

MOCK QUESTION (date of this version: 27/4/2016)

UNIVERSITY OF EDINBURGH
COLLEGE OF SCIENCE AND ENGINEERING
SCHOOL OF INFORMATICS

MOCK QUESTION: EXTREME COMPUTING

Tuesday 1st April 2014

00:00 to 00:00

INSTRUCTIONS TO CANDIDATES

Answer any TWO questions.

All questions carry equal weight.

CALCULATORS MAY NOT BE USED IN THIS EXAMINATION

Year 4 Courses

Convener: ITO-Will-Determine

External Examiners: ITO-Will-Determine

THIS EXAMINATION WILL BE MARKED ANONYMOUSLY

MOCK QUESTION (date of this version: 27/4/2016)

1. (a) In a normal (non-virtualized) system, what ring do user processes run in? [1 mark]
- (b) Which approach typically results in less data being sent from the mappers to reducers, pairs or stripes? Assume the data is distributed like word co-occurrence. [1 mark]
- (c) Table storage layout.
 - i. List two advantages of column stores over row stores. [2 marks]
 - ii. List two advantages of row stores over column stores. [2 marks]
- (d) MapReduce runs combiners in both the mapper and reducer. Explain why would it run combiners in the reducer. [4 marks]
- (e) You are running a newspaper archive site. Visitors are told they can view a reasonable number of articles each month for free. You are afraid that some users will try to copy the entire site. Assume that these users can be identified by their Internet Protocol address and do not change addresses.
 - i. Some users will try to download a large number of articles in a short period of time by load balancing across your web servers. Name a **mathematical approximation** you could use to efficiently catch them. [2 marks]
 - ii. Other users try to slowly download the entire site over a long period. Would you apply online transaction processing or online analytical processing? [1 mark]
 - iii. What can you do if either system cannot handle all the data, besides adding more machines? [3 marks]
- (f) Sketch an efficient distributed algorithm to construct a Bloom filter. Assume that the elements are stored evenly throughout the cluster and the Bloom filter is small enough to fit in RAM. In sketching your algorithm, explain what each machine does and exactly what data it sends to other nodes. You do not need to write pseudo code. Assume machines never fail. [9 marks]

Module Title: Extreme Computing
Mock Exam Question Solutions 2015

1. (a) 3
- (b) Stripes
- (c) (i). *Any two of the following:* read efficiency if only a few columns are being read in bulk, better compression ratio since repeated data is likely to be consecutive, processing with single-instruction multiple-data, uncompress just the data needed to do the query
- (ii). *Any two of the following:* single random access to read or write a row, easier to implement atomic row operations, less remote procedure call traffic for individual accesses
- (d) The reducers merge output from multiple mappers using n -way merge sort. If the number of map outputs is larger than n , this will require multiple passes and temporary files written to disk. If the reducer encounters multiple records with the same key (there can be up to n when n files are being merged) then it should combine them to decrease the size of intermediate files and subsequent merging work.
See <http://ftp.cs.wisc.edu/pub/techreports/1981/TR445.pdf> for more.
- (e) (i). *Any one of the following:* Sticky sampling, lossy counting.
- (ii). Online analytical processing.
- (iii). Sample.
- (f) Aggregation tree. Communication goes up the tree, with each node of the tree aggregating the answers from itself and the children before sending up to its parent. If the count of entries is not known in advance, aggregate by summing the counts from children and the node itself before sending to the parent. Then the root broadcasts the total count so that all nodes agree on Bloom filter size (we can assume hashes etc are agreed in advance). Each node builds a Bloom filter locally on its own data. A node ORs its Bloom filter with those of its children, sending the combined Bloom filter to the parent. Once this inductive process completes, the parent will have a complete Bloom filter containing all the elements.
I do not expect you to know the term aggregation tree. Solutions that sent all the Bloom filters to the same machine would lose a couple marks due to load on one machine.