CME 323: Distributed Algorithms and Optimization, Spring 2016
http://stanford.edu/~rezab/dao.
Instructor: Reza Zadeh, Matroid and Stanford.
Lecture 17, 5/23/2016. Scribed by Xin Jin and Yi-Chun Chen.
Lecture outline:

- Covariance matrices and all-pairs similarity
- Computing $A^{\top} A$
- General convex optimization, where "general" means gradient doesn't necessarily exist

What can be considered a good distributed algorithm? For PRAM models, a good parallel algorithm usually has the same work as the sequential version, and has small, say logarithmic, depth. However, for distributed algorithms, having communication costs that are independent of some input dimensions is often a more important consideration.

### 17.1 Computing $A^{\top} A$

$A$ is an $m \times n$ matrix:

$$
A=\left(\begin{array}{cccc}
a_{1,1} & a_{1,2} & \cdots & a_{1, n} \\
a_{2,1} & a_{2,2} & \cdots & a_{2, n} \\
\vdots & \vdots & \ddots & \vdots \\
a_{m, 1} & a_{m, 2} & \cdots & a_{m, n}
\end{array}\right)
$$

We assume $A$ is tall and skinny, i.e, $m \gg n$. Examples values would be $m=10^{12}$ and $n \in\left\{10^{4}, 10^{6}\right\}$. Moreover, $A$ has sparse rows, each of which has at most $L$ nonzeros. In general, $A$ is stored across hundreds of machines and cannot be streamed through a single machine. In particular, even two consecutive rows may not be on the same machine.

An example of $A$ in real applications would be a Netflix matrix: A lot of users and only a few movies. Rows are all sparse, but some column can be very dense, e.g., the column for movie Godfather.

Our task is to compute $A^{\top} A$ which is $n \times n$, considerably smaller than $A . A^{\top} A$ is in general dense; each entry is simply a dot product of a pair of columns of $A$.

```
Algorithm 17.1 NaiveMapper \(\left(r_{i}\right)\)
    for all pairs \(\left(a_{i j}, a_{i k}\right)\) in \(r_{i}\) do
        Emit \(\left((j, k) \rightarrow a_{i j} a_{i k}\right)\)
    end for
```

We first describe an easy but naive algorithm: since each entry of the product matrix is of the form $\sum_{i} a_{i k} a_{i j}$, we can emit all entry products between column $k$ and $j$ to one reducer and let

```
Algorithm 17.2 NaiveReducer \(\left((i, j),\left\langle v_{1}, \ldots, v_{R}\right\rangle\right)\)
    output \(c_{i}^{\top} c_{j} \rightarrow \sum_{i=1}^{R} v_{i}\)
```

that reducer do the summation. See $\sqrt{17.1}$ and $\sqrt{17.2}$ for the pseudocodes of the Mapper and the Reducer. Easy to see that the shuffle size is $O\left(m L^{2}\right)$ and the largest reduce-key is $O(m)$. Since $m$ is usually very large (e.g., $10^{12}$ ), this algorithm would not work well. It turns out that we can bring down both complexities via clever sampling.

We need the following notion of similarity of two vectors
Definition 17.1 (Cosine Similarity) The cosine similarity between two columns $c_{i}$ and $c_{j}$ is defined as

$$
\cos (i, j)=\frac{c_{i}^{\top} c_{j}}{\left\|c_{i}\right\|\left\|c_{j}\right\|}
$$

```
Algorithm 17.3 DIMSUMMapper \(\left(r_{i}\right)\)
    for all pairs \(\left(a_{i j}, a_{i k}\right)\) in \(r_{i}\) do
        With probability \(\min \left\{1, \frac{\gamma}{\left\|c_{j}\right\|\left\|c_{k}\right\|}\right\}\), emit \(\left((j, k) \rightarrow a_{i j} a_{i k}\right)\)
    end for
```

```
Algorithm 17.4 DIMSUMReducer \(\left((i, j),\left\langle v_{1} \ldots, v_{R}\right\rangle\right)\)
    if \(\frac{\gamma}{\left\|c_{i}\right\|\left\|c_{j}\right\|}>1\) then
        output \(b_{i j} \rightarrow \frac{1}{\left\|c_{j}\right\|\left\|c_{j}\right\|} \sum_{i=1}^{R} v_{i}\)
    else
        output \(b_{i j} \rightarrow \frac{1}{\gamma} \sum_{i=1}^{R} v_{i}\)
    end if
```

The Dimension Independent Matrix Square using MapReduce (DIMSUM) algorithm is described in 17.3 and 17.4 ). The DIMSUM algorithm outputs the cosine similarities (in fact probabilistic estimates of the cosine similarities). Also note that you need to compute the norms of columns beforehand (which requires all-to-all communication). To see why the outputs are estimates of the similarities, define $\tilde{v}_{i}$ to be $v_{i}$ if $v_{i}$ is emitted in 17.3 and 0 if not, and $R$ the number of nonzero products for $c_{i}$ and $c_{j}\left(v_{i}\right.$ and $R$ are the same as in $(17.1)$ and 17.2 , not as in $(17.3)$ and 17.4$)$. Then the expectation of an output from 17.4 is

$$
E\left[\frac{1}{\gamma} \sum_{i=1}^{R} \tilde{v}_{i}\right]=\frac{1}{\gamma} \mathbb{P}\left(\tilde{v}_{i}=v_{i}\right) \sum_{i=1}^{R} v_{i}=\frac{1}{\left\|c_{i}\right\|\left\|c_{j}\right\|} \sum_{i=1}^{R} v_{i}
$$

We shall prove that the shuffle size is $O(n L \gamma)$ and largest reduce-key is $O(\gamma)$. To simplify the proof, we assume $A_{i j} \in\{0,1\}$, and thus $\left\|c_{i}\right\|_{2}=\sqrt{\#\left(c_{i}\right)}$, where $\#\left(c_{i}\right)$ denotes the number of nonzeros
in $c_{i}$. Also we let $\#\left(c_{i}, c_{j}\right)$ denote the number of co-occurrences of nonzeros in $c_{i}$ and $c_{j}$, i.e., the number of $k$ 's for which $c_{i k} c_{j k} \neq 0$.

Theorem 17.2 For $\{0,1\}$ matrices, the expected shuffle size of DIMSUMMapper is $O(n L \gamma)$.
Proof: The expected contribution from each pair of columns will constitute the shuffle size:

$$
\begin{aligned}
\sum_{i=1}^{n} \sum_{j=i+1}^{n} \sum_{k=1}^{\#\left(c_{i}, c_{j}\right)} \mathbb{P}\left(\text { DIMSUMEmit }\left(c_{i}, c_{j}\right)\right) & =\sum_{i=1}^{n} \sum_{j=i+1}^{n} \#\left(c_{i}, c_{j}\right) \mathbb{P}\left(\text { DIMSUMEmit }\left(c_{i}, c_{j}\right)\right) \\
& \leq \sum_{i=1}^{n} \sum_{j=i+1}^{n} \gamma \frac{\#\left(c_{i}, c_{j}\right)}{\sqrt{\#\left(c_{i}\right)} \sqrt{\#\left(c_{j}\right)}} \\
& \leq \frac{\gamma}{2} \sum_{i=1}^{n} \sum_{j=i+1}^{n} \#\left(c_{i}, c_{j}\right)\left(\frac{1}{\#\left(c_{i}\right)}+\frac{1}{\#\left(c_{j}\right)}\right) \quad \text { (by AM-GM) } \\
& \leq \gamma \sum_{i=1}^{n} \frac{1}{\#\left(c_{i}\right)} \sum_{j=1}^{n} \#\left(c_{i}, c_{j}\right) \\
& \leq \gamma \sum_{i=1}^{n} \frac{1}{\#\left(c_{i}\right)} L \#\left(c_{i}\right) \\
& =\gamma L n .
\end{aligned}
$$

On the other hand, the expected reduce-key size is at most

$$
\min \left\{\#\left(c_{j}\right), \#\left(c_{k}\right)\right\} \times \frac{\gamma}{\left\|c_{j}\right\|\left\|c_{k}\right\|} \leq \gamma
$$

which implies the largest reduce-key is $O(\gamma)$.
The key improvement of the DIMSUM algorithm over the naive one is that the complexities now are independent of the dimension $m$. This is achieved by sampling high magnitude columns with low probabilities. The parameter $\gamma$ is a knob that can be used to preserve similarities and singular values:

- With a low setting of $\gamma$, i.e., $\gamma=\Omega(\log (n) / s)$, preserve similar entires of $A^{\top} A$
- With a high setting of $\gamma$, i.e., $\gamma=\Omega\left(n / \epsilon^{2}\right)$, preserve singular values of $A^{\top} A$

Full proofs of the preservation properties can be found in [1].

### 17.2 Gradient descent

$$
\begin{align*}
F(x) & =\sum_{i=1}^{n} F_{i}(x)  \tag{1}\\
\nabla F(x) & =\sum_{i=1}^{n} \nabla F_{i}(x)
\end{align*}
$$

$\nabla F_{i}(x) \in \mathbb{R}^{d}$, and the sum is computed via Reduce. The communication cost is independent of $n$. For general convex optimization problems, you may not have gradients (requires differentiability of objectives) or even subgradients (requires continuity of objectives).

ADMM Introduce $x_{i}$, for which each machine gets one $x_{i}$. Instead of (1), we solve the following problem

$$
\min _{x_{i}, z} \quad \sum_{i=1}^{P} F_{i}\left(x_{i}\right)
$$

such that $\forall 1 \leq i \leq p, x_{i}-z=0$
The alternative formulation can be numerically solved by the following iterative algorithm:

$$
\begin{aligned}
x_{i}^{k+1} & =\underset{x_{i}}{\operatorname{argmin}}\left\{F_{i}\left(x_{i}\right)+\left(y_{i}^{k}\right)^{\top}\left(x_{i}-\bar{x}^{k}\right)+\frac{p}{2}\left\|x_{i}-\bar{x}^{k}\right\|_{2}^{2}\right\} \\
y^{k+1} & =y_{i}^{k}+P\left(x_{i}^{k+1}-\bar{x}^{k+1}\right) \\
\bar{x}^{k+1} & =\frac{1}{P} \sum_{i=1}^{P} x_{i}^{k}
\end{aligned}
$$

Only requires all-to-one communication of size $d$. For details, see [2]. If you have strongly convex functions, you can do parallel SGD:

1. Shuffle data
2. Solve locally on each machine
3. Average results

## References

[1] Zadeh, Reza Bosagh, and Gunnar Carlsson. "Dimension independent matrix square using mapreduce." arXiv preprint arXiv:1304.1467 (2013).
[2] Boyd, Stephen, et al. "Distributed optimization and statistical learning via the alternating direction method of multipliers." Foundations and Trends in Machine Learning 3.1 (2011): 1-122.

