

Homework 3

Student Number:

Name:

Problem 1. (30 points) Assume that machines in MapReduce have 100 GB of disk space each. Assume further that the postings list of the term has a size of 200 GB. Then the MapReduce algorithm as described cannot be run to construct the index. How would you modify MapReduce so that it can handle this case?

Problem 2. (30 points) Estimate the space usage of the Reuters dictionary with blocks of size $k = 8$ and $k = 16$ in blocked dictionary storage.

Problem 3. (40 points) For $n = 15$ splits, $r = 10$ segments, and $j = 3$ term partitions, how long would distributed index creation take for Reuters-RCV1 in a MapReduce architecture? Base your assumptions about cluster machines on Table below.

Symbol	Statistic	Value
s	average seek time	$5ms = 5 \times 10^{-3}s$
b	transfer time per byte	$0.02\mu s = 2 \times 10^{-8}s$
	processor's clock rate	$10^9 s^{-1}$
p	lowlevel operation(e.g., compare & swap a word)	$0.01\mu s = 10^{-8}s$
	size of main memory	several GB
	size of disk space	1TBormore