

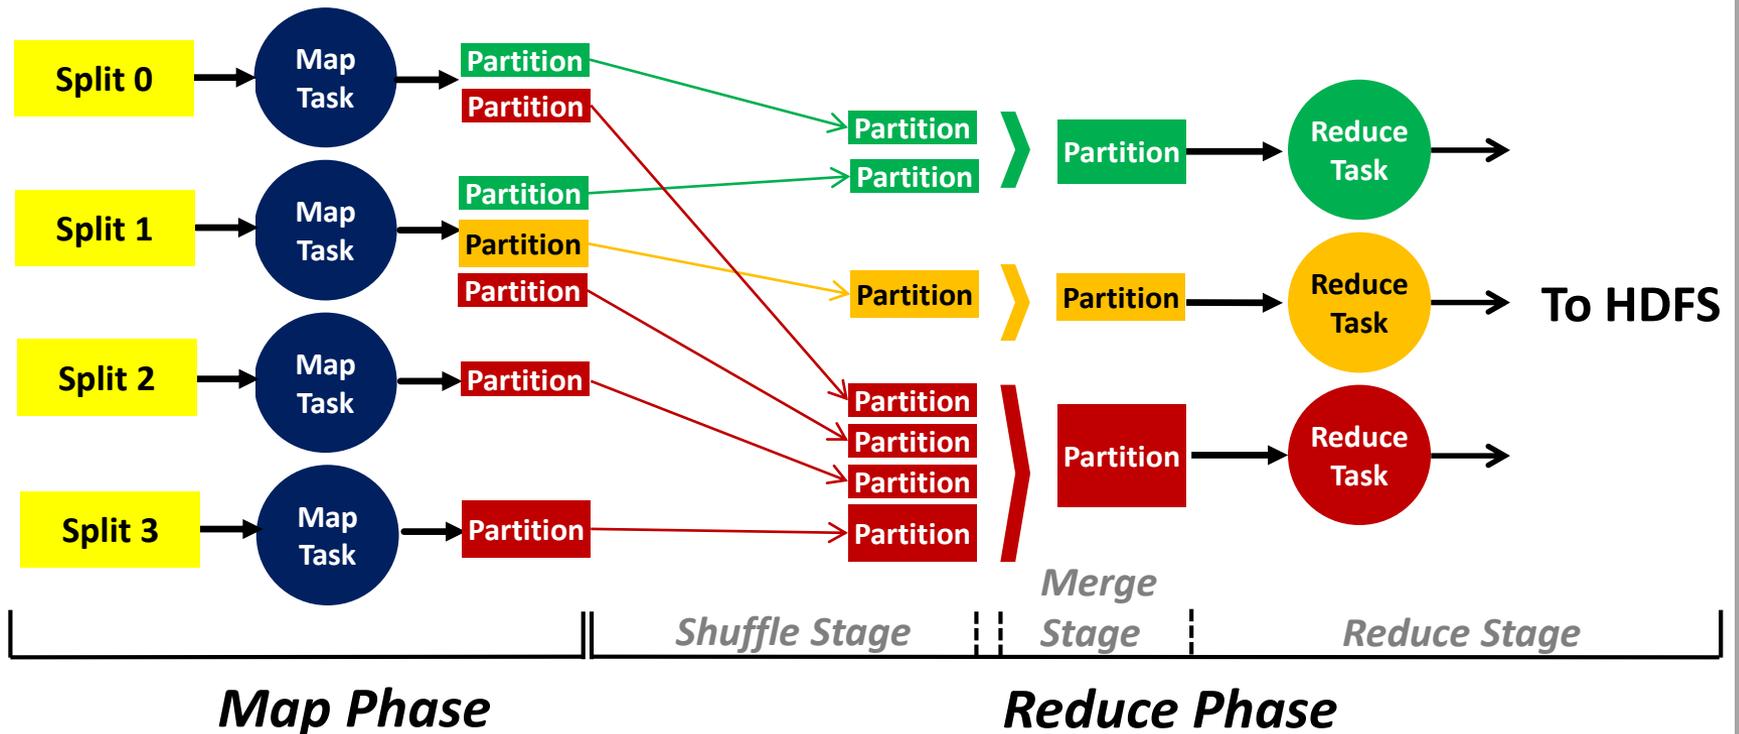
15-440
Distributed Systems
Recitation 11

Tamim Jabban

Project 4

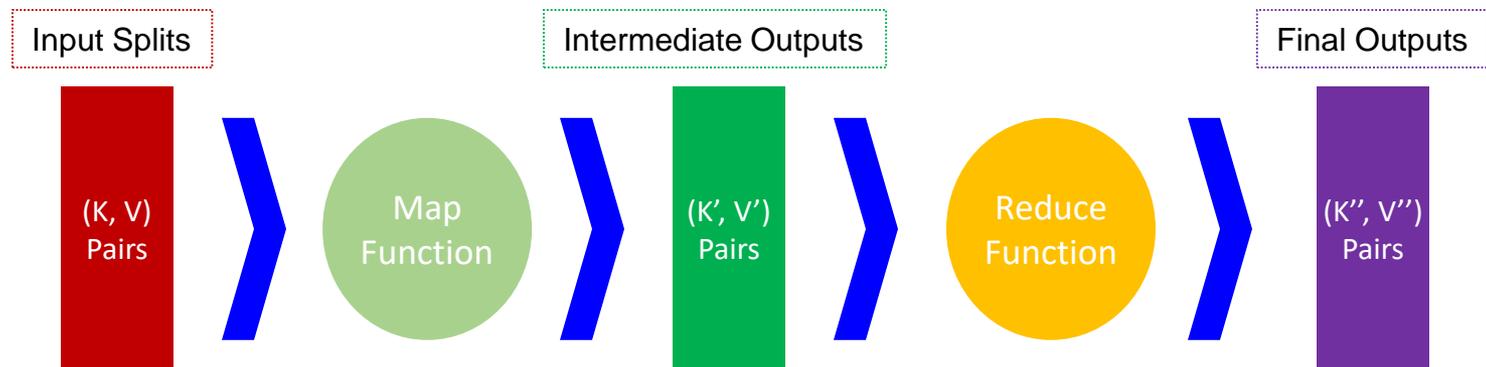
- Apply **MapReduce** to cluster analysis, using the **K-Means** algorithm
- **Due date:** November 30th
 - Please start early!

MapReduce: A Systems View



Data Structure: Keys and Values

- In a MapReduce program, the programmer has to specify two functions: the **Map function** and the **Reduce function** that implement the *Mapper* and the *Reducer*, respectively
- In MapReduce, data elements are always structured as key-value (i.e., (K, V)) pairs
- Therefore, the Map and Reduce functions receive and *emit* (K, V) pairs



MapReduce: An Application View

