CSE 512/CS 554 Homework Assignment 2
Due February 26, 2008

1. Construction of the Empire State Building required about two person-millennia of effort, yet the project was completed in only about one year. How was this possible? What must the average workforce have been? The peak workforce was about 3500. What "efficiency" does this imply for the overall project?

2. Two people can make up a bed in less than half the time it takes for one person to make up the bed. Explain this apparent example of superlinear speedup.

3. Suppose you have misplaced your glasses in an eight room house, and it takes one minute per room for one person to look for them. Is it possible to achieve superlinear speedup if others help by searching multiple rooms simultaneously? If so, give an example; if not, explain why. Assuming that all rooms are equally likely to contain the glasses, is it possible to achieve superlinear speedup in the *expected* time to find the glasses? Explain your answer.

4. (a) For the $n \times n \times z$ 3-D finite difference grid example given in Chapter 4 on *Parallel Performance*, what is the tradeoff point in n for which a 1-D partitioning performs better or worse than a 2-D partitioning. That is, derive an expression for the value of n as a function of t_c, t_s, t_w, z, and p such that a 1-D partitioning performs better for smaller values of n and a 2-D partitioning performs better for larger values of n.

   (b) What is the asymptotic value for the tradeoff point as p becomes very large?

   (c) For the values $t_c = 20$, $t_s = 1000$, $t_w = 5$, and $z = 10$, what are the largest and smallest values for the tradeoff point in n for $9 \leq p < \infty$?

5. (a) Develop a performance model (parameterized by t_s, etc.) for the execution time of the basic single-node broadcast algorithm given in the section of Chapter 1 on *Collective Communication*, for a message of length L on a hypercube with $p = 2^k$ nodes.

   (b) Develop an analogous performance model for a pipelined single-node broadcast algorithm in which the message of length L is broken into q pieces, which are pipelined along the same spanning tree. Determine an expression for the optimal number of pieces q.

   (c) Develop an analogous performance model for single-node broadcast in a hypercube in which the message of length L is broken into k pieces, which are broadcast using k edge-disjoint spanning trees.

   (d) Among the three preceding algorithms, determine the tradeoff points in L, as a function of the remaining parameters, for which each algorithms is best. For the algorithm of (b), assume that the optimal value for q is used.

(e) For each of the three algorithms, plot the time to broadcast a message of length L, 1 Kbyte $\leq L \leq$ 1 Mbyte, in a hypercube having $p = 1024$ processors, a message latency of 5 msec, and a message transfer rate of 100 Mbyte/sec. Use a log-log scale.

(f) Your answers to the preceding questions depend on your assumptions about the degree of concurrency in using communication links. How much difference does it make if a given processor can communicate on only one link at a time or on all links simultaneously?