CS 684 Algorithmic Game Theory	**Scribe: Marek Janicki**
Instructor: Eva Tardos	**28/09/05**
Lecturer: Thomas Wexler	

1 A Facility Location Game(A.Vetta 2002)

1.1 Motivation

The facility location game is related to the facility location problem. The facility location problem is a topology problem. Given a topology one wishes to place facilities. Each facility exists to serve clients, and the problem ask how to minimise both client distance and facility creation. This problem has many applications to the real world. However it presumes that there is only one possible entity that can place facilities. The facility location game simplifies the topology somewhat, but allows multiple players to place facilities.

1.2 The Elements of the Game

The game consists of markets, suppliers and locations. There are m markets, each denoted m_i. There are k suppliers, and each supplier k has an associated set of locations $L^k \subseteq L$, where L is the set of all locations and L_j is a particular location.

Each market m_i has a value Π_i associated with it. This is the value it receives when served. Alternately one can think of it as the highest value that it is willing to pay the suppliers for there supplies. Between each market m_i and each location L_j there is an edge with weight λ_{ij}. This represents the cost of serving supplies to m_i from L_j. Each supplier k may build a single facility at a location L_i contained within L^k. Because only the suppliers have a direct influence on their choice of facility, it is the suppliers who play the game. A solution is then just a mapping of suppliers to locations. Having picked a location, a supplier then tries to maximise his profit from that location.

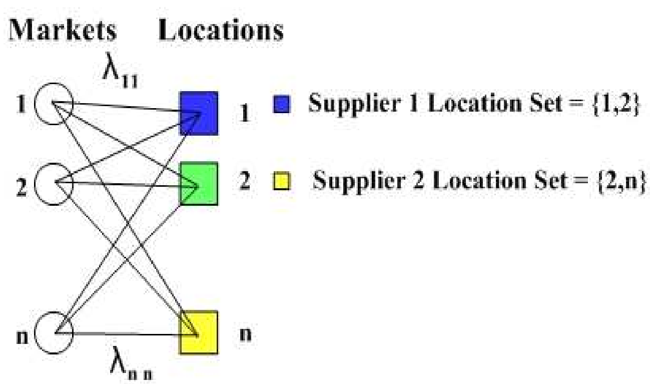

Figure 1: The Facility Location Game

We note that a supplier will only supply a market if he can profit from it. So, anywhere where

$\lambda_{ij} \geq \Pi_i$ the supplier will not attempt to serve the market. This means that $\forall\ \lambda_{ij} \geq \Pi_i$ we can set $\lambda_{ij} = \Pi_i$. We will do so. Thus, λ_{ij} must satisfy: $0 \leq \lambda_{ij} \leq \Pi_i$.

1.3 Characterising the Solution to this Game

Given a solution we would like some way to characterise the quality of this solution. We note that once we have a solution, we can determine what the suppliers will do. To do this we look at a market m_i. In the solution there are $L^i \subseteq L$ open facilities that can serve m_i. Clearly the supplier that will end up serving m_i is the supplier that can serve m_i at cheapest cost, as he can undercut all of his competitor. So we say that m_i is served by $\sigma(i)$, where we define $\sigma(i)$:

$$\sigma(i) = arg(\min_{j \in L^i} \lambda_{ij})$$

The price that this supplier can charge should be equal to the cost that his closest competitor incurs. If he sets it lower, then there is still profit to be made, by raising the price to that level. However if he sets it higher, than his closest competitor can undercut him and he will get no business. So, we say that m_i pays a price p_i where we define p_i:

$$p_i = \min_{\substack{j \in L^i \\ j \neq \sigma(i)}} \lambda_{ij}$$

Note that $\sigma(i)$ defines a location, while p_i defines a cost.

1.4 Quality of the Solution to this Game

To measure the quality of a solution we try and see who benefits when m_i is served at a price p_i. Now, m_i is willing to pay up to Π_i, but it benefits more by paying less. So, a natural benefit function for m_i is $\Pi_i - p_i$. At the same time, the supplier at $\sigma(i)$ is getting p_i but paying $\lambda_{i\sigma(i)}$ so the natural benefit function for $\sigma(i)$ is $p_i - \lambda_{i\sigma(i)}$. The total benefit is just the sum of the benefits of the markets, and the benefits of the suppliers. But we note that at most one supplier is getting any benefit from a given market, which enables us to write.

$$\begin{aligned} TotalBenefit &= \sum_{\text{i} \mid m_i \text{ served}} (\Pi_i - p_i) + \sum_{\text{i} \mid m_i \text{ served}} (p_i - \lambda_{i\sigma(i)}) \\ TotalBenefit &= \sum_{\text{i} \mid m_i \text{ served}} (\Pi_i - \lambda_{i\sigma(i)}) \end{aligned}$$

To make this slightly cleaner we will assume that if m_i is not served by any supplier, it is in fact served by some supplier at cost. That is, $\Pi_i = \lambda_{i\sigma(i)}$. This does not change the value of the total benefit, but it gives us the slightly cleaner form:

$$TotalBenefit = \sum_{\text{all i}} (\Pi_i - \lambda_{i\sigma(i)})$$

2 Nash Equilibria in the Facility Location Game

2.1 Pertinent Questions

There are three questions we typically ask about Nash Equilibria:

- Does a Nash equilibrium exist?
- If it does exist, is it unique?
- If it does exist, how does it compare to the optimal solution?

The answers to these questions follow.

2.2 The Existence of a Nash Equilibrium

Theorem 1 *The Facility Location game is a Potential Game with Potential function* Φ *where*

$$\Phi = \sum_{all\ i} \lambda_{i\sigma(i)}$$

Proof. We need to show that Φ tracks a player's benefit change when he switches. So, let us consider user k. Let us take him out of the game. That is, let him spontaneously decide to leave. Now we consider all of the $\lambda_{i\sigma(i)}$ where $\sigma(i) = k$. We note that by the definition of p_i we get $\lambda_{i\sigma^{new}(i)} = p_i$. So

$$\Delta\Phi = \sum_{i|\sigma(i)=k} (p_i - \lambda_{i\sigma^{old}(i)})$$

This is exactly the loss of profits that user k cost himself when he left the graph. So Φ tracks the losses properly.

Now we consider what happens when k jumps back into the game. At each market he'll either get nothing in which case $\lambda_{i\sigma(i)}$ won't change, or (if $k = \sigma^{new}(i)$) he'll get $p_i^{new} - \lambda_{i\sigma^{new}(i)}$ where p_i^{new} is trivially $\lambda_{i\sigma^{old}(i)}$. But this is exactly the difference in Φ for that market. When we consider all of the markets that the user touches, we note that the benefit to the user is exactly the increase in Φ. So Φ accurately tracks the change in each user's benefit, so Φ is a potential function. ■

Corollary 1 *A Nash equilibrium exists,*

Proof. In a potential game, the minima of the potential function (Φ) are Nash equilibria, and we can clearly get to these minima by forcing users to change if and only if doing so would lessen Φ and forbidding them to do so when it wouldn't. ■

Corollary 2 *The global minimum of* Φ *is the optimal solution. So the best Nash equilibrium is also the best solution.*

Proof. Note that:

$$TotalBenefit = \sum_{\text{all i}} (\Pi_i - \lambda_{i\sigma(i)})$$

But $\sum_{\text{all i}} \Pi_i$ is a constant. So then to maximise the total benefit we have to minimise $\sum_{\text{all i}} \lambda_{i\sigma(i)}$ which is exactly what the global minimum of Φ does. ■

2.3 The Uniqueness of a Nash Equilbrium

The following example, where $L^1 = \{1, 2\}$ and $L^2 = \{3, 4\}$ and $\forall i\ \Pi_i = 1$, shows that there is not a unique Nash Equilibrium.

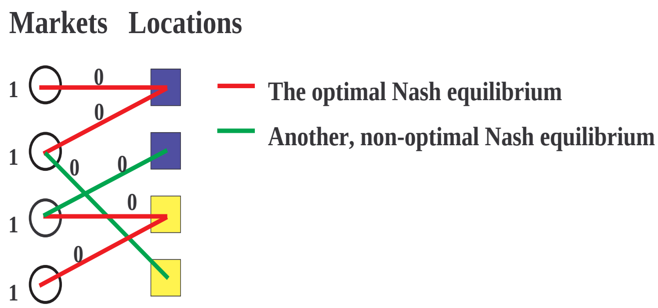

Figure 2: A Facility Location Game with two Nash Equilibria. Edges not shown on the graph have $\lambda_{ij} = 1$, otherwise they have $\lambda_{ij} = 0$.

We note that the red equilibrium has a total benefit of 4, while the green equilibrium has a total benefit of 2, or half that of the optimal. This leads nicely to the next section of the notes, namely:

2.4 The Quality of a Nash Equilibrium

Theorem 2 *The total benefit of any Nash Equilibrium is at least 1/2 of the total benefit of the optimal solution.*

Proof. We want to compare a Nash equilibrium to an optimal solution. To do this we will use regular notation for the Nash equilibrium and primed notation for the optimal solution. As an example $\sigma(i)$ is the location m_i gets assigned to in the Nash equilibrium, and $\sigma'(i)$ is the location m_i gets assigned to in the optimal solution.
Now we introduce some more notation. Let $val(k)$ be the total profit that supplier k gets in the Nash equilibrium, and let $val'(k)$ be the total profit that supplier k gets if everyone else keeps their location in the Nash equilibrium, but supplier k uses his location in the optimal solution. Clearly $val(k) \geq val'(k)$ as $val(k)$ refers to a Nash equilibrium. Finally we define $\delta(i)$ as

$$\delta(i) = \lambda_{i\sigma(i)} - \lambda_{i\sigma'(i)}$$

This gives use the following Lemma:

Lemma 3

$$val'(k) \geq \sum_{all\ i\ |\ k\ supplies\ them\ in\ opt.\ soln.} \delta(i)$$

Proof. Let us consider the portion of $val'(k)$ that comes from a single m_i that k serves in the optimal solution, and compare it to $\delta(i)$. This portion is at least 0. If it is 0 then k is not serving that m_i, so $\lambda_{i\sigma(i)} > \lambda_{i\sigma'(i)}$ in which case $\delta(i) < 0$ which is what we want. Otherwise k is getting some profit from m_i. This profit is $p'_i - \lambda_{i\sigma'(i)}$, where p'_i must be $\lambda_{i\sigma(i)}$. So we get:

$$\begin{aligned} p'_i - \lambda_{i\sigma'(i)} &\geq \delta(i) \\ p'_i - \lambda_{i\sigma'(i)} &\geq \lambda_{i\sigma(i)} - \lambda_{i\sigma'(i)} \\ p'_i &\geq \lambda_{i\sigma(i)} \end{aligned}$$

Thus, each individual element of $val'(k)$ is less than or equal to $\delta(i)$. So the sum over all elements must also have this property and the lemma is proven. ■

But this implies that $\sum_i val(k) \geq \sum_i \delta(i)$. But note that

$$\begin{aligned} \sum_{\text{all i}} \delta(i) &= \sum_{\text{all i}} (\lambda_{i\sigma(i)} - \lambda_{i\sigma'(i)}) \\ \sum_{\text{all i}} \delta(i) &= \sum_{\text{all i}} (\lambda_{i\sigma(i)} - \lambda_{i\sigma'(i)} + \Pi_i - \Pi_i) \\ \sum_{\text{all i}} \delta(i) &= \sum_{\text{all i}} (\Pi_i - \lambda_{i\sigma'(i)}) - \sum_{\text{all i}} (\Pi_i - \lambda_{i\sigma(i)}) \\ \sum_{\text{all i}} \delta(i) &= TotalBenefit(Opt.) - TotalBenefit(Nash) \end{aligned}$$

But we know that $TotalBenefit(Nash) \geq \sum_i val(k) \geq \sum_i \delta(i)$ as the sum over val(k) only considers the benefit to the suppliers. Then we have:

$$\begin{aligned} \sum_{\text{all i}} val(i) &\geq TotalBenefit(Opt.) - TotalBenefit(Nash) \\ TotalBenefit(Nash) &\geq TotalBenefit(Opt.) - TotalBenefit(Nash) \\ 2 * TotalBenefit(Nash) &\geq TotalBenefit(Opt.) \\ TotalBenefit(Nash) &\geq TotalBenefit(Opt.)/2 \end{aligned}$$

■