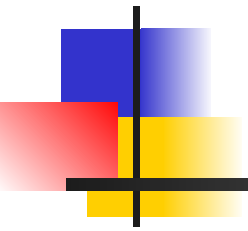

Randomized Numerical Linear Algebra for Interior Point Methods

Gregory Dexter

Department of Computer Science
Purdue University

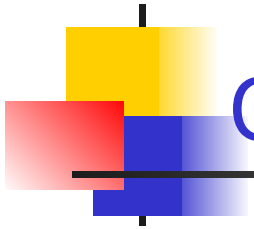

Overview

- Interior Point Methods (IPMs)
  - A class of algorithms to solve linear programs (LPs)
- Randomized Numerical Linear Algebra (RandNLA)
  - Studies randomized algorithms for linear algebra

We focus on using RandNLA to create **practical** fast IPMs to solve large-scale LPs

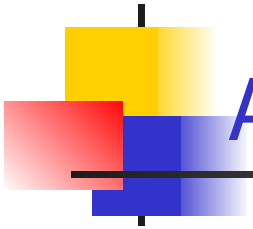

Applications of Linear Programming

- Machine learning
  - Non-negative matrix factorization
  - ℓ1-regularized SVMs
  - Robust linear regression
- Theoretical Computer Science
  - Network flow
  - Multicommodity flow
- Operations Research
  - Logistic routing
  - Balancing electrical grids

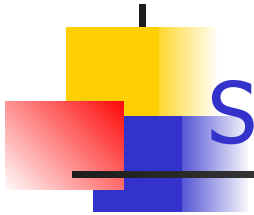

Standard Form Linear Programs

Consider the standard form of the primal LP problem:

$$\min \mathbf{c}^{\mathsf{T}}\mathbf{x}\,,\ \text{subject to } \mathbf{A}\mathbf{x} = \mathbf{b}\,, \mathbf{x} \geq \mathbf{0}$$

The associated dual problem is

$$\max \mathbf{b}^{\mathsf{T}}\mathbf{y}\,,\ \text{subject to } \mathbf{A}^{\mathsf{T}}\mathbf{y} + \mathbf{s} = \mathbf{c}\,, \mathbf{s} \geq \mathbf{0}$$

$\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, and $\mathbf{c} \in \mathbb{R}^n$ are inputs

$\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^m$, and $\mathbf{s} \in \mathbb{R}^n$ are variables

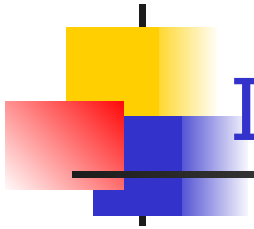

Interior Point Methods

- Path-following IPMs can broadly be categorized as:
  - Long step methods (worse theoretically, fast in practice)
  - Short step methods (better theoretically, slow in practice)
  - Predictor-corrector methods (good theoretically and in practice)

- We focus on the long step and predictor-corrector methods
  - Note: Other IPMs (e.g. potential reduction methods) and non-IPMs (e.g simplex method) exist to solve LPs as well
  - Path-following IPMs are typically best for large LPs

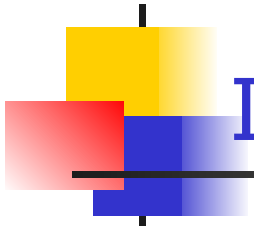

Interior Point Methods cont.

- IPMs work by traversing the interior of the feasible polytope while staying near the "central path"
  - Staying away from the edge of polytope improves convergence theoretically and in practice
- Runtime cost dominated by solving a linear system at each iteration
  - Steps are computed by solving "normal equations"
  - In large scale IPMs, only iterative linear solvers are efficient enough
  - More efficient solvers greatly improve overall runtime

Visualization of Path-Following IPMs

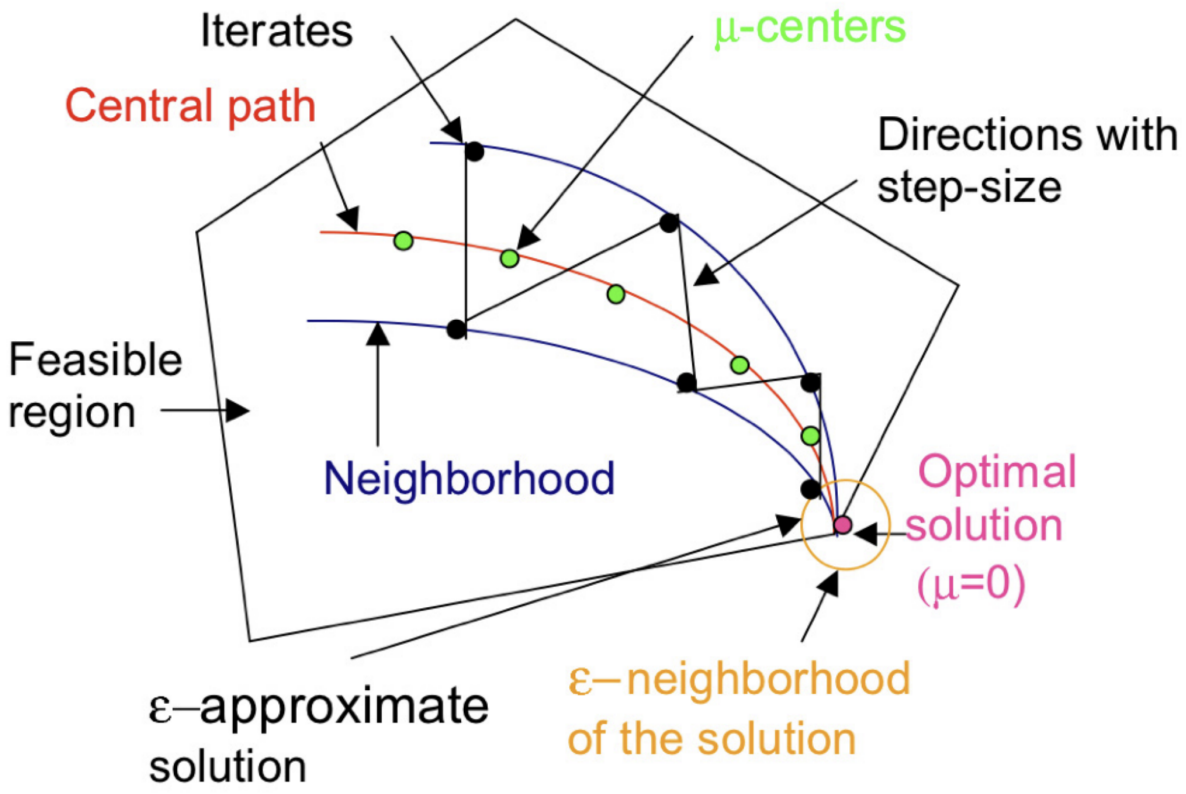

Figure from Lesaja, G. (2009). Introducing interior-point methods for introductory operations research courses and/or linear programming courses. *Open Operational Research Journal*, *3*, 1.

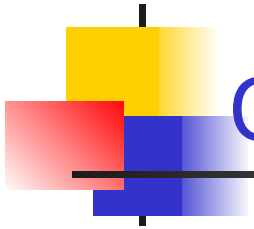

Overview of RandNLA

- Randomized algorithms for linear algebra/matrices
  - Combination of 1) Random matrix theory, 2) Matrix perturbation theory, 3) Numerical linear algebra (NLA)
- Provides many effective improvements to fundamental NLA algorithms
  - Subsampled least-squares regression
  - Approximate matrix multiplication
  - Online/Streaming norm approximation
- Provides theoretical tools to efficiently solve the normal equations with theoretical guarantees

RandNLA for Solving Linear Systems

[slide from H. Avron]

Sketch-and-Solve

1. High success rate
2. Polynomial accuracy dependence (e.g. ϵ^{-2})
3. No iterations

Pros:

1. **Very** fast
2. Deterministic running time

Cons:

1. Only crude accuracy
2. "Monte-Carlo" algorithm

Sketch-to-Precondition

1. High success rate
2. Exponential accuracy dependence (e.g. $\log(1/\epsilon)$)
3. Iterations

Pros:

1. Very high accuracy possible
2. Success = good solution

Cons:

1. Slower than sketch-and-solve
2. Iterations (no streaming)

Challenges to Using Iterative Solvers

Immediate problem: even assuming a feasible starting point, approximate solutions do not lead to feasible updates.

- As a result, standard analyses of the convergence of IPMs are not applicable.
- We use RandNLA approaches to efficiently and provably accurately correct the error induced by the approximate solution and guarantee convergence.

Details: the approximate solution violates critical equalities:

$$\mathbf{A}\mathbf{D}^2\mathbf{A}^\mathsf{T}\hat{\Delta}\mathbf{y} \neq \mathbf{p} \quad \text{and} \quad \mathbf{A}\hat{\Delta}\mathbf{x} \neq -\mathbf{r}_p \,/\, \mathbf{0}_m$$

- The vector r_p is the primal residual; for feasible long-step IPMs, it is the all-zero vector.
- Standard analyses of long-step (infeasible/feasible) IPMs critically need the second inequality to be an equality.
- Without the above equalities, in the case of feasible IPMs, we can not terminate with a feasible solution; we will end up with an approximately feasible solution.

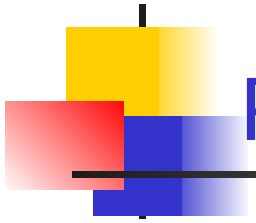

RandNLA & IPMs for LPs

Research Agenda: Explore how approximate, iterative solvers for the normal equations affect the convergence of

(1) long-step (feasible and infeasible) IPMs,

(2) feasible predictor-corrector IPMs.

- We seek to investigate **standard, practical solvers**, such as Preconditioned Conjugate Gradients, Preconditioned Steepest Descent, Preconditioned Richardson's iteration, etc.
- The preconditioner is constructed using RandNLA sketching-based approaches.
- **Remark:** For feasible path-following IPMs, an additional design choice is whether we want the final solution to be feasible or approximately feasible.

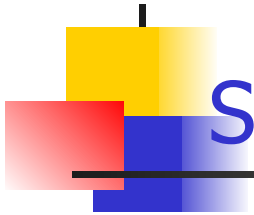

Standard Form Linear Programs

Consider the standard form of the primal LP problem:

$$\min \mathbf{c}^{\mathsf{T}}\mathbf{x}\,,\ \text{subject to } \mathbf{A}\mathbf{x} = \mathbf{b}\,, \mathbf{x} \geq \mathbf{0}$$

The associated dual problem is

$$\max \mathbf{b}^{\mathsf{T}}\mathbf{y}\,,\ \text{subject to } \mathbf{A}^{\mathsf{T}}\mathbf{y} + \mathbf{s} = \mathbf{c}\,, \mathbf{s} \geq \mathbf{0}$$

$\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, and $\mathbf{c} \in \mathbb{R}^n$ are inputs

$\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^m$, and $\mathbf{s} \in \mathbb{R}^n$ are variables

Interior Point Methods (IPMs)

(long-step)

- **Duality measure**:

$$\mu = \frac{\mathbf{x}^\mathsf{T}\mathbf{s}}{n} = \frac{\mathbf{x}^\mathsf{T}(\mathbf{c} - \mathbf{A}^\mathsf{T}\mathbf{y})}{n} = \frac{\mathbf{c}^\mathsf{T}\mathbf{x} - \mathbf{b}^\mathsf{T}\mathbf{y}}{n} \downarrow 0$$

- **Path-following methods:**
  - Let $\mathcal{F}^0 = \{(\mathbf{x}, \mathbf{y}, \mathbf{s}) : \ (\mathbf{x}, \mathbf{s}) > \mathbf{0}, \ \mathbf{A}\mathbf{x} = \mathbf{b}, \ \mathbf{A}^\mathsf{T}\mathbf{y} + \mathbf{s} = \mathbf{c}\}$.
  - Central path: $\mathcal{C} = \{(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{F}^0 \ : \ \mathbf{x} \circ \mathbf{s} = \tau \cdot \mathbf{1}_n; \ \tau > 0\}$
  - Neighborhood: $\mathcal{N}(\gamma) = \left\{(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{F}^0 : \mathbf{x} \circ \mathbf{s} \geq (1 - \gamma)\mu \mathbf{1}_n\right\}, \ \gamma \in (0, 1)$
  - Given the step size $\alpha \in [0, 1]$ and $(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{N}(\gamma)$, it computes the Newton search direction $(\Delta\mathbf{x}, \Delta\mathbf{y}, \Delta\mathbf{s})$ and update the current iterate

$$\boxed{(\mathbf{x}(\alpha), \mathbf{y}(\alpha), \mathbf{s}(\alpha)) = (\mathbf{x}, \mathbf{y}, \mathbf{s}) + \alpha(\Delta\mathbf{x}, \Delta\mathbf{y}, \Delta\mathbf{s}) \in \mathcal{N}(\gamma)}$$

IPM Normal Equations

Path-following IPMs, at every iteration, solve a system of linear equations :

$$\begin{pmatrix} \mathbf{A} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}^\mathsf{T} & \mathbf{I}_n \\ \mathbf{S} & \mathbf{0} & \mathbf{X} \end{pmatrix} \begin{pmatrix} \Delta\mathbf{x} \\ \Delta\mathbf{y} \\ \Delta\mathbf{s} \end{pmatrix} = \begin{pmatrix} -\mathbf{r}_p \\ -\mathbf{r}_d \\ -\mathbf{XS1}_n + \sigma\mu\mathbf{1}_n \end{pmatrix}$$

$\mathbf{D} = \mathbf{X}^{1/2}\mathbf{S}^{-1/2}$ is a diagonal matrix.

normal equations

$$\mathbf{AD}^2\mathbf{A}^\mathsf{T}\Delta\mathbf{y} = \underbrace{-\mathbf{r}_p - \sigma\mu\mathbf{AS}^{-1}\mathbf{1}_n + \mathbf{Ax} - \mathbf{AD}^2\mathbf{r}_d}_{\mathbf{p}},$$

$$\Delta\mathbf{s} = -\mathbf{r}_d - \mathbf{A}^\mathsf{T}\Delta\mathbf{y},$$

$$\Delta\mathbf{x} = -\mathbf{x} + \sigma\mu\mathbf{S}^{-1}\mathbf{1}_n - \mathbf{D}^2\Delta\mathbf{s}.$$

Preconditioning in Interior Point Methods

(H. Avron, A. Chowdhuri, P. Drineas, and P. London, NeurIPS 2020)

Standard form of primal LP:

$$\min \mathbf{c}^{\mathsf{T}}\mathbf{x}, \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}, \quad \mathbf{x} \in \mathbb{R}^n$$

$$\mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{b} \in \mathbb{R}^m, \text{ and } \mathbf{c} \in \mathbb{R}^n$$

Path-following, long-step IPMs: compute the Newton search direction; update the current iterate by following a (long) step towards the search direction.

A standard approach involves solving the normal equations:

$$\mathbf{A}\mathbf{D}^2\mathbf{A}^{\mathsf{T}}\Delta\mathbf{y} = \mathbf{p} \quad \text{where} \quad \mathbf{D} \in \mathbb{R}^{n \times n}, \ \mathbf{p} \in \mathbb{R}^m$$

Vector of m unknowns

Use a preconditioned method to solve the above system: we analyzed preconditioned Conjugate Gradient solvers; preconditioned Richardson's; and preconditioned Steepest Descent, all with randomized preconditioners.

Results

(correction vector idea also in O'Neal and Monteiro 2003)

We construct a "correction" vector $v \in R^n$ s.t.:

$$\begin{aligned}
\mathbf{A}\mathbf{D}^2\mathbf{A}^\mathsf{T}\hat{\Delta}\mathbf{y} &= \mathbf{p} + \mathbf{A}\mathbf{S}^{-1}\mathbf{v}\,, \\
\hat{\Delta}\mathbf{s} &= -\mathbf{r}_d - \mathbf{A}^\mathsf{T}\hat{\Delta}\mathbf{y}\,, \\
\hat{\Delta}\mathbf{x} &= -\mathbf{x} + \sigma\mu\mathbf{S}^{-1}\mathbf{1}_n - \mathbf{D}^2\hat{\Delta}\mathbf{s} - \mathbf{S}^{-1}\mathbf{v}
\end{aligned}$$

Then $\mathbf{A}\hat{\Delta}\mathbf{x} = -\mathbf{r}_p$

The vector r_p is the primal residual; the vector r_d is the dual residual. For feasible long-step IPMs, they are both all-zero vectors.

Results

- Our (sketching-based) "correction" vector $v \in R^n$ works with probability $1 - \delta$ and can be constructed in time

$$\mathcal{O}\left(\text{nnz}(\mathbf{A}) \cdot \log(m/\delta) + m^3 \log(m/\delta)\right)$$

- If sketching-based, randomized preconditioned solvers are used, then **we only need mat-vecs to construct v.**
- Using this "correction" vector $v \in R^n$, analyses of long-step (infeasible/ feasible) IPMs work!

Results: feasible, long-step IPMs

If the constraint matrix $A \in R^{m \times n}$ is short-and-fat $(m \ll n)$, then

- Run $O\left(n \cdot log\left(\frac{1}{\epsilon} \right) \right)$ outer iterations of the IPM solver.
- In each outer iteration, the normal equations are solved by $O(\log n)$ inner iterations of the PCG solver.
- Then, the feasible, long-step IPM converges.
- Can be generalized to (exact) low-rank matrices A with rank $k \ll \min\{m, n\}$.

Thus, approximate solutions suffice; ignoring failure probabilities, each inner iteration needs time

$$\mathcal{O}((\text{nnz}(\mathbf{A}) + m^3) \log n)$$

Results: infeasible, long-step IPMs

If the constraint matrix $A \in R^{m \times n}$ is short-and-fat $(m \ll n)$, then

- Run $O\left(n^2 \cdot log\left(\frac{1}{\epsilon}\right)\right)$ outer iterations of the IPM solver.
- In each outer iteration, the normal equations are solved by $O(\log n)$ inner iterations of the PCG solver.
- Then, the infeasible, long-step IPM converges.
- Can be generalized to (exact) low-rank matrices A with rank $k \ll \min\{m, n\}$.

Thus, approximate solutions suffice; ignoring failure probabilities, each inner iteration needs time

$$\mathcal{O}((\text{nnz}(\mathbf{A}) + m^3) \log n)$$

Feasible Predictor-Corrector IPMs

(joint with H. Avron, A. Chowdhuri, P. Drineas ICML 2022; long paper)

- By oscillating between the following two types of steps at each iteration, Predictor-Corrector (PC) IPMs achieve twofold objective of **(i) reducing duality measure μ** and **(ii) improving centrality :**
  - Predictor step ($\sigma = 0$) to reduce the duality measure **μ**.
  - Corrector steps ($\sigma = 1$) to improve centrality.

- PC obtains the best of both worlds: **(i) the practical flexibility of long-step IPMs** and **(ii) the convergence rate of short-step IPMs**.

Feasible Predictor-Corrector IPMs

(joint with H. Avron, A. Chowdhuri, P. Drineas ICML 2022; long paper)

- By oscillating between the following two types of steps at each iteration, Predictor-Corrector (PC) IPMs achieve twofold objective of **(i) reducing duality measure μ** and **(ii) improving centrality :**
  - Predictor step ($\sigma = 0$) to reduce the duality measure **μ**.
  - Corrector steps ($\sigma = 1$) to improve centrality.

- PC obtains the best of both worlds: **(i) the practical flexibility of long-step IPMs** and **(ii) the convergence rate of short-step IPMs**.
- Our work combines the prototypical PC algorithm (e.g., see Wright (1997)) with (preconditioned) inexact solvers.

Predictor-corrector Algorithm Overview

Alternates between predictor and corrector steps

- Predictor step greatly decreases duality measure, while deviating from the central path (centering parameter $\sigma = 1$).
- Corrector step keeps duality measure constant but returns iterate to near central path (centering parameter $\sigma = 0$).
- Alternates between two neighborhoods of the central path $N_2(0.25)$ and $N_2(0.5)$.

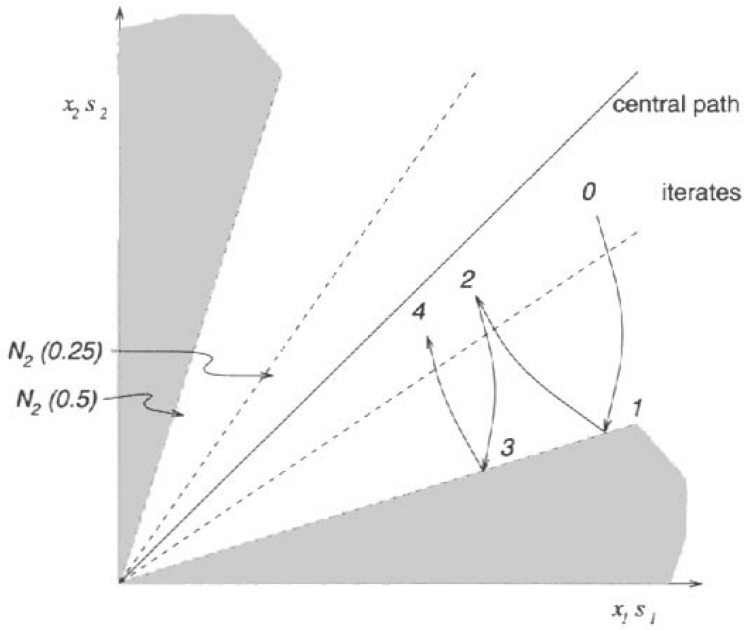

$$\mathcal{N}_2(\theta) = \left\{(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathbb{R}^{2n+m} : \|\mathbf{x} \circ \mathbf{s} - \mu \mathbf{1}_n\|_2 \leq \theta\mu,\ (\mathbf{x}, \mathbf{s}) > 0\right\}.$$

Predictor-Corrector Challenges

Major challenge:

- (Standard analysis breaks; the (feasible) long-step proof was easier; we had to come up with new inequalities for an approximate version of the duality measure.)
- Using a correction vector by itself is insufficient
- Better convergence comes from a more restrictive neighborhood:

$$\mathcal{N}_2(\theta) = \left\{(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathbb{R}^{2n+m} : \|\mathbf{x} \circ \mathbf{s} - \mu \mathbf{1}_n\|_2 \leq \theta\mu, \ (\mathbf{x}, \mathbf{s}) > 0\right\}.$$

Instead of…

$$\mathcal{N}(\gamma) = \left\{(\mathbf{x}, \mathbf{y}, \mathbf{s}) \in \mathcal{F}^0 : \mathbf{x} \circ \mathbf{s} \geq (1-\gamma)\mu \mathbf{1}_n\right\}, \ \gamma \in (0, 1)$$

- Needed sharper inequalities for the inexact steps

Solving the linear system

Again, at each iteration of the Predictor-Corrector IPM, we need to solve the following linear system:

$$\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\Delta\mathbf{y} = \underbrace{-\sigma\mu\mathbf{A}\mathbf{S}^{-1}\mathbf{1}_n + \mathbf{A}\mathbf{x}}_{\mathbf{p}}$$

$$\Delta\mathbf{s} = -\mathbf{A}^\top\Delta\mathbf{y}$$

$$\Delta\mathbf{x} = -\mathbf{x} + \sigma\mu\mathbf{S}^{-1}\mathbf{1}_n - \mathbf{D}^2\Delta\mathbf{s}.$$

Note that the last two equations only involve matrix-vector products. Therefore, we only focus on solving the first equation efficiently.

In the predictor step: ($\sigma = 0$); In the corrector step: ($\sigma = 1$)

Structural Conditions for Inexact PC

- Let $\Delta\tilde{y}$ be an approximate solution to the normal equations $\left(AD^2A^T\right) \cdot \Delta y = p$.
- If $\Delta\tilde{y}$ satisfies (sufficient conditions):

$$\|\Delta\tilde{\mathbf{y}} - \Delta\mathbf{y}\|_{\mathbf{AD^2A}^T} \leq \Theta\left(\frac{\epsilon}{\sqrt{n}\log 1/\epsilon}\right) \quad \|\mathbf{AD}^2\mathbf{A}^T\Delta\tilde{\mathbf{y}} - \mathbf{p}\|_2 \leq \Theta\left(\frac{\epsilon}{\sqrt{n}\log 1/\epsilon}\right)$$

- Then, we prove that the Inexact PC method converges in $O\left(\sqrt{n}\cdot\log\left(\frac{1}{\epsilon}\right)\right)$ iterations, as expected.
- The final solution (and all intermediate iterates) are only approximately feasible.

Structural Conditions for Inexact PC using a correction vector v

- We modified the PC method **<u>using a correction vector v</u>** to make iterates exactly feasible.
- Let $\Delta\tilde{y}$ be an approximate solution to the normal equations $\left(AD^2A^T\right) \cdot \Delta y = p$.
- If $\Delta\tilde{y}$ and v satisfy (sufficient conditions):

$$\mathbf{AS}^{-1}\mathbf{v} = \mathbf{AD}^2\mathbf{A}^T\Delta\tilde{\mathbf{y}} - \mathbf{p}$$

$$\|\mathbf{v}\|_2 < \Theta(\epsilon)$$

- Then, we prove that this modified Inexact PC method converges in $O\left(\sqrt{n} \cdot \log\left(\frac{1}{\epsilon}\right)\right)$ iterations, as expected.
- The final solution (and all intermediate iterates) are exactly feasible.

Satisfying the structural conditions

- We analyzed Preconditioned Conjugate Gradients (PCG) solvers with randomized preconditioners for constraint matrices $A \in R^{n \times n}$ that are: short-and-fat $(m \ll n)$, tall-and-thin $(m \gg n)$ or have exact low-rank $\mathrm{k} \ll \min\{m, n\}$.

- **Satisfying the structural conditions for "standard" Inexact PC:** the PCG solver needs $O\left(\log\left(\frac{n \cdot \sigma_1(AD)}{\epsilon}\right)\right)$ iterations (inner iterations).

- **Satisfying the structural conditions for the "modified" Inexact PC:** the PCG solver needs $O\left(\log\left(\frac{n}{\epsilon}\right)\right)$ iterations (inner iterations).

- Notice that using the error-adjustment vector v in the modified Inexact PC eliminates the dependency on the largest singular value of the matrix AD.

Details: the preconditioned equation

Corresponding preconditioned equation is given by

$$\mathbf{Q}^{-1/2}\,\mathbf{A}\mathbf{D}^2\,\mathbf{A}^\top \underbrace{\mathbf{Q}^{-1/2}\mathbf{z}}_{\Delta\mathbf{y}} = \mathbf{Q}^{-1/2}\mathbf{p}$$

Here $Q \in R^{m\times m}$ is the preconditioner.

Clearly, we need a matrix Q which is "easily" invertible.

(Will come back to this later.)

Satisfying the sufficient conditions for Inexact Predictor-Corrector IPMs (no correction vector)

For an accuracy parameter $\zeta \in (0,1)$, it can be shown that the following two conditions on the preconditioner $Q^{-1/2}$ of the iterative solver can be used to derive the sufficient conditions:

(C1) All singular values $\sigma_i,\ i = 1 \ldots m$ of the preconditioned matrix $Q^{-1/2}\, A\, D$ satisfy:

$$\frac{2}{2+\zeta} \leq \sigma_i^2\left(\mathbf{Q}^{-\frac{1}{2}}\mathbf{AD}\right) \leq \frac{2}{2-\zeta}$$

(C2) As the number of iterations t of the iterative solver increase, the residual norm w.r.t the preconditioned system decreases monotonically:

$$\left\|\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^t - \mathbf{Q}^{-1/2}\mathbf{p}\right\|_2 \leq \zeta^t\left\|\mathbf{Q}^{-1/2}\mathbf{p}\right\|_2$$

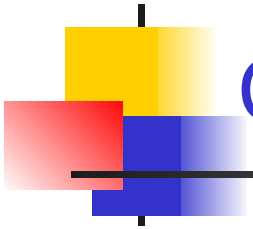

Constructing our preconditioner

- For a suitable sketching matrix $W \in R^{n \times w}$ with $w \ll n$ let $Q = ADWW^T DA^T$.
- To invert Q, it is sufficient to compute the SVD of ADW, which takes $O(m^2 w)$ time.
- **Choice of the sketching matrix W**:
  - W could be the CountSketch matrix with $w = O(m \log m)$ and $\log m$ non-zero entries per row.
  - Many, many other choices exist (random Gaussians, fast randomized transforms, etc.)
  - ADW can be computed in $O(\log m \cdot nnz(A))$ time.
- We can compute $Q^{-1/2}$ in time:

$$\mathcal{O}\left(\mathrm{nnz}(\mathbf{A}) \cdot \log m + m^3 \log m\right)$$

Iterative solver: summary

Input: $\mathbf{AD} \in \mathbb{R}^{m \times n}$ with $m \ll n, \mathbf{p} \in \mathbb{R}^m$, sketching matrix $\mathbf{W} \in \mathbb{R}^{n \times w}$, iteration count t;

Step 1. Compute $\mathbf{ADW}$ and its SVD. Let $\mathbf{U_Q} \in \mathbb{R}^{m \times m}$ be the matrix of its left singular vectors and let $\mathbf{\Sigma_Q^{1/2}} \in \mathbb{R}^{m \times m}$ be the matrix of its singular values;

Step 2. Compute $\mathbf{Q}^{-1/2} = \mathbf{U_Q \Sigma_Q^{-1/2} U_Q^\top}$;

Step 3. Initialize $\tilde{\mathbf{z}}^0 \leftarrow \mathbf{0}_m$ and run standard CG on $\mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^T\mathbf{Q}^{-1/2}\tilde{\mathbf{z}} = \mathbf{Q}^{-1/2}\mathbf{p}$ for t iterations;

Output: return $\Delta\tilde{\mathbf{y}} = \mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^t$.

- Approximate solution $\Delta\tilde{\mathrm{y}}$ can be found by pre-multiplying the solution by the preconditioner.
- Instead of Conjugate Gradients (CG), one can use other iterative solvers, namely, Chebyshev iteration, Steepest descent etc.

Satisfying condition **C1**: Bounding the condition number of the preconditioned matrix

Let $\mathbf{Q} = \mathbf{ADWW}^\top\mathbf{DA}^\top$ and if the sketching matrix $\mathbf{W}$ satisfies $\left\|\mathbf{V}^\top\mathbf{WW}^\top\mathbf{V} - \mathbf{I}_m\right\|_2 \leq \zeta/2$, then, for all $i = 1 \ldots m$

$$(1+\zeta/2)^{-1} \leq \sigma_i^2\left(\mathbf{Q}^{-1/2}\mathbf{AD}\right) \leq (1-\zeta/2)^{-1}$$

- Here V is the matrix of the right singular vectors of A (thin SVD, containing only the singular vectors corresponding to non-zero singular values.
- This is the so-called ℓ_2-subspace embedding condition and implies that the condition number of $Q^{-1/2}AD$ remains small.
- Our Wsatisfies the ℓ_2-subspace embedding condition with high probability.

Satisfying condition **C2**: the residual norm w.r.t the preconditioned system decreases monotonically

Given our preconditioner $\mathbf{Q}^{-1/2} = \left(\mathbf{ADWW}^\top\mathbf{DA}^\top\right)^{-1/2}$ satisfying condition **(C1)** for an accuracy parameter $\zeta \in (0,1)$ and all $\mathbf{i} = 1, 2 \ldots \mathbf{m}$, our iterative solver satisfies

$$\left\|\mathbf{Q}^{-1/2}\mathbf{AD}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^t - \mathbf{Q}^{-1/2}\mathbf{p}\right\|_2 \leq \zeta^t \left\|\mathbf{Q}^{-1/2}\mathbf{p}\right\|_2$$

Here $\tilde{\mathbf{z}}^t$ is the approximate solution returned by the CG iterative solver after t iterations.

- Residual drops exponentially fast as the number of iterations t increases.
- The above guarantee holds for various iterative solvers including CG, Chebyshev iteration, Steepest descent etc.

Satisfying condition **C2** using conjugate gradient

Result (Theorem 8 of Bouyouli et al. (2009)):

Let $\tilde{\mathbf{f}}^{(j)} = \mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}^2\mathbf{A}^\top\mathbf{Q}^{-1/2}\tilde{\mathbf{z}}^j - \mathbf{Q}^{-1/2}\mathbf{p}$ be the residual of by the CG solver at steps j. Then,

$$\|\tilde{\mathbf{f}}^{(j)}\|_2 \leq \frac{\kappa^2(\mathbf{Q}^{-1/2}\mathbf{A}\mathbf{D}) - 1}{2}\|\tilde{\mathbf{f}}^{(j-1)}\|_2$$

- Note that, in general, an energy norm error on the approximate solution derived via CG does not ensure that the residual norms decrease monotonically (even if the energy norm error decreases monotonically).
- From **(C1),** we already have a bound on the condition number of $Q^{-1/2}AD$.
- If we combine the above inequality with the recursion, we get **(C2)**.
- Therefore, our preconditioner ensures the CG residual decreases monotonically.

Satisfying condition **C2** using Chebyshev iteration

Result (Theorem 1.6.2 of Gutknecht (2008)):

The residual norm reduction of the Chebyshev iteration, when applied to an SPD system whose condition number is upper bounded by $\mathcal{U}$, is bounded according to

$$\frac{\|\tilde{\mathbf{f}}^{(t)}\|_2}{\|\tilde{\mathbf{f}}^{(0)}\|_2} \leq 2\left[\left(\frac{\sqrt{\mathcal{U}}+1}{\sqrt{\mathcal{U}}-1}\right)^t + \left(\frac{\sqrt{\mathcal{U}}-1}{\sqrt{\mathcal{U}}+1}\right)^t\right]^{-1}$$

- Chebyshev iteration avoids the computation of the communication intensive inner products which is typically needed for CG or other non-stationary methods.
- Therefore, this solver is convenient in parallel or distributed settings.
- Due to **(C1)**, we already have a bound for U. Using this, we establish **(C2)**.

Other solvers

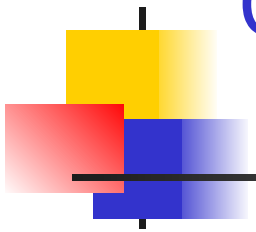

- Similarly, our preconditioner also satisfies **(C2)** with respect to other two popular iterative solvers, namely Steepest descent and Richardson iteration.

- The proofs for both the solvers rely on the fact that due to the efficient preconditioning the residuals of the preconditioned system decrease monotonically.

Constructing the vector v

- Any iterative solver solves the system approximately. Therefore, due to the approximation error caused by the solver, the iterates of our predictor-corrector IPM lose feasibility right after the first iteration.
- As already discussed, for our inexact corrected predictor-corrector, we introduce a correction vector v in order to maintain feasibility at each iteration of the IPM.
- v must satisfy the following invariant at each iteration:

$$\mathbf{AS}^{-1}\mathbf{v} = \mathbf{AD}^{2}\mathbf{A}^{\mathbf{T}}\Delta\tilde{\mathbf{y}} - \mathbf{p}$$

Recall that $\Delta\tilde{y}$ is the solution returned by the iterative solver.

(A solution originally proposed by **Monteiro & O'Neal (2003)** is expensive)

Constructing the vector v

- **Our solution:**

$$\mathbf{v} = (\mathbf{XS})^{1/2}\mathbf{W}(\mathbf{ADW})^{\dagger}\left(\mathbf{AD}^2\mathbf{A}^{\top}\Delta\tilde{\mathbf{y}} - \mathbf{p}\right)$$

- Inspired by work on sketching for **under-constrained regularized regression problems**.
- We use the same sketching matrix **W** that we used for constructing our preconditioner.
- Due to the "good" preconditioner we used, we can show that the norm of v is nicely bounded and thus the sufficient conditions are satisfied.
- Other constructions might be possible and perhaps better in theory and/or practice.

Time to compute the correction vector

- **Recall our solution:**

$$\mathbf{v} = (\mathbf{XS})^{1/2}\mathbf{W}(\mathbf{ADW})^{\dagger}\left(\mathbf{AD}^2\mathbf{A}^{\top}\Delta\tilde{\mathbf{y}} - \mathbf{p}\right)$$

- We have already computed the pseudoinverse of ADW when constructing our preconditioner.
- Pre-multiplying by W takes $O\left(nnz(A)\cdot \log m\right)$ time, assuming $nnz(A) \geq n$.
- $X,\ S$ are diagonal matrices.
- Therefore, computing v takes $O\left(nnz(A)\cdot \log m\right)$ time.

Overall running time (per iteration)

Accounting for the number of iterations of the solver, as well as the failure probability $\eta \in (0,1)$, the per-iteration cost of our approaches is given by:

- **Without** a correction vector:

$$\mathcal{O}\left(\mathrm{nnz}(\mathbf{A}) \cdot \log m/\eta + m^3 \log m/\eta + m \log \frac{\sigma_{\max}(\mathbf{AD})n\mu}{\delta} + \mathrm{nnz}(\mathbf{A}) \cdot \log \frac{\sigma_{\max}(\mathbf{AD})n\mu}{\delta}\right)$$

- **With** a correction vector:

$$\mathcal{O}\left(\mathrm{nnz}(\mathbf{A}) \cdot \log m/\eta + m^3 \log m/\eta + m \log \frac{n\mu}{\delta} + \mathrm{nnz}(\mathbf{A}) \cdot \log \frac{n\mu}{\delta}\right)$$

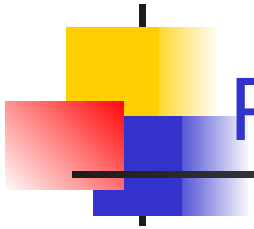

Recap

- Infeasible/Feasible inexact long step method using a correction vector maintains prior outer iteration complexity
- Structural conditions for inexact feasible predictor-corrector methods to maintain prior outer iteration complexity
  - Approximately feasible solution using standard predictor-correct method with inexact solver
  - Exactly feasible solution when using a correction vector
- Fast iterative linear solver for matrices with low exact rank by using RandNLA
  - Fulfills necessary conditions for above IPM convergence analysis
  - Takes advantage of sparsity in constrain matrix
  - Efficient and **practical** large-scale linear programming

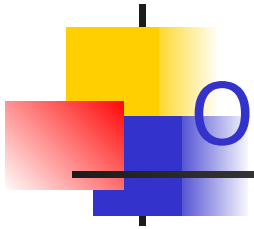

Open problems

- Can we prove similar results for infeasible predictor-corrector IPMs? Recall that such methods need $O(n)$ outer iterations (Yang & Namashita 2018).
- Are our structural conditions necessary? Can we derive simpler conditions? Is a lower precision solver sufficient?
- Could our structural conditions change from one iteration to the next? Could we use dynamic preconditioning or reuse preconditioners from one iteration to the next (e.g., low-rank updates of the preconditioners)?
- Will a similar approach work for more general optimization problems e.g., Quadratic Programming (QP) or Semidefinite Programming (SDP)

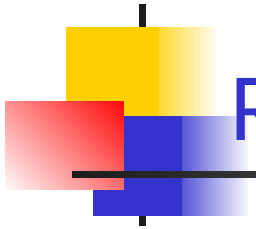

Relevant literature

G. Dexter, A. Chowdhuri, H. Avron, and P. Drineas, On the convergence of Inexact Predictor-Corrector Methods for Linear Programming, ICML 2022.

A. Chowdhuri, G. Dexter, P. London, H. Avron, and P. Drineas, Faster Randomized Interior Point Methods for Tall/Wide Linear Programs, Under Review by JMLR, 2022.

A. Chowdhuri, P. London, H. Avron, and P. Drineas, Speeding up Linear Programming using Randomized Linear Algebra, NeurIPS 2020.

R. Monteiro and J. O'Neal, Convergence analysis of a long-step primaldual infeasible interior-point LP algorithm based on iterative linear solvers, 2003.

D. Woodruff, Sketching as a Tool for Numerical Linear Algebra, FTTCS 2014.

M. W. Mahoney and P. Drineas, RandNLA: Randomized Numerical Linear Algebra, CACM 2016.

P. Drineas and M. W. Mahoney, Lectures on Randomized Numerical Linear Algebra, Amer. Math. Soc., 2018.