + Q$ is non-diagonal so $A[\Theta^{-1} + Q]^{-1}A^T$ is less likely to be sparse

Quadratic programming problems

$$\begin{array}{ll} \text{minimize} & f = \boldsymbol{c}^T\boldsymbol{x} + \frac{1}{2}\boldsymbol{x}^T Q \boldsymbol{x} \\ \text{subject to} & A\boldsymbol{x} = \boldsymbol{b} \quad \boldsymbol{x} \geq \boldsymbol{0} \end{array}$$

- Solved by natural extension to interior point methods for LP
- Each iteration
  - Form

$$(A[\Theta^{-1} + Q]^{-1} A^T)^{-1} \boldsymbol{r} \qquad \text{or} \qquad \begin{bmatrix} -\Theta^{-1} - Q & A^T \\ A & 0 \end{bmatrix}^{-1} \boldsymbol{r}$$

  - Matrix $\Theta^{-1} + Q$ is non-diagonal so $A[\Theta^{-1} + Q]^{-1} A^T$ is less likely to be sparse
- **Challenge:** nonconvex QP (with Q indefinite)

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A\Theta A^T$ encourages the use of iterative methods

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A\Theta A^T$ encourages the use of iterative methods
- $A[\Theta^{-1} + Q]^{-1}A^T$ is positive definite: consider preconditioned conjugate gradients (PCG)

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A\Theta A^T$ encourages the use of iterative methods
- $A[\Theta^{-1} + Q]^{-1}A^T$ is positive definite: consider preconditioned conjugate gradients (PCG)
- Exciting and fruitful cross-disciplinary activity between optimization and PDEs

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A\Theta A^T$ encourages the use of iterative methods
- $A[\Theta^{-1}+Q]^{-1}A^T$ is positive definite: consider preconditioned conjugate gradients (PCG)
- Exciting and fruitful cross-disciplinary activity between optimization and PDEs

Interior point methods

- Convex QP has Q positive definite in

$$\begin{bmatrix} -\Theta^{-1}-Q & A^T \\ A & 0 \end{bmatrix}$$

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A\Theta A^T$ encourages the use of iterative methods
- $A[\Theta^{-1} + Q]^{-1} A^T$ is positive definite: consider preconditioned conjugate gradients (PCG)
- Exciting and fruitful cross-disciplinary activity between optimization and PDEs

Interior point methods

- Convex QP has Q positive definite in

$$\begin{bmatrix} -\Theta^{-1} - Q & A^T \\ A & 0 \end{bmatrix}$$

- Tailored PCG is possible

  Bergamaschi, Gondzio and Zilli won COAP prize for 2004

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A\Theta A^T$ encourages the use of iterative methods
- $A[\Theta^{-1}+Q]^{-1}A^T$ is positive definite: consider preconditioned conjugate gradients (PCG)
- Exciting and fruitful cross-disciplinary activity between optimization and PDEs

Interior point methods

- Convex QP has Q positive definite in

$$\begin{bmatrix} -\Theta^{-1}-Q & A^T \\ A & 0 \end{bmatrix}$$

- Tailored PCG is possible

  Bergamaschi, Gondzio and Zilli won COAP prize for 2004

- **Challenge:** Improved preconditioners

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A\Theta A^T$ encourages the use of iterative methods
- $A[\Theta^{-1} + Q]^{-1}A^T$ is positive definite: consider preconditioned conjugate gradients (PCG)
- Exciting and fruitful cross-disciplinary activity between optimization and PDEs

Interior point methods

- Convex QP has Q positive definite in

$$\begin{bmatrix} -\Theta^{-1} - Q & A^T \\ A & 0 \end{bmatrix}$$

- Tailored PCG is possible

  Bergamaschi, Gondzio and Zilli won COAP prize for 2004

- **Challenge:** Improved preconditioners

Saddle point problems from PDEs

- Elliptic PDE has H positive definite in

$$\begin{bmatrix} -H & A^T \\ A & C \end{bmatrix}$$

Iterative methods for linear systems in interior point methods

- Potential loss of sparsity when forming $A \Theta A^T$ encourages the use of iterative methods
- $A[\Theta^{-1} + Q]^{-1} A^T$ is positive definite: consider preconditioned conjugate gradients (PCG)
- Exciting and fruitful cross-disciplinary activity between optimization and PDEs

Interior point methods

- Convex QP has Q positive definite in

$$\begin{bmatrix} -\Theta^{-1} - Q & A^T \\ A & 0 \end{bmatrix}$$

- Tailored PCG is possible

  Bergamaschi, Gondzio and Zilli won COAP prize for 2004

- **Challenge:** Improved preconditioners

Saddle point problems from PDEs

- Elliptic PDE has H positive definite in

$$\begin{bmatrix} -H & A^T \\ A & C \end{bmatrix}$$

- Many special preconditioners

Object-Oriented Parallel Solver (OOPS)

- Gondzio, Grothey and Sarkissian
- Exploit nested block-angular structure

Object-Oriented Parallel Solver (OOPS)

- Gondzio, Grothey and Sarkissian
- Exploit nested block-angular structure
- Results for asset liability management problems with mean-variance QP formulation

problem	scenarios	constraints	variables	iterations	time (s)	processors	machine
ALM8	12,831,873	64,159,366	153,982,477	42	3923	512	BlueGene
ALM9	6,415,937	96,239,056	269,469,355	39	4692	512	BlueGene
ALM10	12,831,873	179,646,223	500,443,048	45	6089	1024	BlueGene
ALM11	16,039,809	352,875,799	1,010,507,968	53	3020	1280	HPCx

Object-Oriented Parallel Solver (OOPS)

- Gondzio, Grothey and Sarkissian
- Exploit nested block-angular structure
- Results for asset liability management problems with mean-variance QP formulation

problem	scenarios	constraints	variables	iterations	time (s)	processors	machine
ALM8	12,831,873	64,159,366	153,982,477	42	3923	512	BlueGene
ALM9	6,415,937	96,239,056	269,469,355	39	4692	512	BlueGene
ALM10	12,831,873	179,646,223	500,443,048	45	6089	1024	BlueGene
ALM11	16,039,809	352,875,799	1,010,507,968	53	3020	1280	HPCx

- ALM11 is the largest optimization problem ever solved directly!

Conclusions

- Healthy competition between simplex and interior point methods

Conclusions

- Healthy competition between simplex and interior point methods
- University of Edinburgh group is playing a leading role in addressing the **major challenges:**
  - Better sparse direct methods
  - Preconditioned iterative solvers for interior point methods

Conclusions

- Healthy competition between simplex and interior point methods
- University of Edinburgh group is playing a leading role in addressing the **major challenges:**
  - Better sparse direct methods
  - Preconditioned iterative solvers for interior point methods
  - Parallel simplex solver?

Conclusions

- Healthy competition between simplex and interior point methods
- University of Edinburgh group is playing a leading role in addressing the **major challenges:**
  - Better sparse direct methods
  - Preconditioned iterative solvers for interior point methods
  - Parallel simplex solver?
  - LP solver with simplex iteration cost and interior point iteration count??

Conclusions

- Healthy competition between simplex and interior point methods
- University of Edinburgh group is playing a leading role in addressing the **major challenges:**
  - Better sparse direct methods
  - Preconditioned iterative solvers for interior point methods
  - Parallel simplex solver?
  - LP solver with simplex iteration cost and interior point iteration count??

`http://www.maths.ed.ac.uk/hall/Talks`

Thank you