15-440: Distributed Systems
Problem Solving Assignment 4

School of Computer Science
Carnegie Mellon University, Qatar
Fall 2015

Assigned date: November 4, 2015
Due date: November 19, 2015

I) Replication and Consistency (40 Points):

1. A design issue in content replication is whether updates to data items are pulled or pushed by or to the clients from servers. As discussed in the class, the former strategy is commonly known as *pull-based* content replication, while the latter one is denoted as *push-based* content replication.

   a) Discuss one main advantage of pull-based versus push-based replications.
   b) Discuss one main disadvantage of pull-based versus push-based replications.
   c) Devise a scenario where neither push-based nor pull-based replication will work well by itself.
   d) Suggest an algorithm that combines both strategies (or something totally different) so as to effectively tackle your presented scenario.
   e) When using your algorithm, is it necessary that the clocks of a client and the server are tightly synchronized? Explain.
   f) Would your algorithm necessarily lead to better performance? Explain.

2. Explain which strategy scales better, totally-ordered multicasting using Lamport's logical clocks or the sequencer approach discussed in the class under active replication.

3. An operation ***OP1*** upon an object ***B1*** causes ***B1*** to invoke an operation ***OP2*** upon another object ***B2***. If ***B1*** is replicated but not ***B2***, what problem(s) will emerge? Suggest a solution for such (a) problem(s).

4. Assume two applications ***A1*** and ***A2*** and two clusters ***C1*** and ***C2***. Suppose that ***A1*** and ***A2*** have high and low read-write ratios, respectively. Suppose that ***C1*** is characterized with scarce network bandwidth while ***C2*** exhibits exactly the opposite. For two types of consistency, sequential and eventual, which application will you deploy at which cluster? Justify your answer.

II) Programming Models (40 Points):

1. A 50-node Hadoop cluster has a very slow node. Tasks that run on that node always end up being marked as stragglers by Hadoop. A job with 12.5 GB of input data is running on a Hadoop cluster using the default configuration. Assume that each node has 2 Map slots and that Map tasks arrive in waves evenly across all the cluster nodes.

   a) How many Map tasks will be launched in total?

   b) How many of those Map tasks will successfully complete and provide inputs to the Reduce tasks?

2. Which of MapReduce, Pregel, or GraphLab is best suited to compute All-Pairs Shortest Paths (APSP) on a weighted, fully-connected Graph? Explain.

3. Given the following graph, assume that vertices are labeled as v_i where *i* is the ID of a vertex. The red, blue and green areas represent the scopes of the vertices v_2, v_3 and v_6, respectively. Assume an initial vertex *schedule* $\tau = (v_2, v_3, v_6)$, wherein v_2 was added to the schedule before v_3, and v_3 was added before v_6.

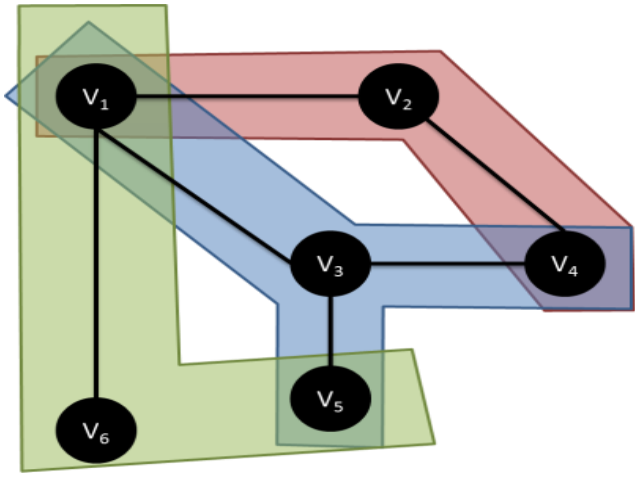

The function $f(v, S_v)$ is defined as $f(v, S_v)$ = the minimum vertex ID among all vertices in S_v. Given that the GraphLab execution engine is run over only one CPU, and can use two types of *schedulers*: (i) FIFO and (ii) Random, what would the content of τ and the returned value be after the first 3 iterations for each scheduler? Demonstrate the results after every iteration, and discuss which scheduler will guarantee a full coverage of all the vertices after exactly 3 iterations.

III) Research Paper Critique (20 Points):

1. Write a critique for the research paper: "**PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs**" by Y. Low *et al.* Please refer to the course webpage on how to write a critique for a research article.