Deep Data Partitioning in MapReduce

Akshai Sarma, Josh Walters

May 19, 2015

What is a data partition?

- Groups data into logical sets.
- Allows for smaller data reads.
- Faster queries.

YAHOO!

Example

- Customers database, partitioned on country.
- School database, partitioned by class code.

What is a deep data partition?

- Also known as a “nested” partition.
- Multiple partitions within each partition.
- Greatly minimize data reads.

Example

- A user database has a `state`, `first_name`, and `last_name` column.
- Layout: `state/last_name/first_name`
- Example: `shire/baggins/bilbo`

YAHOO!

Difficulties with deep data partitions

- Partition cardinality.
- Data locality during processing.

Naive implementation

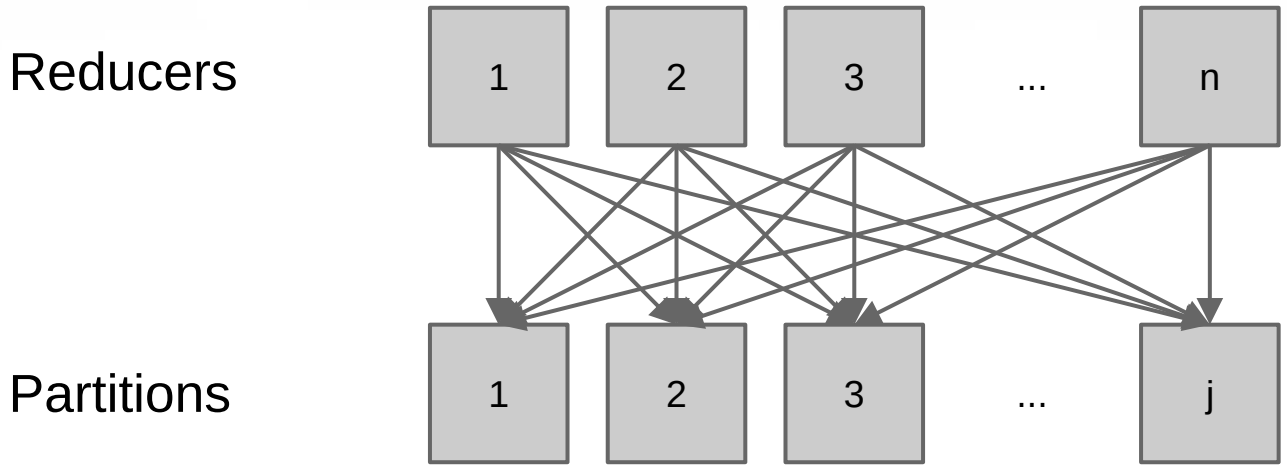

YAHOO!

Dynamic Fractional Partitioning

1. Collect metrics in previous MapReduce job.

1. Compute partition size relative to data set.

1. Assign reducers by partition size ratio.

Partition metrics

1. Count full partition paths.
    a. Eg: `shire/baggins/bilbo == 1`

1. Sum them all up to get a total per partition.

1. Find the ratio of each partition to the whole.

Dynamic Fractional Partitioning

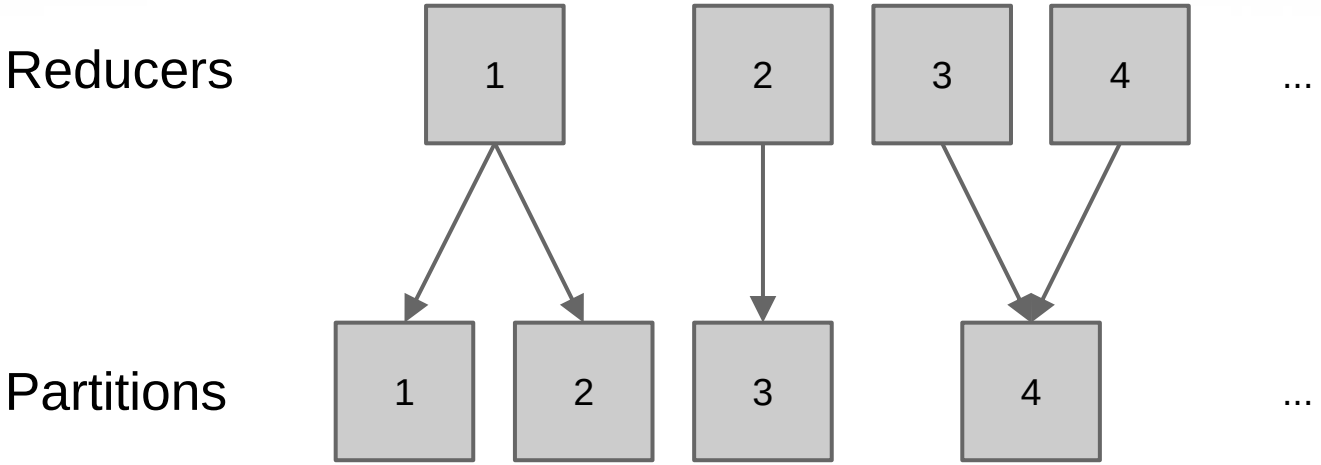

YAHOO!

Thanks!

- Akshai Sarma, asarma@yahoo-inc.com
- Josh Walters, josh@joshwalters.com

YAHOO!