Hotelling Games on Disconnected Graphs

Mateo Espinosa Zarlenga
Department of Computer Science
Cornell University
Ithaca, NY 14850
me326@cornell.edu

Michael Luo
Department of Computer Science
Cornell University
Ithaca, NY 14850
mrl233@cornell.edu

Abstract

Hotelling games are a well-studied family of games for modeling resource competition when both users and service providers are spatially distributed across some topology. In this paper, we formalize a general framework for Hotelling games in undirected weighted graphs where the underlying network of locations can be disconnected and players can be arbitrarily distributed across nodes. Although traditional Hotelling games are constant sum in nature, in this more general setting the social welfare achieved by different strategy profiles can vary widely. This implies that studying the social welfare of equilibrium solutions in these games becomes not only meaningful but can lead to further understanding of how selfish behavior affects the overall good in markets where spatial competition in dominant. Our main contribution is proving a novel tight upper bound for the price of anarchy with respect to the social welfare in this family of games. This upper bound is dependent only on the number of sellers and is asymptotically bounded by two.

1 Introduction

1.1 Motivation

Hotelling games [1, 2] are a widely studied model for understanding competition across users when a service is spatially distributed across some topology. These games attempt to model how suppliers of a single service act when they are given the option to provide that service in one of many possible locations. Although each seller will always prefer choosing a location close to as many customers as possible, she must also avoid other sellers with whom she would have to compete for nearby customers. The essence of Hotelling games is then to capture these dynamics through a network of connected sites and a distribution of clients across those sites.

Previous work has examined equilibria in Hotelling games and the dynamics that lead to them [3]. The existance of equilibria in these games brings forth the question of understanding how good these equilibria are for the players. This concept is perfectly quantified with the *price of anarchy* [4], where the objective function at the social optimum is compared with the objective function at equilibria. While previous work has managed to obtain price of anarchy bounds for Hotelling games [5], they have used fairness rather than social welfare as their objective function. This metric is arguably less meaningful than social welfare (which represents not only how much utility the players receive, but also how many customers are served). In addition, previous work on understanding the price of anarchy in these games presumes that the network of locations is connected, i.e., that every location can be reached from every other location. This assumption disregards the fact that, in real-world markets, networks are often disconnected and not all locations can be reached by all clients.

1.2 Overview of Results

Our work extends the existing work for Hotelling games and attempts to address both of the primary concerns described above. We first formalize a more general and realistic framework for Hotelling games where the underlying network of locations can be disconnected. In this variation, we will allow players to be arbitrarily distributed across the nodes and the graph edges to be weighted. Then, using this framework, we note that the social welfare can vary widely from strategy profile to strategy profile and therefore one can meaningfully analyze the price of anarchy with respect to the total social welfare. We conclude our work by providing a proof for an upper bound on the price of anarchy, and we provide a simple construction to prove its tightness.

2 Background

2.1 Pure Nash Equilibrium

In a game with multiple players, a Pure Nash equilibrium (or PNE for short) is a strategy profile such that no player has an incentive to deviate [6]. This means that the utility that any player receives from the Nash equilibrium is at least as great as the utility she would receive from any strategy profile achieved by a unilateral deviation in strategy on her part.

Formally speaking, if a strategy profile $\boldsymbol{s}$ is a PNE and we let $(\boldsymbol{s}_{-l}, a)$ be the strategy profile where player l changes her strategy to a while all other players keep their strategies as in $\boldsymbol{s}$, then for all players l and all allowed strategies a for player l we must have that

$$u_l(\boldsymbol{s}) \geq u_l(\boldsymbol{s}_{-l}, a)$$

where $u_l(\boldsymbol{s})$ is the utility player l gets in strategy profile $\boldsymbol{s}$.

Pure Nash Equilibria are inherent properties of a game and can be reached in real-world scenarios through selfish behavior from players [7]. Therefore, understanding the nature of these solutions can help understanding the dynamics of selfish players in the underlying game.

2.2 Price of Anarchy

In a utility game with an arbitrary objective function $f(\boldsymbol{s})$ and at least one pure Nash equilibrium, the Price Of Anarchy (POA) [4] is the maximum possible value of the ratio between the objective function evaluated at its optimal maximizing strategy and the objective function evaluated at a PNE strategy. Formally, if OPT is the optimal value of the objective function and $\boldsymbol{s}$ is an arbitrary PNE, the POA is defined as

$$\text{POA} \stackrel{\text{def}}{=} \max_{\boldsymbol{s}} \left(\frac{\text{OPT}}{f(\boldsymbol{s})} \right)$$

Note that to maximize this quantity, one should find a PNE that minimizes the provided objective function. That means that this ratio quantifies the worst-case scenario cost of allowing players to selfishly act pursuing individual utility maximization, rather than force them into an optimal configuration that may not be necessarily an equilibrium.

2.3 Hotelling Games

Hotelling games first came as a generalization of a thought experiment proposed in 1929 by Harold Hotelling [1]. In his dissertation, Hotelling studied an scenario where two sellers want to provide the same service on a street where costumers are uniformly distributed [1, 8]. Through his mental experiment, Hotelling exposed the clear difference between the socially optimal configuration for the providers and the configuration that is achieved through iterative strategic behavior by players [8]. While this game is very simple in nature, in the years since it was proposed, Hotelling's work has been extended [9, 10, 11] and generalized into a full family of games: the so-called Hotelling games [2, 12].

A common formulation of Hotelling games is a game in which players need to select a location in a given graph from they will offer a service to customers. Formally, the game is defined by a graph G across which a fixed N number of customers are distributed. There exists k providers (i.e., players) which wish to offer a service to customers in the graph, and each provider is allowed to pick one and only one location on the network where she will operate her business. Every customer patronizes the nearest provider, with ties broken uniformly at random. Intuitively, providers would like to attract as many customers as possible (in expectation) with their picked location. It is assumed that the graph is connected, so every customer must pick some provider. This means that no matter how the providers are configured, the social welfare is constant and maximal [5]. Thus in traditional analyses, the objective function is not social welfare but rather fairness, the minimum number of customers any provider gets. In [5], the authors prove that the price of anarchy for Hotelling games when fairness is the objective function is $(2k-2)/k$.

3 Model Description

For this paper we work with the following generalized version of a Hotelling game. We are given a graph $G(V, E)$ that has a total of n locations (i.e., $|V| = n$) and every location $v \in V$ has a total of n_v clients associated with it. Intuitively, the number of clients in a node is a factor that determines how good a location is for a given seller. In our model we allow our network to be weighted so that every edge $(u, v) \in E$ has a non-negative weight $w_{u,v}$ associated with it. Note that we specifically do not restrict G to be a connected graph. This is one of the main differences between our results and the results obtained concerning the price of anarchy for Hotelling games in [5]. The fact that G is not necessarily connected implies that G can be formed of one or more connected components that are disconnected from each other. Given this graph, a Hotelling game develops as follows:

- There are a total of k players $\{1, 2, ..., k\}$ that want to provide the same service to all the clients in G.
- Each player is allowed to select a node in the graph to provide her service on that site. These selections define a strategy profile $\boldsymbol{s}$ where $\boldsymbol{s}_i \in V$ is the site selected by player i. For simplicity we define $\mathcal{S}$ as the set of sites selected by all players. Note that this set could have a cardinality between 1 and $\min\big(n, k\big)$; however it can never be empty.
- For all locations $v \in V$, all n_v clients in this node are allocated to the locations in $\mathcal{S}$ that are connected to and closest to v. In this case we define the distance between u and v along path P as the summation of the weights of the edges in P. When there is no path between locations u and v then we say that the distance between these two nodes is infinite. If there is some location v whose distance to all the sites in $\mathcal{S}$ is infinite, then the n_v clients in this node will not be provided the service given by the sellers and these clients will not be allocated to any locations in $\mathcal{S}$. Finally, if there are multiple sellers equally as close to v, we evenly distribute n_v among all these players.
- The utility of player i, called $u_i(\boldsymbol{s})$, is given by the total number of clients that are assigned to player i using the distribution scheme defined above.

Note that allowing this graph to be disconnected models the idea that in real markets some locations are not accessible to all clients and thus sellers must consider this when choosing where to provide their service. As an example we could think of every connected component of G as a different city and sellers are deciding in which city (and location within the city) they would like to provide their service, knowing that this service would not be available to clients outside of this city.

4 Results

In this section we analyze the social welfare and price of anarchy on the generalized Hotelling games defined above. We show that given any arbitrary equilibrium, the social welfare of this solution is no more than a factor of $\frac{2k-1}{k}$ times smaller than the optimal social welfare. Before getting into our proof, however, we justify the use of the social welfare as our objective function.

4.1 Objective Function: Social Welfare in Disconnected Graphs

In this paper we analyze how bad social welfare can be when players reach pure Nash equilibria. One important thing to note here is that while previous work on analyzing Hotelling games used other objective functions such as fairness (the minimum utility over all players), we use social welfare. This is because, in our generalized model, social welfare becomes a good measure of how good a strategy profile is. To see why this is true we note that when the graph is connected, as is assumed for [5], the social welfare of an arbitrary strategy profile $\boldsymbol{s}$ is given by

$$\text{SW}(\boldsymbol{s}) = \sum_i u_i(\boldsymbol{s}) = \sum_{v \in V} n_v$$

where the last equality holds because every site $v \in V$ is contributing exactly n_v to the social welfare, as all clients are added to the utilities of the players with the closest sites to them. Because this summation is a constant and completely independent of the given strategy profile, analyzing the price of anarchy in connected graphs when the objective function is the social welfare is pointless. Nevertheless this is not necessarily true when the graph is not required to be connected.

Assume that the given graph G is not necessarily connected and thus is can be partitioned into m disjoint connected components $C_1, C_2, ..., C_m \subseteq V$. Furthermore, for all $1 \leq i \leq m$, define $\mathcal{N}_i$ as

$$\mathcal{N}_i \stackrel{\text{def}}{=} \sum_{v \in C_i} n_v$$

This is simply the total number of clients contained in component C_i. We also define N as

$$N \stackrel{\text{def}}{=} \sum_{v \in V} n_v$$

Now consider an arbitrary strategy profile $\boldsymbol{s}$. The social welfare of this solution can be expressed as

$$\text{SW}(\boldsymbol{s}) = \sum_i u_i(\boldsymbol{s}) = \sum_{j=1}^{m} \sum_{i | \boldsymbol{s}_i \in C_j} u_i(\boldsymbol{s})$$

Note that the last equality is simply a consequence that every player can select exactly one location and that location has to belong to some component. Now we notice, as we did before for the case of the connected graph, that a site $v \in C_i$ contributes exactly n_v to the utilities of the players whose sites are in component C_i. This site cannot contribute to the utilities of players in other components as this would mean that there is a path from C_i to another component; contradicting the fact that all components are disconnected. This means that if a component C_i has at least one player whose selected site is in C_i, then the social welfare increases by $\mathcal{N}_i$ regardless of how many players selected their sites within component C_i. This means that the summation above can be rewritten as

$$\text{SW}(\boldsymbol{s}) = \sum_{j=1}^{m} \mathcal{N}_j * \mathbb{I}_j(\boldsymbol{s}) \tag{1}$$

where we define the indicator variable $\mathbb{I}_j$ as

$$\mathbb{I}_j(\boldsymbol{s}) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } \exists i, \text{ s.t. } \boldsymbol{s}_i \in C_j \\ 0 & \text{otherwise} \end{cases}$$

In other words, the social welfare of a strategy profile $\boldsymbol{s}$ is going to be an indicator of how many clients are being offered the service in their vicinity. This means that the total number of clients that are being denied the service provided by the sellers with strategy profile $\boldsymbol{s}$ is given by $N - \text{SW}(\boldsymbol{s})$.

Once we allow the graph to be discontinuous, the social welfare of a solution is no longer trivially determined. In fact, it is an important metric to measure the quality of our solution because it not only represents the total utility of the sellers, but also the total number of customers who have access to the service. Thus, in this paper we will use social welfare as the objective function with which to analyze the strategy profiles.

4.2 Bound for the Price of Anarchy

Now that we have given intuition as to why we decided to use the social welfare as our objective function in analyzing Hotelling games, we proceed to show a tight upper bound on the price of anarchy for this family of games. Before getting to the main result of this paper, however, we first show a couple of theorems:

Theorem 1. *Consider a Hotelling game with k players and graph G. Let $C_1, C_2, ..., C_m$ be the distinct connected components that form G where, without loss of generality, we assume that $\mathcal{N}_1 \geq \mathcal{N}_2 \geq \cdots \geq \mathcal{N}_m$. Furthermore let $\boldsymbol{s}$ be a pure Nash equilibrium and let i^* be the number of components in which $\boldsymbol{s}$ has allocated sellers. If we define the **multiset** $A(\boldsymbol{s})$ as*

$$A(\boldsymbol{s}) \stackrel{\text{def}}{=} \{\mathcal{N}_i \mid \exists\, j \text{ s.t. } \boldsymbol{s}_j \in C_i\}$$

and let $i^* = |A(\boldsymbol{s})|$ then we must have that

$$A(\boldsymbol{s}) = \{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_{i^*}\}$$

Where $\{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_{i^*}\}$ is a multiset.

Proof. For this proof, whenever we use sets we will assume we are working with multisets. This means that the elements within the defined sets can have a multiplicity greater than 1.
We will proceed to show this theorem by showing that for all $\mathcal{N}_i \in A(\boldsymbol{s})$ and every $\mathcal{N}_j \in \{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_m\}/A(\boldsymbol{s})$ we must have that $\mathcal{N}_i \geq \mathcal{N}_j$. To see why this is true let's assume, for the sake of contradiction, that there exists an $\mathcal{N}_i \in A(\boldsymbol{s})$ and a $\mathcal{N}_j \in \{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_m\}/A(\boldsymbol{s})$ such that $\mathcal{N}_i < \mathcal{N}_j$. Now, by definition of $\mathcal{N}_j$, we know that solution $\boldsymbol{s}$ did not allocate any sellers in component C_j however it did allocate at least one player, call it l, in component C_i; again by definition of $\mathcal{N}_i$. Consider what happens when player l switches to a random location $t \in C_j$ while every other player keeps the same strategy as in $\boldsymbol{s}$. In this case player l will be the only player allocated to component C_j. This means that all clients in this component will contribute to l's utility as there are no other players in this component offering the service (i.e., player l will hold a monopoly in C_j). Thus the utility of player l in this new strategy profile will be exactly equal to $\mathcal{N}_j$. Furthermore clearly the maximum utility player l could obtain in the original profile $\boldsymbol{s}$ is $\mathcal{N}_i$ as this is the total number of clients that are available in the component containing l's location in solution $\boldsymbol{s}$. These two facts give us the following inequality

$$\begin{aligned} u_l(\boldsymbol{s_{-l}}, t) &= \mathcal{N}_j \\ &> \mathcal{N}_i \\ &\geq u_l(\boldsymbol{s}) \end{aligned}$$

where the first inequality comes from the fact that $\mathcal{N}_j > \mathcal{N}_i$. This result, however, implies that player l would benefit from deviating from her current strategy. This contradicts the fact that $\boldsymbol{s}$ is a Pure Nash Equilibrium. Hence we can conclude that for every $\mathcal{N}_i \in A(\boldsymbol{s})$ and every $\mathcal{N}_j \in \{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_m\}/A(\boldsymbol{s})$ we must have that $\mathcal{N}_i \geq \mathcal{N}_j$.

With the result shown above we see that if a PNE allocates sellers to a total of i^* components, then these components must have the largest i^* total number of clients as otherwise the claim we just proved could not hold. Given that we labeled the components such that the number of clients of the largest i^* components (with respect to the total number of people in them) are $\{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_{i^*}\}$, we can conclude that the multisets $A(s)$ and $\{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_{i^*}\}$ must be equal. This is what we originally wanted to show.

□

Intuitively this proof is saying that when there are multiple disconnected components where sellers could potentially provide a service, any seller i will always prefer setting up a monopoly in an isolated component except when there is a component with a large enough clientele to support healthy competition amongst all providers in that component. The following claim is an immediate corollary of theorem 1 and equation (1):

Corollary 1.1. *If $\boldsymbol{s}$ is a PNE for a Hotelling game with graph G and connected components $C_1, C_2, ..., C_m$ such that $\mathcal{N}_1 \geq \mathcal{N}_2 \geq \cdots \geq \mathcal{N}_m$, then the social welfare of $\boldsymbol{s}$ will be given by*

$$\text{SW}(\boldsymbol{s}) = \sum_{i=1}^{i^*} \mathcal{N}_i$$

Where i^* is the number of distinct connected components that have at least one seller in them in solution $\boldsymbol{s}$.

Before we dive into looking at the POA of our variation of Hotelling games, we first show a theorem that describes how good a solution can be, in the best case scenario, for any given game setup:

Theorem 2. *Given an optimal solution $\boldsymbol{s}^*$ of a Hotelling game with k players, graph G, and connected components $C_1, C_2, ..., C_m$, where we order these components such that $\mathcal{N}_1 \geq \mathcal{N}_2 \geq \cdots \geq \mathcal{N}_m$, the social welfare so such solution will be equal to*

$$\text{OPT} = \text{SW}(\boldsymbol{s}^*) = \sum_{i=1}^{\min(k,m)} \mathcal{N}_i$$

Proof. For this proof we say that the clients in location $v \in V$ are satisfied in solution $\boldsymbol{s}$ if and only if there exist $1 \leq i \leq k$ and $1 \leq j \leq m$ such that $\boldsymbol{s}_i \in C_j$ and $v \in C_j$. Intuitively, the clients in location v are satisfied if they have a path to a provider and can thus access the service. We will further refer to the "total number of satisfied players" by strategy profile $\boldsymbol{s}$ as the sum of all clients that are satisfied by $\boldsymbol{s}$.

Recall from (1) that the social welfare of a solution $\boldsymbol{s}$ can be expressed as

$$\text{SW}(\boldsymbol{s}) = \sum_{j=1}^{m} \mathbb{I}_j(\boldsymbol{s})\mathcal{N}_j$$

Note that the number of components that have at least one site with a seller in it can be expressed as the summation $\sum_{j=1}^{m} \mathbb{I}_j(\boldsymbol{s})$. Clearly, this summation has to be at least one. This would correspond to the case where all players select sites within the same component. Moreover, this summation cannot be more than $\min(k, m)$. It is obvious that we could never have more sellers in distinct components than the total number of unique components m. Similarly, we cannot have more sellers in unique components than the number of players.

Hence for any strategy profile $\boldsymbol{s}$, the number of unique components that have at least one seller in them can never exceed $\min(k, m)$. This means that the total number of satisfied clients by strategy profile $\boldsymbol{s}$ can never be greater than

$$\sum_{i=1}^{\min(k,m)} \mathcal{N}_i$$

As the components are numbered such that $\mathcal{N}_1 \geq \mathcal{N}_2 \geq \cdots \geq \mathcal{N}_m$. Note that we could build a solution that satisfies exactly by $\sum_{i=1}^{\min(k,m)} \mathcal{N}_i$ clients by "assigning" the first $\min(k, m)$ players to the largest $\min(k, m)$ components and then "assigning" any remaining players to a locations at random.

Finally, from our discussion above, it happens that the total number of satisfied clients by strategy profile $\boldsymbol{s}$ corresponds to the value of the social welfare of $\boldsymbol{s}$. Thus the maximum social welfare any solution could ever achieve is $\sum_{i=1}^{\min(k,m)} \mathcal{N}_i$ as this is the maximum number of clients any solution can ever satisfy. This is exactly what we wanted to show. □

The result above implies that in an optimal solution, players will attempt to "cover" as many "large" components as possible by selecting at least one site in each component of the largest components. This maximizes the number of people that are being offered the service provided by the sellers.

Given the results shown above, we show a tight bound for the POA in our version of Hotelling games:

Theorem 3. *If $\boldsymbol{s}$ is a pure Nash equilibrium in a Hotelling game with graph G and k players, and $\boldsymbol{s}^*$ is a strategy profile with the maximum social welfare achievable for the same game, then we must have that*

$$\frac{\text{SW}(\boldsymbol{s}^*)}{\text{SW}(\boldsymbol{s})} \leq \frac{2k-1}{k}$$

In other words, the price of anarchy on Hotelling games is bounded above by $\frac{2k-1}{k}$.

Proof. Without loss of generality, let $C_1, C_2, ..., C_m$ be the distinct connected components of G ordered such that $\mathcal{N}_1 \geq \mathcal{N}_2 \geq \cdots \geq \mathcal{N}_m$. We start by noticing that because $\boldsymbol{s}$ is a PNE, by Corollary 1.1 there must exist a $1 \leq i^* \leq m$ such that

$$\text{SW}(\boldsymbol{s}) = \sum_{i=1}^{i^*} \mathcal{N}_i$$

Similarly we know from Theorem 2 that the social welfare of $\boldsymbol{s}^*$ is

$$\text{SW}(\boldsymbol{s}^*) = \sum_{i=1}^{\min(k,m)} \mathcal{N}_i$$

Thus we can obtain the following expression:

$$\frac{\text{SW}(\boldsymbol{s}^*)}{\text{SW}(\boldsymbol{s})} = \frac{\sum_{i=1}^{\min(k,m)} \mathcal{N}_i}{\sum_{i=1}^{i^*} \mathcal{N}_i}$$

Because $\boldsymbol{s}^*$ is an optimal solution, we must have that $i^* \leq \min(k, m)$ (as otherwise the social welfare of $\boldsymbol{s}$ would be greater than the social welfare of the optimal). This implies that we can rewrite the expression above as

$$\frac{\text{SW}(\boldsymbol{s}^*)}{\text{SW}(\boldsymbol{s})} = \frac{\sum_{i=1}^{i^*} \mathcal{N}_i + \sum_{i=i^*+1}^{\min(m,k)} \mathcal{N}_i}{\sum_{i=1}^{i^*} \mathcal{N}_i} = 1 + \frac{\sum_{i=i^*+1}^{\min(m,k)} \mathcal{N}_i}{\sum_{i=1}^{i^*} \mathcal{N}_i}$$

Again, using Corollary 1.1 we can rewrite the denominator to be a different way to express social welfare, the sum of every player's utility

$$\frac{\text{SW}(\boldsymbol{s}^*)}{\text{SW}(\boldsymbol{s})} = 1 + \frac{\sum_{i=i^*+1}^{\min(m,k)} \mathcal{N}_i}{\sum_{i=1}^{k} u_i(\boldsymbol{s})}$$

At this point we apply the what we know from Theorem 1 which is that

$$A(\boldsymbol{s}) = \left\{\mathcal{N}_i \mid \exists\, j \text{ s.t. } \boldsymbol{s}_j \in C_i\right\} = \left\{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_{i^*}\right\}$$

Consider the multiset of all of the component values $U = \{\mathcal{N}_i | 1 \leq i \leq m\}$, and the set difference $U/A(\boldsymbol{s}) = \{\mathcal{N}_i | i^* + 1 \leq i \leq m\}$. The values in the sum $\sum_{i=i^*+1}^{\min(m,k)} \mathcal{N}_i$ in the numerator are clearly elements of $U/A(\boldsymbol{s})$ because the indices of summation are $i > i^*$. Using this fact, we can upper bound our ratio with

$$\frac{\text{SW}(\boldsymbol{s}^*)}{\text{SW}(\boldsymbol{s})} \leq 1 + \frac{(\min(m, k) - i^*) \max_{x \in U/A(\boldsymbol{s})} x}{\sum_{i=1}^{k} u_i(\boldsymbol{s})}$$

Consider the least number of customers any seller receives in the Nash solution. This is exactly $h = \min_i u_i(\boldsymbol{s})$. We know that this value is at least as great as the number of customers in any component not picked by the Nash solution, $\max_{x \in U/A(\boldsymbol{s})} x$. To see why this is true, recall that $U/A(\boldsymbol{s})$ consists of the values of the components not chosen by Nash. If any of these values is greater than the minimum utility, then the player receiving the minimum utility would unilaterally deviate to a location on this untapped better component and receive all of the customers on that component. The player benefits from this deviation, which is a contradiction that $\boldsymbol{s}$ is Nash, so we conclude that every element in $U/A(\boldsymbol{s})$ (even the maximum) is less than or equal to the utility each seller receives at Nash (even the minimum). Using this minimum utility h, we simplify our inequalities as

$$\begin{aligned}
\frac{\text{SW}(\boldsymbol{s^*})}{\text{SW}(\boldsymbol{s})} &= 1 + \frac{(\min(m,k) - i^*)\max_{x \in U/A(\boldsymbol{s})} x}{\sum_{i=1}^{k} u_i(\boldsymbol{s})} \\
&\leq 1 + \frac{(\min(m,k) - i^*)h}{\sum_{i=1}^{k} u_i(\boldsymbol{s})} \\
&\leq 1 + \frac{(\min(m,k) - i^*)h}{kh} \\
&= 1 + \frac{\min(m,k) - i^*}{k} \\
&\leq 1 + \frac{k - i^*}{k} \\
&\leq 1 + \frac{k - 1}{k} \\
&= \frac{2k - 1}{k}
\end{aligned}$$

Where the last two inequalities follow from the facts that i^* has to be at least equal to 1 and $\min(k, m)$ is obviously less than or equal to k. This is the bound we wanted to show.

We furthermore show that this bound is tight. To see this consider the scenario where we have k players and k isolated nodes $\{1, 2, ..., k\}$ (i.e., there are a total of k connected components and location in G) such that $n_1 = k$ and $n_i = 1$ for all $i > 1$. The optimal solution in this case will allocate one seller to each location and get a social welfare of $k + (k - 1) = 2k - 1$. Now consider the solution $\boldsymbol{s}$ where every seller selects location 1. This is clearly a pure Nash equilibrium as every player gets utility 1 and changing to another location can give you at most utility 1. The social welfare of this solution is equal to k. Thus we will have that

$$\frac{\text{OPT}}{\text{SW}(\boldsymbol{s})} = \frac{2k - 1}{k}$$

Which is exactly the value of the given bound.

□

4.3 Discussion of Results

We see that our bound is bounded above by 2 in the asymptotic case when the number of players is very large. This is very similar to the result obtained in [8] when the graph was constrained to be connected. This is certainty interesting as our results used a completely different objective function than the one used in [8]. Furthermore, this result is independent of the weights used in the graph and is also independent of the metric used to determine how close two locations are (as long as such metric assigns a value of "infinity" to the distance between any two locations in distinct connected components of G). Similarly, this result will also hold if player i is only allowed to select a subset of locations $S_i \subseteq V$ as long as this subset contains at least one location in every connected component of G.

5 Future Directions

While our results are quite general in the sense that there are very few assumptions made on the topology of the underlying graph G and on the metrics used, there are still something we would like to explore as a continuation of this project. A few things that are worth looking into are

1. Could we find a similar POA bound when we use different objective functions like fairness or the summation of the distance traveled by clients as in [8] and [5], respectively ?
2. What can we say about the POA when players are only allowed to selecte their sites from an arbitrary subset of locations in V (not necessarily one from each connected component)?

3. What could we say about the POA when players are allowed to select multiple locations at the same time?
4. Under what conditions can we construct a PNE in our variation of Hotelling games? We started looking into this question and found an answer that depends on the existence of a PNE in Hotelling games where G is connected. Maybe looking deeper into this could give us some interesting results.

These are all interesting opportunities for future research in this area which could generalize the results in this paper even further.

References

[1] H Hotelling. Stability in competition economic journal vol. 39. 1929.

[2] Marios Mavronicolas, Burkhard Monien, Vicky G Papadopoulou, and Florian Schoppmann. Voronoi games on cycle graphs. In *International Symposium on Mathematical Foundations of Computer Science*, pages 503–514. Springer, 2008.

[3] Nicholas Economides. Hotelling's "main street" with more than two competitors. *Journal of Regional Science*, 33(3):303–319, 1993.

[4] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413. Springer, 1999.

[5] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 373–382, 2008.

[6] John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

[7] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20:1729–1736, 2007.

[8] Gaëtan Fournier and Marco Scarsini. Hotelling games on networks: efficiency of equilibria. 2014.

[9] Edward Hastings Chamberlin. *Theory of monopolistic competition: A re-orientation of the theory of value*. Oxford University Press, London, 1949.

[10] Abba P Lerner and Hans W Singer. Some notes on duopoly and spatial competition. *Journal of Political Economy*, 45(2):145–186, 1937.

[11] B Curtis Eaton and Richard G Lipsey. The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *The Review of Economic Studies*, 42(1):27–49, 1975.

[12] Amir Epstein, Michal Feldman, and Yishay Mansour. Strong equilibrium in cost sharing connection games. *Games and Economic Behavior*, 67(1):51–68, 2009.