

DM538 – Weekly Note 13

Lecture, week 50

Monday, December 10

Kleinberg and Tardos: Algorithm Design (photocopies)

- Section 13.6: Hashing: A Randomized Implementation of Dictionaries
- Section 13.9: Chernoff Bounds
- Section 13.10: Load Balancing

Exercises, week 50

Thursday, December 13

1. Exam January 2011 Problem 5
2. *Load balancing* algorithms for parallel or distributed systems seek to spread out collections of computing jobs over multiple machines. The jobs may come from diverse sources that cannot coordinate.

Suppose there are k machines, and k jobs show up for processing. Consider the following randomized algorithm for load balancing.

Each job is assigned to one of the k machines independently at random (with each machine equally likely).

- (a) Let $N(k)$ denote the expected number of machines that do not receive a job. Give a formula for

$$\lim_{k \rightarrow \infty} \frac{N(k)}{k}$$

- (b) Assume that the machines cannot queue up jobs. Thus, if a machine receives more than one job, it rejects all but the first one. What is the expected total number of rejected jobs?
- (c) Now assume that the machines can each buffer one job. Hence, a machine will reject jobs only if it receives more than two jobs. Let $R(k)$ denote the expected number of rejected jobs. Give a formula for

$$\lim_{k \rightarrow \infty} \frac{R(k)}{k}$$