Pregel Execution

1. Many copies of the program begin executing on a cluster
2. The master assigns a partition of input (vertices) to each worker
   - Each worker loads the vertices and marks them as *active*
3. The master instructs each worker to perform a iteration
   - Each worker loops through its active vertices & computes for each vertex
   - Messages can be sent whenever, but need to be delivered before the end of the iteration (i.e., the barrier)
   - When all workers reach iteration barrier, master starts next iteration
4. Computation halts when, in some iteration: no vertices are active and when no messages are in transit
5. Master instructs each worker to save its portion of the graph