# Implementing Randomized Matrix Algorithms in Parallel and Distributed Environments 

Michael W. Mahoney

Stanford University<br>( For more info, see:<br>http:// cs.stanford.edu/people/mmahoney/<br>or Google on "Michael Mahoney")

July 2012

## Outline

(1) Randomized matrix algorithms and large-scale environments
(2) Solving $\ell_{2}$ regression using MPI

- Preconditioning
- Iteratively solving
(3) Solving $\ell_{1}$ regression on MapReduce
- Problem formulation and preconditioning
- Computing a coarse solution
- Computing a fine solution


## Motivation: very large-scale "vector space analytics"

Small-scale and medium-scale:

- Model data by graphs and matrices
- Compute eigenvectors, correlations, etc. in RAM

Very large-scale:

- Model data with flat tables and the relational model
- Compute with join/select and other "counting" in, e.g., Hadoop

Can we "bridge the gap" and do "vector space computations" at very large scale?

- Not obviously yes: exactly computing eigenvectors, correlations, etc. is subtle and uses lots of comminication.
- Not obviously no: lesson from random sampling algorithms is you can get $\epsilon$-approximation of optimal with very few samples.


## Over-determined/over-constrained regression problems

An $\ell_{p}$ regression problem is specified by a design matrix $A \in \mathbb{R}^{m \times n}$, a response vector $b \in \mathbb{R}^{m}$, and a norm $\|\cdot\|_{p}$ :

$$
\operatorname{minimize}_{x \in \mathbb{R}^{n}}\|A x-b\|_{p}
$$

Assume $m \gg n$, i.e., many more "constraints" than "variables." Given an $\epsilon>0$, find a $(1+\epsilon)$-approximate solution $\hat{x}$ in relative scale, i.e.,

$$
\|A \hat{x}-b\|_{p} \leq(1+\epsilon)\left\|A x^{*}-b\right\|_{p}
$$

where $x^{*}$ is a/the optimal solution.

- $p=2$ : Least Squares Approximation: Very widely-used, but highly non-robust to outliers.
- $p=1$ : Least Absolute Deviations: Improved robustness, but at the cost of increased complexity.


## Strongly rectangular data

Some examples:

| SNP | m | n |
| :---: | :---: | :---: |
| number of SNPs $\left(10^{6}\right)$ <br> Tinylmages <br> PDE <br> sensor network <br> NLP | number of pixels in each image $\left(10^{3}\right)$ <br> number of degrees of freedom <br> size of sensing data <br> number of words and $n$-grams | number of subjects $\left(10^{3}\right)$ <br> number of images $\left(10^{8}\right)$ <br> number of time steps <br> number of sensors |
| number of principle components |  |  |

More generally:

- Over-constrained $\ell_{1} / \ell_{2}$ regression is good model for implementing other regression algorithms in large-scale settings.
- Best advances for low-rank matrix problems come by considering the underlying regression problem.


## Traditional algorithms

- for $\ell_{2}$ regression:
- direct methods: QR, SVD, and normal equation ( $O\left(m n^{2}+n^{2}\right)$ time $)$
$\star$ Pros: high precision \& implemented in LAPACK
$\star$ Cons: hard to take advantage of sparsity \& hard to implement in parallel environments
- iterative methods: CGLS, LSQR, etc.
$\star$ Pros: low cost per iteration, easy to implement in some parallel environments, \& capable of computing approximate solutions
$\star$ Cons: hard to predict the number of iterations needed
- for $\ell_{1}$ regression:
- linear programming
- interior-point methods (or simplex, ellipsoid? methods)
- re-weighted least squares
- first-order methods

Nearly all traditional algorithms for low-rank matrix problems, continuous optimization problems, etc. boil down to variants of these methods.

## Why randomized matrix algorithms?

Traditional algorithms are designed to work in RAM and their performance is measured in floating-point operations per second (FLOPS).

- Traditional algorithms are NOT well-suited for:
- problems that are very large
- distributed or parallel computation
- when communication is a bottleneck
- when the data must be accessed via "passes"
- Randomized matrix algorithms are:
- faster: better theory
- simpler: easier to implement
- inherently parallel: exploiting modern computer architectures
- more scalable: modern massive data sets

Big success story in high precision scientific computing applications!
Can they really be implemented in parallel and distributed environments?

## Parallel environments and how they scale

- Shared memory
- cores: $\left[10,10^{3}\right]^{*}$
- memory: [100GB, 100TB]
- Message passing
- cores: $\left[200,10^{5}\right]^{\dagger}$
- memory: [1TB, 1000TB]
- CUDA cores: $\left[5 \times 10^{4}, 3 \times 10^{6}\right]^{\ddagger}$
- GPU memory: [500GB, 20TB]
- MapReduce
- cores: $\left[40,10^{5}\right]^{8}$
- memory: [240GB, 100TB]
- storage: $[100 \mathrm{~TB}, 100 \mathrm{~PB}]^{\top}$
- Distributed computing
- cores: $\left[-, 3 \times 10^{5}\right]^{11}$.

```
\({ }^{*}\) http://www.sgi.com/pdfs/4358.pdf
\(\dagger_{\text {http: //www.top500.org/list/2011/11/100 }}\)
\({ }^{+}\)ht
op500.org/site/50310
\(\S_{\text {http: //www.cloudera.com/blog/2010/04/pushing-the-limits-of-distributed-processing/ }}\)
http://hortonworks.com/blog/an-introduction-to-hdfs-federation/
\(\|_{\text {http://fah-web.stanford.edu/cgi-bin/main.py?qtype=osstats }}\)
```


## Two important notions: leverage and condition

(Mahoney, "Randomized Algorithms for Matrices and Data," FnTML, 2011.)

- Statistical leverage. (Think: eigenvectors. Important for low-precision.)
- The statistical leverage scores of $A$ (assume $m \gg n$ ) are the diagonal elements of the projection matrix onto the column span of $A$.
- They equal the $\ell_{2}$-norm-squared of any orthogonal basis spanning $A$.
- They measure:
» how well-correlated the singular vectors are with the canonical basis
* which constraints have largest "influence" on the LS fit
« a notion of "coherence" or "outlierness"
- Computing them exactly is as hard as solving the LS problem.
- Condition number. (Think: eigenvalues. Important for high-precision.)
- The $\ell_{2}$-norm condition number of $A$ is $\kappa(A)=\sigma_{\max }(A) / \sigma_{\min }^{+}(A)$.
- $\kappa(A)$ bounds the number of iterations; for ill-conditioned problems (e.g., $\kappa(A) \approx 10^{6} \gg 1$ ), the convergence speed is very slow.
- Computing $\kappa(A)$ is generally as hard as solving the LS problem.

These are for the $\ell_{2}$-norm. Generalizations exist for the $\ell_{1}$-norm.

## Meta-algorithm for $\ell_{2}$-norm regression (1 of 2)

(Drineas, Mahoney, etc., 2006, 2008, etc., starting with SODA 2006; Mahoney FnTML, 2011.)

1: Using the $\ell_{2}$ statistical leverage scores of $A$, construct an importance sampling distribution $\left\{p_{i}\right\}_{i=1}^{m}$.
2: Randomly sample a small number of constraints according to $\left\{p_{i}\right\}_{i=1}^{m}$ to construct a subproblem.
3: Solve the $\ell_{2}$-regression problem on the subproblem.

A naïve version of this meta-algorithm gives a $1+\epsilon$ relative-error approximation in roughly $O\left(m n^{2} / \epsilon\right)$ time (DMM 2006, 2008). (Ugh.)

## Meta-algorithm for $\ell_{2}$-norm regression (2 of 2 )

(Drineas, Mahoney, etc., 2006, 2008, etc., starting with SODA 2006; Mahoney FnTML, 2011. ${ }^{\ddagger \ddagger}$ )
A naïve version of this meta-algorithm gives a $1+\epsilon$ relative-error approximation in roughly $O\left(m n^{2} / \epsilon\right)$ time (DMM 2006, 2008). (Ugh.)

## But, we can make this meta-algorithm "fast" in RAM:**

- This meta-algorithm runs in $O(m n \log n / \epsilon)$ time in RAM if:
- we perform a Hadamard-based random random projection and sample uniformly sampling in the randomly rotated basis, or
- we quickly computing approximations to the statistical leverage scores and using those as an importance sampling distribution.
And, we can make this meta-algorithm "high precision" in RAM:
- This meta-algorithm runs in $O(m n \log n \log (1 / \epsilon))$ time in RAM if:
- we use the random projection/sampling basis to construct a preconditioner and couple with a traditional iterative algorithm.

[^0]
## Outline

(1) Randomized matrix algorithms and large-scale environments
(2) Solving $\ell_{2}$ regression using MPI

- Preconditioning
- Iteratively solving
(3) Solving $\ell_{1}$ regression on MapReduce
- Problem formulation and preconditioning
- Computing a coarse solution
- Computing a fine solution


## Algorithm LSRN (for strongly over-determined systems)

1: Choose an oversampling factor $\gamma>1$, e.g., $\gamma=2$. Set $s=\lceil\gamma n\rceil$.
2: Generate $G=\operatorname{randn}(s, m)$, a Gaussian matrix.
3: Compute $\tilde{A}=G A$.
4: Compute $\tilde{A}$ 's economy-sized SVD: $\tilde{U} \tilde{\Sigma} \tilde{V}^{T}$.
5: Let $N=\tilde{V} \tilde{\Sigma}^{-1}$.
6: Iteratively compute the min-length solution $\hat{y}$ to

$$
\operatorname{minimize}_{y \in \mathbb{R}^{r}} \quad\|A N y-b\|_{2} .
$$

7: Return $\hat{x}=N \hat{y}$.

## Why we choose Gaussian random projection

(Meng, Saunders, and Mahoney 2011)

Gaussian random projection

- has the best theoretical result on conditioning,
- can be generated super fast,
- uses level 3 BLAS on dense matrices,
- speeds up automatically on sparse matrices and fast operators,
- still works (with an extra "allreduce" operation) when $A$ is partitioned along its bigger dimension.

So, although it is "slow" (compared with "fast" Hadamard-based projections i.t.o. FLOPS), it allows for better communication properties.

## Theoretical properties of LSRN

(Meng, Saunders, and Mahoney 2011)

- In exact arithmetic, $\hat{x}=x^{*}$ almost surely.
- The distribution of the spectrum of $A N$ is the same as that of the pseudoinverse of a Gaussian matrix of size $s \times r$.
- $\kappa(A N)$ is independent of all the entries of $A$ and hence $\kappa(A)$.
- For any $\alpha \in(0,1-\sqrt{r / s})$, we have

$$
\mathcal{P}\left(\kappa(A N) \leq \frac{1+\alpha+\sqrt{r / s}}{1-\alpha-\sqrt{r / s}}\right) \geq 1-2 e^{-\alpha^{2} s / 2}
$$

where $r$ is the rank of $A$.
So, if we choose $s=2 n \geq 2 r$, we have $\kappa(A N)<6$ w.h.p., and hence we only need around 100 iterations to reach machine precision.

## Implementation of LSRN

- Shared memory (C++ with MATLAB interface)
- Multi-threaded ziggurat random number generator (Marsaglia and Tsang 2000), generating $10^{9}$ numbers in less than 2 seconds using 12 CPU cores.
- A naïve implementation of multi-threaded dense-sparse matrix multiplications.
- Message passing (Python)
- Single-threaded BLAS for matrix-matrix and matrix-vector products.
- Multi-threaded BLAS/LAPACK for SVD.
- Using the Chebyshev semi-iterative method (Golub and Varga 1961) instead of LSQR.


## Solving real-world problems

| matrix | $m$ | $n$ | nnz | rank | cond | DGELSD | $A \backslash b$ | Blendenpik | LSRN |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| landmark | 71952 | 2704 | 1.15 e 6 | 2671 | 1.0 e 8 | 29.54 | $0.6498^{*}$ | - | 17.55 |
| rail4284 | 4284 | 1.1 e 6 | 1.1 e 7 | full | 400.0 | $>3600$ | $1.203^{*}$ | OOM | 136.0 |
| tnimg_1 | 951 | 1 e 6 | 2.1 e 7 | 925 | - | 630.6 | $1067^{*}$ | - | 36.02 |
| tnimg_2 | 1000 | 2 e 6 | 4.2 e 7 | 981 | - | 1291 | $>3600^{*}$ | - | 72.05 |
| tnimg_3 | 1018 | 3 e 6 | 6.3 e 7 | 1016 | - | 2084 | $>3600^{*}$ | - | 111.1 |
| tnimg_4 | 1019 | 4 e 6 | 8.4 e 7 | 1018 | - | 2945 | $>3600^{*}$ | - | 147.1 |
| tnimg_5 | 1023 | 5 e 6 | 1.1 e 8 | full | - | $>3600$ | $>3600^{*}$ | OOM | 188.5 |

Table: Real-world problems and corresponding running times. DGELSD doesn't take advantage of sparsity. Though MATLAB's backslash may not give the min-length solutions to rank-deficient or under-determined problems, we still report its running times. Blendenpik either doesn't apply to rank-deficient problems or runs out of memory (OOM). LSRN's running time is mainly determined by the problem size and the sparsity.

## LSQR (Paige and Saunders 1982)

Code snippet (Python):

```
u = A.matvec(v) - alpha*u
beta = sqrt(comm.allreduce(np.dot(u,u)))
v = comm.allreduce(A.rmatvec(u)) - beta*v
```

Cost per iteration:

- two matrix-vector multiplications
- two cluster-wide synchronizations


## Chebyshev semi-iterative (CS) method (Golub and Varga 1961)

The strong concentration results on $\sigma^{\max }(A N)$ and $\sigma^{\min }(A N)$ enable use of the CS method, which requires an accurate bound on the extreme singular values to work efficiently.

Code snippet (Python):

```
v = comm.allreduce(A.rmatvec(r)) - beta*v
x += alpha*v
r -= alpha*A.matvec(v)
```

Cost per iteration:

- two matrix-vector multiplications
- one cluster-wide synchronization


## LSQR vs. CS on an Amazon EC2 cluster

(Meng, Saunders, and Mahoney 2011)

| solver | $N_{\text {nodes }}$ | $N_{\text {processes }}$ | $m$ | $n$ | nnz | $N_{\text {iter }}$ | $T_{\text {iter }}$ | $T_{\text {total }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| LSRN w/ CS |  |  | 1024 |  | 8.4e7 | 106 | 34.03 | 170.4 |
| LSRN w/ LSQR | 2 | 4 | 1024 | 4 e 6 | 8.4 e 7 | 84 | 41.14 | 178.6 |
| LSRN w/ CS |  |  |  |  |  | 106 | 50.37 | 193.3 |
| LSRN w/ LSQR | 5 | 10 | 1024 | 1 e 7 | 2.1 e 8 | 84 | 68.72 | 211.6 |
| LSRN w/ CS | 10 | 2 | 1024 |  |  | 106 | 73.73 | 220.9 |
| LSRN w/ LSQR | 10 | 20 | 1024 | 2 e 7 | 4.2 e 8 | 84 | 102.3 | 249.0 |
| LSRN w/ CS |  |  |  |  |  | 106 | 102.5 | 255.6 |
| LSRN w/ LSQR | 20 | 40 | 1024 | 4 e 7 | 8.4 e 8 | 84 | 137.2 | 290.2 |

Table: Test problems on an Amazon EC2 cluster and corresponding running times in seconds. Though the CS method takes more iterations, it actually runs faster than LSQR by making only one cluster-wide synchronization per iteration.

## Outline

## (1) Randomized matrix algorithms and large-scale environments

(2) Solving $\ell_{2}$ regression using MPI

- Preconditioning
- Iteratively solving
(3) Solving $\ell_{1}$ regression on MapReduce
- Problem formulation and preconditioning
- Computing a coarse solution
- Computing a fine solution


## Problem formulation

We use an equivalent formulation of $\ell_{1}$ regression, which consists of a homogeneous objective function and an affine constraint:

$$
\begin{array}{ll}
\operatorname{minimize}_{x \in \mathbb{R}^{n}} & \|A x\|_{1} \\
\text { subject to } & c^{T} x=1
\end{array}
$$

Assume that $A \in \mathbb{R}^{m \times n}$ has full column rank, $m \gg n$, and $c \neq 0$.

Condition number, well-conditioned bases, and leverage scores for the $\ell_{1}$-norm

- A matrix $U \in \mathbb{R}^{m \times n}$ is $(\alpha, \beta, p=1)$-conditioned if $|U|_{1} \leq \alpha$ and $\|x\|_{\infty} \leq \beta\|U x\|_{1}, \forall x$; and $\ell_{1}$-well-conditioned if $\alpha, \beta=\operatorname{poly}(n)$.
- Define the $\ell_{1}$ leverage scores of an $m \times n$ matrix $A$, with $m>n$, to be the $\ell_{1}$-norms-squared of the rows of an $\ell_{1}$-well-conditioned basis of A. (Only well-defined up to poly $(n)$ factors.)
- Define the $\ell_{1}$-norm condition number of $A$, denoted by $\kappa_{1}(A)$, as:

$$
\kappa_{1}(A)=\frac{\sigma_{1}^{\max }(A)}{\sigma_{1}^{\min }(A)}=\frac{\max _{\|x\|_{2}=1}\|A x\|_{1}}{\min _{\|x\|_{2}=1}\|A x\|_{1}}
$$

This implies: $\sigma_{1}^{\min }(A)\|x\|_{2} \leq\|A x\|_{1} \leq \sigma_{1}^{\max }(A)\|x\|_{2}, \quad \forall x \in \mathbb{R}^{n}$.

## Meta-algorithm for $\ell_{1}$-norm regression

1: Using an $\ell_{1}$-well-conditioned basis for $A$, construct an importance sampling distribution $\left\{p_{i}\right\}_{i=1}^{m}$ from the $\ell_{1}$-leverage scores.
2: Randomly sample a small number of constraints according to $\left\{p_{i}\right\}_{i=1}^{m}$ to construct a subproblem.
3: Solve the $\ell_{1}$-regression problem on the subproblem.

A naïve version of this meta-algorithm gives a $1+\epsilon$ relative-error approximation in roughly $O\left(m n^{5} / \epsilon^{2}\right)$ time (DDHKM 2009). (Ugh.)

But, as with $\ell_{2}$ regression:

- We can make this algorithm run much faster in RAM by
- approximating the $\ell_{1}$-leverage scores quickly, or
- performing an " $\ell_{1}$ projection" to uniformize them approximately.
- We can make this algorithm work at higher precision in RAM at large-scale by coupling with an iterative algorithm.


## Subspace-preserving sampling

(Dasgupta, et al. 2009)

Theorem (Dasgupta, Drineas, Harb, Kumar, and Mahoney 2009)
Given $A \in \mathbb{R}^{m \times n}$ and $\epsilon<\frac{1}{7}$, let $s \geq 64 n^{1 / 2} \kappa_{1}(A)\left(n \ln \frac{12}{\epsilon}+\ln \frac{2}{\delta}\right) / \epsilon^{2}$. Let $S \in \mathbb{R}^{m \times m}$ be a diagonal "sampling matrix" with random diagonals:

$$
S_{i i}= \begin{cases}\frac{1}{p_{i}} & \text { with probability } p_{i} \\ 0 & \text { otherwise }\end{cases}
$$

where

$$
p_{i} \geq \min \left\{1, \frac{\left\|A_{i *}\right\|_{1}}{|A|_{1}} \cdot s\right\}, \quad i=1, \ldots, m .
$$

Then, with probability at least $1-\delta$, the following holds, for all $x \in \mathbb{R}^{n}$ :

$$
(1-\epsilon)\|A x\|_{1} \leq\|S A x\|_{1} \leq(1+\epsilon)\|A x\|_{1} .
$$

## Computing subsampled solutions

(Dasgupta, et al. 2009)

Let $\hat{x}$ be an optimal solution to the subsampled problem:

$$
\begin{array}{ll}
\operatorname{minimize} & \|S A x\|_{1} \\
\text { subject to } & c^{T} x=1
\end{array}
$$

Then with probability at least $1-\delta$, we have

$$
\|A \hat{x}\|_{1} \leq \frac{1}{1-\epsilon}\|S A \hat{x}\|_{1} \leq \frac{1}{1-\epsilon}\left\|S A x^{*}\right\|_{1} \leq \frac{1+\epsilon}{1-\epsilon}\left\|A x^{*}\right\|_{1} .
$$

- It is hard to follow the theory closely on the sample size.
- We determine the sample size based on hardware capacity, not on $\epsilon$.


## $\ell_{1}$-norm preconditioning via oblivious projections

Find an oblivious (i.e., independent of $A$ ) projection matrix $\Pi \in \mathbb{R}^{\mathcal{O}(n \log n) \times m}$, such that

$$
\|A x\|_{1} \leq\|\Pi A x\|_{1} \leq \kappa_{\Pi}\|A x\|_{1}, \quad \forall x
$$

Compute $R=\operatorname{qr}(\Pi A)$.
Then,

$$
\frac{1}{\kappa_{\Pi}}\|y\|_{2} \leq\left\|A R^{-1} y\right\|_{1} \leq \mathcal{O}\left(n^{1 / 2} \log ^{1 / 2} n\right)\|y\|_{2}, \quad \forall y
$$

Therefore, $A R^{-1}$ is $\ell_{1}$-well-conditioned: $\kappa_{1}\left(A R^{-1}\right)=\mathcal{O}\left(n^{1 / 2} \log ^{1 / 2} n \cdot \kappa_{\Pi}\right)$.

| Constructions for $\Pi$ | time | $\kappa \Pi$ |
| :---: | :---: | :---: |
| Cauchy (Sohlier and Woodruff 2011) | $\mathcal{O}\left(m n^{2} \log n\right)$ | $\mathcal{O}(n \log n)$ |
| Fast Cauchy (Clarkson, Drineas, Magdon-Ismail, | $\mathcal{O}(m n \log n)$ | $\mathcal{O}\left(n^{2} \log ^{2} n\right)$ |

## Evaluation on large-scale $\ell_{1}$ regression problem (1 of 2 ).

|  | $\left\\|x-x^{*}\right\\|_{1} /\left\\|x^{*}\right\\|_{1}$ | $\left\\|x-x^{*}\right\\|_{2} /\left\\|x^{*}\right\\|_{2}$ | $\left\\|x-x^{*}\right\\|_{\infty} /\left\\|x^{*}\right\\|_{\infty}$ |
| :---: | :---: | :---: | :---: |
| CT (Cauchy) | $[0.008,0.0115]$ | $[0.00895,0.0146]$ | $[0.0113,0.0211]$ |
| GT (Gaussian) | $[0.0126,0.0168]$ | $[0.0152,0.0232]$ | $[0.0184,0.0366]$ |
| NOCD | $[0.0823,22.1]$ | $[0.126,70.8]$ | $[0.193,134]$ |
| UNIF | $[0.0572,0.0951]$ | $[0.089,0.166]$ | $[0.129,0.254]$ |

Table: The first and the third quartiles of relative errors in $1-, 2-$, and $\infty$-norms on a data set of size $10^{10} \times 15$. CT clearly performs the best. (FCT performs similarly.) GT follows closely. NOCD generates large errors, while UNIF works but it is about a magnitude worse than CT.

## Evaluation on large-scale $\ell_{1}$ regression problem (2 of 2 ).



Figure: The first (solid) and the third (dashed) quartiles of entry-wise absolute errors on a data set of size $10^{10} \times 15$. CT clearly performs the best. (FCT performs similarly.) GT follows closely. NOCD and UNIF are much worse.

## $\ell_{1}$-norm preconditioning via ellipsoidal rounding

Find an ellipsoid $\mathcal{E}=\left\{x \mid x^{\top} E^{-1} x \leq 1\right\}$ such that

$$
\frac{1}{\kappa_{1}} \mathcal{E} \subseteq \mathcal{C}=\left\{x \mid\|A x\|_{1} \leq 1\right\} \subseteq \mathcal{E}
$$

Then we have

$$
\|y\|_{2} \leq\left\|A E^{1 / 2} y\right\|_{1} \leq \kappa_{1}\|y\|_{2}, \quad \forall y .
$$

|  | time | $\kappa_{1}$ | passes |
| :---: | :---: | :---: | :---: |
| Löwner-John ellipsoid | $($ exists $)$ | $n^{1 / 2}$ |  |
| Clarkson 2005 (Lovász 1986) | $\mathcal{O}\left(m n^{5} \log m\right)$ | $n$ | multiple |
| Meng and Mahoney 2012 | $\mathcal{O}\left(m n^{3} \log m\right)$ | $2 n$ | multiple |
|  | $\mathcal{O}\left(m n^{2} \log \frac{m}{n}\right)$ | $2 n^{2}$ | single |
|  | $\mathcal{O}\left(m n \log \frac{m}{n^{2}}\right)$ | $\mathcal{O}\left(n^{5 / 2} \log ^{1 / 2} n\right)$ | single |

## Fast ellipsoidal rounding

(1) Partition $A$ into sub-matrices $A_{1}, A_{2}, \ldots, A_{M}$ of size $\mathcal{O}\left(n^{3} \log n\right) \times n$.
(2) Compute $\tilde{A}_{i} \in \mathbb{R}^{\mathcal{O}(n \log n) \times n}=\operatorname{FJLT}\left(A_{i}\right)$, for $i=1, \ldots, M$.
(3) Compute an ellipsoid $\mathcal{E}$, which gives a $2 n$-rounding of $\tilde{\mathcal{C}}=\left\{x \mid \sum_{i=1}^{M}\left\|\tilde{A}_{i} x\right\|_{2} \leq 1\right\}$.
$\rightarrow$ By a proper scaling, $\mathcal{E}$ gives an $\mathcal{O}\left(n^{5 / 2} \log ^{1 / 2} n\right)$-rounding of $\mathcal{C}$.

Can use this to get a "one-pass conditioning" algorithm!

## A MapReduce implementation

- Inputs: $A \in \mathbb{R}^{m \times n}$ and $\kappa_{1}$ such that

$$
\|x\|_{2} \leq\|A x\|_{1} \leq \kappa_{1}\|x\|_{2}, \quad \forall x
$$

$c \in \mathbb{R}^{n}$, sample size $s$, and number of subsampled solutions $n_{x}$.

- Mapper:
(1) For each row $a_{i}$ of $A$, let $p_{i}=\min \left\{s\left\|a_{i}\right\|_{1} /\left(\kappa_{1} n^{1 / 2}\right), 1\right\}$.
(2) For $k=1, \ldots, n_{x}$, emit $\left(k, a_{i} / p_{i}\right)$ with probability $p_{i}$.
- Reducer:
(1) Collect row vectors associated with key $k$ and assemble $A_{k}$.
(2) Compute $\hat{x}_{k}=\arg \min _{c^{T} x=1}\left\|A_{k} x\right\|_{1}$ using interior-point methods.
(3) Return $\hat{x}_{k}$.

Note that multiple subsampled solutions can be computed in a single pass.

## Iteratively solving

If we want to have a few more accurate digits from the subsampled solutions, we may consider iterative methods.

|  | passes | extra work per pass |
| :---: | :---: | :---: |
| subgradient (Clarkson 2005) | $\mathcal{O}\left(n^{4} / \epsilon^{2}\right)$ |  |
| gradient (Nesterov 2009) | $\mathcal{O}\left(m^{1 / 2} / \epsilon\right)$ |  |
| ellipsoid (Nemiroski and Yudin 1972) | $\mathcal{O}\left(n^{2} \log \left(\kappa_{1} / \epsilon\right)\right)$ |  |
| inscribed ellipsoids | $\mathcal{O}\left(n \log \left(\kappa_{1} / \epsilon\right)\right)$ | $\mathcal{O}\left(n^{7 / 2} \log n\right)$ |
| (Tarasov, Khachiyan, and Erikh 1988) |  |  |

## The Method of Inscribed Ellipsoids (MIE)

MIE works similarly to the bisection method, but in a higher dimension.
It starts with a search region $\mathcal{S}_{0}=\{x \mid S x \leq t\}$ which contains a ball of desired solutions described by a separation oracle. At step $k$, we first compute the maximum-volume ellipsoid $\mathcal{E}_{k}$ inscribing $\mathcal{S}_{k}$. Let $y_{k}$ be the center of $\mathcal{E}_{k}$. Send $y_{k}$ to the oracle, if $y_{k}$ is not a desired solution, the oracle returns a linear cut that refines the search region $\mathcal{S}_{k} \rightarrow \mathcal{S}_{k+1}$.

Why do we choose MIE?

- Least number of iterations
- Initialization using all the subsampled solutions
- Multiple queries per iteration


## Constructing the initial search region

Given any feasible $\hat{x}$, let $\hat{f}=\|A \hat{x}\|_{1}$ and $\hat{g}=A^{T} \operatorname{sign}(A \hat{x})$. we have

$$
\left\|x^{*}-\hat{x}\right\|_{2} \leq\left\|A\left(x^{*}-\hat{x}\right)\right\|_{1} \leq\left\|A x^{*}\right\|_{1}+\|A \hat{x}\|_{1} \leq 2 \hat{f}
$$

and, by convexity,

$$
\left\|A x^{*}\right\|_{1} \geq\|A \hat{x}\|_{1}+\hat{g}^{T}\left(x^{*}-\hat{x}\right)
$$

which implies $\hat{g}^{T} x^{*} \leq \hat{g}^{T} \hat{x}$.
Hence, for each subsampled solution, we have a hemisphere that contains the optimal solution.

We use all these hemispheres to construct the initial search region $\mathcal{S}_{0}$.

## Computing multiple $f$ and $g$ in a single pass

On MapReduce, the cost of input/output may dominate the cost of the actual computation, which requires us to design algorithms that could do more computations in a single pass.

- A single query:

$$
f(x)=\|A x\|_{1}, \quad g(x)=A^{T} \operatorname{sign}(A x)
$$

- Multiple queries:

$$
F(X)=\operatorname{sum}(|A X|, 0), \quad G(X)=A^{T} \operatorname{sign}(A X)
$$

An example on a 10 -node Hadoop cluster:

- $A: 10^{8} \times 50,118.7 \mathrm{~GB}$.
- A single query: 282 seconds.
- 100 queries in a single pass: 328 seconds.


## MIE with sampling initialization and multiple queries



Figure: Comparing different MIEs on an $\ell_{1}$ regression problem of size $10^{6} \times 20$.

## MIE with sampling initialization and multiple queries



Figure: Comparing different MIEs on an $\ell_{1}$ regression problem of size $5.24 e 9 \times 15$.

## Conclusion

- Implementations of randomized matrix algorithms for $\ell_{p}$ regression in large-scale parallel and distributed environments.
- Includes Least Squares Approximation and Least Absolute Deviations as special cases.
- Scalability comes due to restricted communications.
- Randomized algorithms are inherently communication-avoiding.
- Look beyond FLOPS in large-scale parallel and distributed environments.
- Design algorithms that require more computation than traditional algorithms, but that have better communication profiles.
- On MPI: Chebyshev semi-iterative method vs. LSQR.
- On MapReduce: Method of inscribed ellipsoids with multiple queries.


[^0]:    ** (Sarlós 2006; Drineas, Mahoney, Muthu, Sarlós 2010; Drineas, Magdon-Ismail, Mahoney, Woodruff 2011.)
    $\dagger \dagger$ (Rokhlin \& Tygert 2008; Avron, Maymounkov, \& Toledo 2010; Meng, Saunders, \& Mahoney 2011.)
    $\ddagger \ddagger$ (Mahoney, "Randomized Algorithms for Matrices and Data," FnTML, 2011.)

