CS 684 Algorithmic Game Theory	Scribe: Xin Qi
Instructor: Eva Tardos	February 25, 2004

In this lecture, we continue to consider the randomized load balancing game. We will evaluate the quality of randomized Nash equilibria with respect to the optimum, but for a different quality measure.

1 Setting of the problem

There are n jobs and m servers. Each job j has weight w_j, and an available subset S_j of servers to choose from. We will use p_i^j to denote the probability for job j to choose server i, so a solution to this game is a set $\{p_j\}_j$ of probability distributions. Each job wants to experience as small load as possible.

The expected load experienced by job j is described as:

$$c_j = \sum_i p_i^j (w_j + \sum_{k \neq j} p_i^k w_k) \tag{1}$$

in which the term $(w_j + \sum_{k \neq j} p_i^k w_k)$ is the load experienced by job j if it goes to server i.

Alternatively, we can use a random variable X_j to denote the load experienced by job j. Then expected load can be expressed as:

$$c_j = Exp(X_j) \tag{2}$$

2 Quality measures

The following are some natural measures of social cost that enable us comparing Nash equilibria with social optima.

1. $\sum_j c_j$
2. $\max_j c_j = \max_j Exp(X_j)$
3. $Exp(\max_j X_j)$

Before, we have concentrated on the second measure. Last time, we discussed that if $S_j = S$ $(\forall j)$ then $\max_j c_j \leq 2 \cdot opt$. In homework, we saw the a special case of pure game where $w_j = 1$ $(\forall j)$. In that case, we know $\max_j c_j \leq O(\log m) \cdot opt$. Actually, that bound also holds for jobs with non-uniform weights, and for randomized games.

In this lecture, we will talk about the third measure. Although it looks like the second one, where the order of max and Exp is exchanged, it is very different. There might be some instances that every job has a very small probability of experiencing large delay, but for any solution, there are always some jobs experiencing large delay. For this kind of instances, the second objective function is very small, while the third one has a large value.

3 An instance with bad Nash equilibrium

In this case, we have $S_j = S$ $(\forall j)$ and $w_j = 1$ $(\forall j)$. We also let $m = n$. There is only one deterministic Nash, in which each job chooses separate server. It is also the optimal solution to the second and the third measures. However, there are many randomized Nash equilibria.

Claim 1 *The solution that every job chooses all servers uniformly randomly, i.e.* $p_i^j = \frac{1}{m}$ $(\forall j, \forall i)$ *is a Nash equilibrium.*

Proof. It is not hard to see this, since every server seems to be equal. To be a little more formal, let us take a look at the load experienced by job j if it goes to server i:

$$w_j + \sum_{k \neq j} p_i^k w_k = w_j + \frac{1}{m} \sum_{k \neq j} w_k = 1 + \frac{n-1}{n} \tag{3}$$

So each server i appears the same to job j, and therefore job j will be indifferent to choose among servers, provided all the other jobs choose servers uniformly randomly. ■

From the above proof, we also know that

$$c_j = 1 + \frac{n-1}{n} < 2 \ \ (\forall j) \tag{4}$$

So the second quality measure, which is the maximum expected load, is good for this solution. However, the third quality measure, the expectation of maximum load, is very bad. According to the results of Balls & Bins problem, which we have learned in CS681, maximum load on a server is $O(\log n / \log \log n)$ with high probability.

The remaining part of this lecture will focus on the proof of the above result, using Chernoff bound.

Chernoff bound (the stronger version) There are n constant numbers $a_1, a_2, \ldots, a_n$, s.t. $0 \leq a_1, a_2, \ldots, a_n \leq 1$, and n random variables $y_1, y_2, \ldots, y_n$, s.t. $y_i \in \{0, 1\}$ $(\forall i)$. Let

$$Y = \sum_{i=1}^{n} y_i a_i \tag{5}$$

and

$$\mu \geq Exp(Y) = \sum_i a_i Pr(y_i = 1) \tag{6}$$

We have the following bound:

$$Pr(Y > (1+\delta)\mu) \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}} \right)^{\mu} \tag{7}$$

Bound on the quality of randomized Nash equilibria Our goal is to use the above Chernoff bound to prove the following result:

$$Exp(\max_j X_j) \le O(\log m) \cdot opt \tag{8}$$

With the stronger version of Chernoff bound described above, we can slightly generalize our problem to the case of non-uniform weights of jobs. Suppose job j has weight w_j. We can assume w.l.o.g., that $w_j \le 1$ $(\forall j)$. This is true as we can always divide all weights by some factor of at least $\max_j w_j$. Actually, we will scale by *op*, which is clearly at least $\max_j w_j$. From last lecture, we know that under this assumption,

$$Exp(y^i) \le 2 \tag{9}$$

We have the following random variables:

$$y_j^i = \begin{cases} 1 & \text{if job } j \text{ is on server } i \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

Then the load of server i can be described as:

$$y^i = \sum_j y_j^i w_j \tag{11}$$

Claim 2 *For any server i, and some constant c,*

$$Pr(y^i > c \cdot \log m) \le \frac{1}{m^2} \tag{12}$$

Why is claim 2 useful? Because it implies:

$$Pr(\max_i y^i > c \cdot \log m) \le \sum_i Pr(y^i > c \cdot \log m) \le \frac{1}{m} \tag{13}$$

and therefore

$$Exp(\max_i y^i) \le \frac{m-1}{m}(\log m + \frac{1}{m}\sum_j w_j) \le c \cdot \log m + 1 \tag{14}$$

Now we will give the proof of claim 2.

Proof. Let $\mu = 2$, since $Exp(y^i) \le 2$, and let $1 + \delta = \log m$. We can then apply Chernoff bound to y^i:

$$\begin{aligned}
Pr(y^i > 2\log m) &= Pr(y^i > (1+\delta)\mu) \\
&\le \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\mu = \left(\frac{e^{\log m - 1}}{(\log m)^{\log m}}\right)^2 \\
&\approx \left(\frac{e^{\log m}}{(\log m)^{\log m}}\right)^2 \\
&= \left(\frac{m}{m^{\log\log m}}\right)^2 \\
&= \left(\frac{1}{m^{\log\log m - 1}}\right)^2
\end{aligned}$$

For m that is big enough, obviously we have $m^{\log\log m - 1} \leq m$, which implies

$$Pr(y^i > 2\log m) \leq \frac{1}{m^2} \tag{15}$$

which is claim 2.

$\blacksquare$