6.895 Theory of Parallel Systems	Due: Wednesday, November 12th

Problems 7–11

Problem 7. In this problem we develop an alternative analysis for work stealing. We use a potential-function argument rather than a delay-sequence argument. The advantage is that we need not use the augmented DAG G', but instead we use more mathematics. Let $d(u)$ be the depth of u in G. Define the *weight* of thread u to be

$$w(u) = T_\infty - d(u) \ .$$

Define the *potential* of thread u to be

$$\Phi(u) = 3^{w(u)} \ .$$

Define the potential of a processor p to be the sum of the potentials of the nodes on p (i.e., in its deque or executing). That is:

$$\Phi(p) = \sum_{\text{u on p}} \Phi(u) \ .$$

Define the potential of the system to be the sum of the potentials of the processors. That is,

$$\Phi(S) = \sum_{i=1}^{p} \Phi(i) \ .$$

Now divide the execution into *rounds* as in class.

Problem 7.a. Show that a *constant fraction* of the potential on a processor is at the top of the deque. What fraction?

Problem 7.b. Show that there is a *constant probability* in each round that the potential decreases by at least a constant fraction. Calculate the constants.

Problem 7.c. Assume you have a coin that comes up heads with probability p. Show that with error probability exponentially small in n, if you flip n times, there will be $\Theta(np)$ heads. Calculate constant.

Problem 7.d. Conclude that with probability at least $1-\epsilon$, only $(T_\infty + \log(1/\epsilon))$ rounds (where ϵ is polynomial in P) are necessary before all threads are complete.

Hint: You will need linearity of expectation, 0/1 random variables, and Markov's Inequality. See CLRS.

Problem 8. For the Cilk scheduler, show that the *expected* number of steals is $O(T_\infty \cdot P)$.

Problem 9. Show that the MESI protocol, as presented in Lecture 13, guarantees sequential consistency.

Problem 10. (Extra Credit) Show that the directory based protocol is correct. Note that this involves more carefully specifying how the protocol works.

Problem 11. Recall from class the SPIN-BLOCK PROBLEM: in order to acquire some resource, you can either *spin*, which has cost proportional to how long you spin, or you can *block*, which has fixed cost c.

In class, we saw for the 1-Random-Choice Spin Block Problem that if we block at time $c/2$ with probability $p = 2/5$, then the competitive ratio is 1.8.

Suppose that we get to make one random choice at time $t = \beta \cdot c$ with probability p to block. (The rest of the protocol must be determinisitc.) What are the optimal values of β and p that minimize the competitive ratio? What is the resulting competitive ratio?