CS 684 Algorithmic Game Theory **Scribe: Craig Bowles**
Instructor: Eva Tardos **October 21, 2005**

Randomized Strategies

We continue with a load balancing framework but now consider strategies that are randomized, i.e. involve users assigning probabilities to various machines rather than deterministically choosing one. Note that this is a strictly larger set of player strategies than what we have previously considered; the deterministic game can be viewed as a player assigned a probability of 1 to a single machine and 0 to all others. This game is looked at in greater detail by Koutsoupias and Papadimitriou.

1 Framework

- Jobs $j = 1, \ldots, n$ where each job has non-negative weight w_j.
- Identical machines $i = 1, \ldots, m$.
- Strategy for job j: probability $p_{ij} \geq 0$ is the probability of job j going to machine i, where $\sum_{i=1}^{m} p_{ij} = 1$.
- Random variables X_{ij} which is 1 if j chooses machine i and 0 otherwise.
- Load of a machine is random variable $L_i = \sum_{j=1}^{n} X_{ij} w_j$.
- C_j is load on machine that j chooses.

In this framework we do not think of jobs having process times but weights because we do not consider jobs being given an order. All jobs assigned to a machine experience the same delay, which we can think of as just being the sum of job weights on that machine.

2 Nash Equilibria

What becomes more difficult, given this framework, is creating a definition of a Nash in this randomized setting. Obviously a player, i.e. a job, has the goal of choosing a machine with a small load. Though it is completely debatable, we will use the standard probability expectation function as it does make, at least to some degree, intuitive sense. In particular, a job j wishes to use machines i that minimize

$$\mathcal{E}(L_i | j \text{ is on } i) = w_j + \sum_{l \neq j} p_{il} w_l. \tag{1}$$

Given this, since w_j is a constant, we define a Nash equilibrium as follows.

Definition 1 *A randomized solution is a* Nash equilibrium *if, for all jobs j and machines i, $p_{ij} > 0$ if and only if $\sum_{l \neq j} p_{il} w_l$ is minimal among all machines.*

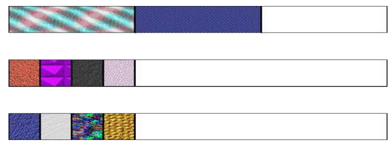

Figure 1: Nash equilibrium in which the jobs shown are effectively deterministic, i.e. all probabilities are either 1 or 0.

Example 2 *In Figure 1 is the "bad" Nash from before where probabilities are either 1 or 0. This is still a Nash under the new settings and definitions. The smaller the jobs are on the second and third machines in the example, the closer the price of anarchy gets to a ratio of 2.*

Example 3 *Consider the uniformly independently random solution, i.e.* $p_{ij} = \frac{1}{m}$ *for all* i, j. *This is referred to just as "Balls and Bins" because everything is so random. This is indeed a Nash because* $\sum_{l \neq j} p_{il} w_l$ *is the same for all machines.*

Good News: Randomized/Mixed Nash Equilibria always exist (see next lecture).
Bad News: Perhaps they can be worse since it's a bigger class.

3 Quality of Solution

Koutsoupias and Papadimitriou consider the quality to be measured by the expected max load on the machines, i.e. $\mathcal{E}(\max_j C_j)$, but as we will see, such a metric is not perfect. We will compare with instead measuring quality by the max expected load on machines, i.e. $\max_j \mathcal{E}(C_j)$. The following example highlights the difference between the two.

Example 4 *Consider again the "Balls and Bins" version, i.e. all jobs and machines are identical with* $w_j = 1$ *for all* j. *Suppose also that* $n = m$. *It is easy to see that the* only *deterministic Nash is where every job is assigned a unique machine, in which case the max load is 1, which is optimal.*

Consider instead the uniformly independently random solution from before, where $p_{ij} = \frac{1}{m} = \frac{1}{n}$ *for all* i, j. *In this case* $\mathcal{E}(C_j) = w_j + \sum_{l \neq j} p_{il} w_l = 1 + \frac{n-1}{n}$, *and thus* $\max_j \mathcal{E}(C_j) \leq 2$. *Compare this to* $\mathcal{E}(\max_j C_j) \approx \log n / \log \log n$. *(The details of this derivation we leave to the paper.)*

The bound of 2 is the same as in the deterministic game. Is this just unique the the Balls and Bins situation? It is not, actually.

Theorem 5 *Max expected load,* $\max_j \mathcal{E}(C_j)$, *in a Nash is at most 2 times that of an optimal solution.*

Proof. This proof is completely analogous to that for the deterministic game. Consider any job j with strategies, i.e. probabilities p_{ij}, in a Nash equilibrium. Let OPT be the optimal (minimal) value of $\max_j \mathcal{E}(C_j)$ achievable by any solutions. Note the following two facts.

- $w_j \leq OPT$, and

- $\frac{1}{m}\sum_l w_l \leq OPT$.

The first point is intuitively clear since no job can get around having to wait long enough for its own completion. The second point is clear as well because it says that OPT is at least the size of the average job size, which for the same reason makes sense.

With those small facts realized, all that is left is to remember that $\mathcal{E}(C_j) = w_j + \sum_{l\neq j} p_{il}w_l$ if i is a machine with $p_{ij} > 0$, and that $p_{ij} > 0$ can only happen for the machine i that minimizes this expression.

Now compare the minimum $\sum_{l\neq j} p_{il}w_l$ to the average over all machines, which is:

$$\frac{1}{m}\sum_i \sum_{l\neq j} p_{il}w_l = \frac{1}{m}\sum_{l\neq j} w_l,$$

as the sum of $\sum_i p_{il} = 1$ for all machines i. So the two lower bounds together imply that

$$\mathcal{E}(C_j) = w_j + \sum_{l\neq j} p_{il}w_l \leq 2OPT,$$

as claimed.