Advanced Computer Architecture: Google 2

Jeremy Bradley

8 March 2004

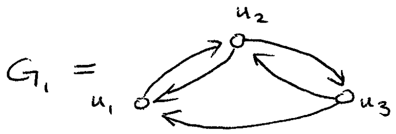

Figure 1: Graph G_1

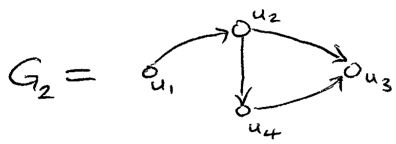

Figure 2: Graph G_2

1. The graph G_1 represents a set of connected web pages. Taking $c = 4/5$, write down:
   (a) The transition matrix P
   (b) The personalisation vector $\vec{p}$
   (c) The modified transition matrix A
   (d) Two iterations of $\vec{x}_{(k+1)} = \vec{x}_{(k)}A$, i.e. $\vec{x}_{(1)}, \vec{x}_{(2)}$ with $\vec{x}_{(0)} = (0, 1, 0)$
   (e) Two iterations of the PageRank algorithm for comparison
   (f) The value $\delta = ||\,|\vec{x}_{(k+1)} - \vec{x}_{(k)}|\,||_1$ for each iteration

   If you have access to a machine you can implement the PageRank algorithm on slide 17 using an ϵ-value of 0.01, instead of doing parts (d) and (e) by hand. How many iterations does it take to converge?

2. Repeat question 1 for the graph G_2, this time with $\vec{x}_{(0)} = (1, 0, 0, 0)$. How does the answer/convergence vary if you alter c?

3. Suggest scalable techniques for implementing each of the vector operations below across several processors, where as necessary each processor has the same partitioned set of rows for each vector $\vec{v}_1$ and $\vec{v}_2$
   (a) $v_1 + v_2$
   (b) $||v_1||_1$
   (c) αv_1 for some scalar multiplier, α