## All-Pairs-Shortest-Paths in Spark

# Charles Zheng, Jingshu Wang and Arzav Jain 

Stanford University

$$
\text { June 1, } 2015
$$

## Problem

- Weighted graph $G=(V, E)$ with $n$ vertices
- Compute $n \times n$ matrix of distances $S$ where

$$
S_{i j}=\text { weight of shortest path from } i \text { to } j
$$



## Floyd-Warshall: Single Core

$S_{i j}^{k}$ - shortest path distance from $i$ to $j$ using intermediate nodes 1 to $k$

$$
S_{i j}^{k}=\left\{\begin{array}{cc}
w_{i j} & k=0 \\
\min \left(S_{i j}^{k-1}, S_{i k}^{k-1}+S_{k j}^{k-1}\right) & k>0
\end{array}\right.
$$

$$
S \leftarrow \min (S, S(:, k) \otimes S(k,:))
$$

## Floyd-Warshall

## Initial input



## Floyd-Warshall

## Iteration 1

$$
\left(\begin{array}{llllll}
0 & 1 & & & & 1 \\
1 & 0 & 1 & & & 2 \\
& 1 & 0 & 1 & & \\
& & 1 & 0 & 1 & \\
& & & 1 & 0 & 1 \\
1 & 2 & & & 1 & 0
\end{array}\right)
$$



## Floyd-Warshall

## Iteration 2

$$
\left(\begin{array}{llllll}
0 & 1 & 2 & & & 1 \\
1 & 0 & 1 & & & 2 \\
2 & 1 & 0 & 1 & & 3 \\
& & 1 & 0 & 1 & \\
& & & 1 & 0 & 1 \\
1 & 2 & 3 & & 1 & 0
\end{array}\right)
$$



## Floyd-Warshall

## Iteration 3

$$
\left(\begin{array}{llllll}
0 & 1 & 2 & 3 & & 1 \\
1 & 0 & 1 & 2 & & 2 \\
2 & 1 & 0 & 1 & & 3 \\
3 & 2 & 1 & 0 & 1 & 4 \\
& & & 1 & 0 & 1 \\
1 & 2 & 3 & 4 & 1 & 0
\end{array}\right)
$$



## Floyd-Warshall

Iteration 4

$$
\left(\begin{array}{llllll}
0 & 1 & 2 & 3 & 4 & 1 \\
1 & 0 & 1 & 2 & 3 & 2 \\
2 & 1 & 0 & 1 & 2 & 3 \\
3 & 2 & 1 & 0 & 1 & 4 \\
4 & 3 & 2 & 1 & 0 & 1 \\
1 & 2 & 3 & 4 & 1 & 0
\end{array}\right)
$$



## Floyd-Warshall

Iteration 5

$$
\left(\begin{array}{llllll}
0 & 1 & 2 & 3 & 4 & 1 \\
1 & 0 & 1 & 2 & 3 & 2 \\
2 & 1 & 0 & 1 & 2 & 3 \\
3 & 2 & 1 & 0 & 1 & 2 \\
4 & 3 & 2 & 1 & 0 & 1 \\
1 & 2 & 3 & 2 & 1 & 0
\end{array}\right)
$$



## Floyd-Warshall

Iteration 6, (terminate)

$$
\left(\begin{array}{llllll}
0 & 1 & 2 & 3 & 2 & 1 \\
1 & 0 & 1 & 2 & 3 & 2 \\
2 & 1 & 0 & 1 & 2 & 3 \\
3 & 2 & 1 & 0 & 1 & 2 \\
2 & 3 & 2 & 1 & 0 & 1 \\
1 & 2 & 3 & 2 & 1 & 0
\end{array}\right)
$$



## Floyd-Warshall

- Cost: $O\left(n^{3}\right)$ operations (single-core)
- Takes $n$ sequential iterations


## Problems with Floyd-Warshall

- FW updates by considering 1 new vertex at a time
- Result: $n$ iterations
- High \# iterations = latency in distributed setting
- Solomonik et al. (2013) show how to "block" FW iterates
- We modify their block-based approach for Spark


## Block APSP

Number of vertices $=n$, Block Size $=b$


## Block APSP

Iteration 1A: Compute APSP within $V_{1}$ (block 1 on diagonal)


## Block APSP

Iteration 1B: Update weights of all paths to/from $V_{1}$


## Block APSP

Iteration 1C: Update weights of all paths starting and ending in $V_{-1}$ using

$$
S_{i j} \leftarrow \min \left(S_{i j}, S_{i k} \otimes S_{k j}\right) \quad \text { where } k=1
$$



## Block APSP

Iteration 2: Do the same for block 2 on the diagonal


## Block APSP: Single-core

- Block size $b, n / b$ iterations
- A-step (all paths within block): $O\left(b^{3}\right)$
- B-step (all paths to/from block): $O\left(n b^{2}\right)$
- C-step (all paths through block): $O\left(n^{2} b\right)$
- Iteration: $O\left(n^{2} b+n b^{2}+b^{3}\right)$
- Total: $O\left(\frac{n}{b}\left(n^{2} b+n b^{2}+b^{3}\right)\right)=O\left(n^{3}+n^{2} b+n b^{2}\right)$
- The case $b=1$ is almost the same as Floyd-Warshall


## Distributing Block APSP

Problem setup

- Input format: Given by dense adjacency matrix, stored as BlockMatrix $S$ with block size $b$
- Number of vertices $n$ is large
- Output format: same
- Each block fits in memory


## Distributing Block APSP

For $i=1, \ldots, n / b$

- A-step: (update all paths within block)
- One-to-one communication
- Computation $O\left(b^{3}\right)$
- Bandwidth $O\left(b^{2}\right)$
- Runtime $O\left(b^{3}\right)$



## Distributing Block APSP

- B-step: (update all paths to/from block)
- One-to-all communication
- Computation per worker: $O\left(n b^{2} / \sqrt{p}\right)$
- Bandwidth $O\left(b^{2} \sqrt{p}\right)$
- Runtime $O\left(\log (p) b^{2}+b^{2} n / \sqrt{p}\right)$



## Distributing Block APSP

- C-step: (update all paths through block)
- All-to-all communication
- Computation per worker: $O\left(n^{2} b / p\right)$
- Bandwidth: $O(n b \sqrt{p})$
- Runtime: $O\left(n^{2} b / p+n b\right)$



## Distributing Block APSP

Overall cost:

- Total computational cost is $O\left(n^{3}+n^{2} b\right)$ divided evenly among workers plus $O\left(n b^{2}\right)$ on driver
- Total communication cost: $O\left(n^{2} \sqrt{p}\right)$
- Total runtime: $O\left(\frac{n^{3}}{p}+\frac{n^{2} b}{\sqrt{p}}+n^{2}+n b^{2}+n b \log (p)\right)$

Ignoring latency, optimal $b=1$
With latency, runtime is

$$
\frac{n}{b} L+K\left(n b^{2}+\left(n \log (p)+\frac{n^{2}}{\sqrt{p}}\right) b+\frac{n^{3}}{p}+n^{2}\right)
$$

so $b \neq 1$ may be optimal

## More Implementation details on Spark

- Grid Partitioner
- Checkpointing


## Results

$$
\mathrm{n}=500, \mathrm{p}=4 ; \text { Local }[4] 8 \mathrm{~GB}
$$



## Thank you!

