

Forming Better Stable Solutions in Group Formation Games Inspired by Internet Exchange Points (IXPs)

Elliot Anshelevich, Wennan Zhu

Rensselaer Polytechnic Institute, Troy, NY 12180
eanshel@cs.rpi.edu, zhuw5@rpi.edu

Abstract

We study a coordination game motivated by the formation of Internet Exchange Points (IXPs), in which agents choose which facilities to join. Joining the same facility as other agents you communicate with has benefits, but different facilities have different costs for each agent. Thus, the players wish to join the same facilities as their “friends”, but this is balanced by them not wanting to pay the cost of joining a facility. We first show that the Price of Stability (PoS) of this game is at most 2, and more generally there always exists an α -approximate equilibrium with cost at most $\frac{2}{\alpha}$ of optimum. We then focus on how better stable solutions can be formed. If we allow agents to pay their neighbors to prevent them from deviating (i.e., a player i voluntarily pays another player j so that j joins the same facility), then we provide a payment scheme which stabilizes the solution with minimum social cost s^* , i.e. PoS is 1. In our main technical result, we consider how much a central coordinator would have to pay the players in order to form good stable solutions. Let Δ denote the total amount of payments needed to be paid to the players in order to stabilize s^* , i.e., these are payments that a player would lose if they changed their strategy from the one in s^* . We prove that there is a tradeoff between Δ and the Price of Stability: $\frac{\Delta}{cost(s^*)} \leq 1 - \frac{2}{5} PoS$. Thus when there are no good stable solutions, only a small amount of extra payment is needed to stabilize s^* ; and when good stable solutions already exist (i.e., PoS is small), then we should be happy with those solutions instead. Finally, we consider the computational complexity of finding the optimum solution s^* , and design a polynomial time $O(\log n)$ approximation algorithm for this problem.

Introduction

We study a coordination game motivated by the formation of Internet Exchange Points (IXPs). In this game, there are m facilities available, and the players (modeling ISPs, or more generally entities which wish to exchange traffic with each other) choose which facilities to join. Joining a facility f_k has a cost for player i , which we call the “connection cost” and denote by $w(i, f_k)$; this cost can be different for different players and facilities. The reason why players are willing to pay such costs is because joining the same facility as other players is beneficial: a pair of players i and j which do not

connect to the same facility must pay a cost $w(i, j)$, but if they share a facility then this cost disappears. Finally, the facilities themselves have costs $c(f_k)$ which must be paid for by the players using these facilities. In summary, the players wish to join the same facilities as their “friends” in order to avoid paying the costs $w(i, j)$, but this is counterbalanced by them not wanting to pay the cost of joining a facility.

While our game is quite general, and models general group formation (e.g., facilities are clubs or groups people can join, and they wish to join the same clubs as their friends), this game is specifically inspired by the formation of IXPs in the Internet. IXPs are facilities where Internet Service Providers (ISPs) can exchange Internet traffic with high speed; a large fraction of total Internet traffic flows through such hubs (Ager et al. 2012). If two ISPs join the same IXP (and pay their cost for joining, which can depend on many factors including the pricing scheme and the physical location(s) of the IXP), then they gain the benefit of mutual high speed communication. If, however, two ISPs do not use the same IXP, they must use alternate means of exchanging traffic with each other (e.g., through their providers or private peering), which we model by them incurring an extra cost $w(i, j)$.

Coordination games have been widely studied in various situations where agents gain utility by forming coalitions with other agents. Even with the large amount of existing work on both coordination games and group formation, the questions we consider in this paper have not been studied before for our game (see Related Work). Like many such games, ours can be represented by a graph, in which each node stands for a player and the edges between them have weights representing the disconnection cost for them not belonging to the same facility. One major difference between our game and much (although certainly not all) of existing work is that the facilities (i.e., groups that players can join) are not identical: their quality for a player i depends not only on who else has joined the same group (as in hedonic games (Aziz and Savani 2016)), but also on the specific facility being joined, as quantified by the cost $w(i, f_k)$. This immediately changes a lot about equilibrium structure: it is no longer the case that everyone being in the same group is an equilibrium solution which minimizes social cost; instead equilibrium solutions involve players balancing their cost for joining facilities with their cost of being separated from their

friends. Other coordination games look at cases where only a limited number of facilities can be open, or when players have both “friends” and “enemies” (i.e., $w(i, j)$ can be negative); for the types of settings we consider, however, all facilities can open as long as players are able to pay for them, and there is never any additional cost from two players joining the same facility (i.e., $w(i, j) \geq 0$). Moreover, unlike most other coordination games, we assume that facilities have a cost which must be shared among the players using it, which adds a significant layer of complexity to our results (for example, our game is no longer a potential game (Monderer and Shapley 1996)). For more details and comparison with existing work, see the Related Work section.

Our Contributions

In this paper, we study a coordination game where a strategy of an agent i is to choose a facility f_k to join, by paying a connection cost $w(i, f_k)$ (or to not join any facility). If two agents i, j do not use the same facility, then both of them are charged a disconnection cost $w(i, j)$. In addition, there is a fixed facility cost $c(f_k)$ for each open facility, which is split among all agents using f_k according to an arbitrary pricing rule. An agent’s total social cost is the sum of its connection cost, disconnection cost, and its share of the facility cost. An assignment with a pricing rule is stable if it is budget balanced (each $c(f_k)$ is fully paid by all agents using f_k), and no agent wants to switch facilities, i.e., it is a Nash equilibrium.

We study the quality of equilibrium solutions for this game, as well as ways to create new stable solutions. We first show that while the Price of Anarchy can be arbitrarily high, the Price of Stability (PoS) is at most 2, and more generally there always exists an α -approximate equilibrium with cost at most $\frac{2}{\alpha}$ of optimum. While we use potential arguments to prove this (Tardos and Wexler 2007), note that this game is not a potential game due to facility costs, and thus new proof techniques are needed beyond simply defining a potential function. We then focus on how better stable solutions can be formed. If we allow agents to pay their neighbors to prevent them from deviating (i.e., a player i voluntarily pays another player j so that j joins the same facility), then we provide a payment scheme which stabilizes the solution with minimum social cost s^* , i.e. PoS is 1. This is essentially what occurs, for example, in paid peering (Shrimali and Kumar 2006), where two ISPs have different incentives, and so one ISP pays the other in order to form a peering connection. Finally, for our main result, we consider how much a central coordinator would have to pay the players in order to form good stable solutions, similarly to (Anshelevich and Sekar 2014; Bachrach et al. 2009). Let Δ denote the total amount of payments needed to be paid to the players in order to stabilize s^* , i.e., these are payments that a player would lose if they changed their strategy from the one in s^* . We prove that there is a tradeoff between Δ and the Price of Stability: $\frac{\Delta}{cost(s^*)} \leq 1 - \frac{2}{5}PoS$. Thus when there are no good stable solutions, only a small amount of extra payment is needed to stabilize s^* ; and when good stable solutions already exist (i.e., PoS is small), then we should be happy with those solutions instead! This result is proven by

forming several solutions where specific subsets of players perform their best responses, and then showing that when a small amount of payment is not enough to stabilize s^* , then at least one of these solutions is guaranteed to be better than s^* , giving a contradiction. The difficulty here results from the fact that letting any single player move to their best response strategy from s^* could still result in solutions worse than s^* ; to get a contradiction and form a solution strictly better than s^* requires changing the strategy of many players simultaneously.

The results above are for the setting where each agent can join at most one facility at a time. In Section , we study the setting where each agent is allowed to use multiple facilities simultaneously. Many of the results above still hold for this general mode, but only under the assumption that a player can only switch their strategy by leaving one facility at a time (although it is allowed to join multiple new facilities at once).

Finally, we consider the computational complexity of finding the optimum solution s^* . We prove that computing it is NP-Hard (and in fact inapproximable to better than $\Omega(\log n)$ unless $P=NP$), and design a polynomial time approximation algorithm that gives a $\min\{m + 1, O(\log n)\}$ -approximation to the optimal solution (with n being the number of players, and m the number of facilities). We also provide a simple 2-approximation algorithm when all facility costs are zero.

Related Work

There is a very large amount of work on both group formation and coordination games, which is too large to survey here. Hedonic games (Aziz and Savani 2016; Dreze and Greenberg 1980) is an important class of games related to coordination games, in which the agents form groups, and each agent’s utility only depends on the other agents in its own group, but is not affected by how agents are arranged in other groups. The objectives are usually maximizing social welfare (Apt et al. 2014; Aziz et al. 2019; Aziz, Brandt, and Seedig 2013; Brânzei and Larson 2009; Feldman, Lewin-Eytan, and Naor 2015; Gairing and Savani 2010) or minimizing social cost (Feldman, Lewin-Eytan, and Naor 2015). Often, although not always, all players in a group have the same cost or utility. In much of the work, the number of groups is fixed (Bhalgat, Chakraborty, and Khanna 2010; Feldman, Lewin-Eytan, and Naor 2015; Gourvès and Monnot 2010; Hofer 2007). There are also various utility/cost functions which have been studied, with the most common one being that an agent’s utility is the total utility gained from being with all other agents in its group. In fractional hedonic games (Aziz et al. 2019; Bilò et al. 2014, 2015), an agent’s utility is the average value of its presence to every other agent in the group. More generally, there are also other types of related group formation games, e.g., congestion games (Christodoulou and Koutsoupias 2005; Rosenthal 1973) and profit sharing games (see (Augustine et al. 2011) and references therein), where an agent’s utility only depends on the size of the group.

While coordination games can be considered a special case of general hedonic games, usually coordination games

involve players with some sort of graph structure, where for a pair of players, being in the same group gives them both a benefit if they are “friends” (or a penalty if they are “enemies”). This is in contrast to many hedonic games, where all players in a group have the same utility, or the total utility of a group is somehow shared among its participants. In most related work, either the objective functions of the players are very different from ours (e.g., they depend on the number of players in their group) (Apt et al. 2014; Aziz et al. 2019; Brânzei and Larson 2011), or there are players who specifically don’t want to be in the same group (“enemies”, negative-weight edges) (Auletta et al. 2016b; Bhalgat, Chakraborty, and Khanna 2010; Feldman and Friedler 2015), or all groups are identical and the optimum solution would correspond to either everyone joining the same group or everyone forming a group on their own (Apt et al. 2014; Aziz, Brandt, and Seedig 2013; Bhalgat, Chakraborty, and Khanna 2010; Brânzei and Larson 2009; Feldman, Lewin-Eytan, and Naor 2015; Feldman and Friedler 2015). In contrast, our work is motivated by settings where everyone would like to form one group together to reduce the disconnection cost, but the complexity in the solution structure comes from the players trading this desire off with their individual connection costs to (non-identical) facilities.

As discussed in the Introduction, general coordination games include other settings in which an agent’s utility or cost also depends on which group it joins (i.e., the groups are not identical). Our work is more closely related to this type of game. Using a graph representation, one can think of such games as either Max-Uncut (maximize the weight of edges to friends in your group) or Min-Cut (minimize the weight of edges to friends not in your group) objectives, but with additional utility or cost depending on which group a player joins (which can be modeled using additional “anchor nodes” which *must* belong to a specific group, see e.g., (Anshelevich and Sekar 2014)). Work on such coordination games with non-identical groups or facilities includes (Anshelevich and Sekar 2014; Auletta et al. 2017; Chierichetti, Kleinberg, and Oren 2013). In k -Coloring games (Carosi and Monaco 2019) each agent gains utility by choosing a certain color/facility, and loses utility by choosing the same color as other adjacent agents, i.e., they are anti-coordination games in which all agents want to be in different groups if possible (see references in (Carosi and Monaco 2019) for more discussion of such games). In generalized Discrete Preference Games (Auletta et al. 2016b), there are exactly two groups, and the players could be friends or enemies. Similar to hedonic games, the research in this area usually focused on properties of stable solutions, e.g., (Auletta et al. 2016a) studies how a single agent could affect the Nash Equilibria converged from best responses, and (Auletta et al. 2017) compares the prices of anarchy and stability under different objective functions.

Perhaps the most related work to ours is (Chierichetti, Kleinberg, and Oren 2013), as it is also a Min-Cut game with non-identical groups. The main differences between our work and (Chierichetti, Kleinberg, and Oren 2013) are: it is assumed in (Chierichetti, Kleinberg, and Oren 2013) that every agent has a favorite group, and an agent’s cost

depends on the distance between its current group and favorite group, and the distances to its neighbors. We do not bind each agent’s cost with a group in our setting. Suppose two agents have the same “favorite group” f_k , which is the group with the lowest connection cost to them; in our model their cost to any other group f'_k could be very different. (Chierichetti, Kleinberg, and Oren 2013) also focuses on the setting where the group locations form a general metric or tree metric, while we do not have such assumptions. Last but not least, unlike in the works mentioned above, we assume there is a facility cost to open each facility, with different facilities having different costs. We study stable states where the facility cost is split among the agents using it, so each facility is paid for, and each agent is stable with three types of costs: connection cost to facility, its own share of facility cost, and disconnection cost to its neighbors that use different facilities. We also study the case that each agent can join multiple groups. Finally, parts of our work are also closely related to (Anshelevich and Sekar 2014), which shows that an optimal solution could be stabilized by providing a reasonable amount of payments to the agents, just as we do. Their model, however, involves maximizing utility instead of minimizing costs (which changes the equilibrium structure and all approximation factors like PoS and cost of stabilization entirely), and does not include any facility costs.

Model and Preliminaries

We are given a set of m facilities $\mathcal{F} = \{f_1, f_2, \dots, f_m\}$ and a set of n agents (which we will also call “players”) $\mathcal{A} = \{1, 2, \dots, n\}$. An agent i can use any facility f_k by paying a connection cost $w(i, f_k)$. A pair of agents (i, j) can form connections through facility f_k if they are both using f_k . However, if i and j do not use the same facility, then both of them are charged a disconnection cost $w(i, j)$. A facility f_k is open if and only if there exists an agent using it. There is a fixed facility cost of $c(f_k) \geq 0$ for any open facility f_k .

In much of this paper, we assume each agent uses at most one facility, so the strategy set of an agent consists of \mathcal{F} together with the empty set. A facility assignment $s = \{s_1, s_2, \dots, s_n\}$ denotes the facilities that each agent uses: s_i denotes the facility that agent i uses in assignment s . In the case that agent i does not use any facility, let $s_i = \emptyset$ and $w(i, s_i) = 0$. A pricing strategy $\gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ assigns the price for using each facility f_k to every agent i . $\gamma_i(f_k)$ is a non-negative number that denotes agent i ’s share of the facility cost for using f_k . (s, γ) is a state with assignment s and pricing strategy γ . Note that agent i only pays its share of the facility cost to f_k if i uses f_k , i.e., $\gamma_i(f_k) > 0$ only if $s_i = f_k$.

To summarize, the total cost of agent i in a state (s, γ) is the sum of the following three parts:

1. If i uses facility f_k , then there is a connection cost $w(i, f_k)$ to i .
2. For each agent j that do not use s_i , i.e., $s_i \neq s_j$, there is a disconnection cost $w(i, j)$ to both i and j . A special case is when both i and j are not using any facility: although $s_i = \emptyset$ and $s_j = \emptyset$, we still say that $s_i \neq s_j$ in this case

to make it consistent that the disconnection cost $w(i, j)$ is charged to both i and j if $s_i \neq s_j$.

3. If i uses facility f_k , then there is a facility cost $\gamma_i(f_k)$ to i .

We denote the total cost of agent i as $c_i(s, \gamma)$. Summing up the three types of cost mentioned above:

$$c_i(s, \gamma) = w(i, s_i) + \sum_{j|s_i \neq s_j} w(i, j) + \gamma_i(s_i)$$

For convenience, we denote the cost of agent i without facility cost as $\tilde{c}_i(s)$:

$$\tilde{c}_i(s) = w(i, s_i) + \sum_{j|s_i \neq s_j} w(i, j)$$

In this paper, we are interested in the social cost of stable states. The total social cost of a state (s, γ) equals the sum of $c_i(s, \gamma)$, plus the total cost of all open facilities. For each facility f_k , the cost is $c(f_k)$ minus the sum of $\gamma_i(f_k)$ of each agent using f_k , i.e., $c(f_k) - \sum_{i|s_i=f_k} \gamma_i(f_k)$. In other words, one can think of each facility as an agent with cost $c(f_k)$, and with other agents paying it the prices $\gamma_i(f_k)$ for using it. The sum of $\gamma_i(f_k)$ cancels out, and the total social cost is actually the sum of $\tilde{c}_i(s)$ plus the sum of $c(f_k)$ of open facilities:

$$\begin{aligned} c(s) &= \sum_{f_k \in \mathcal{F}, f_k \text{ is open}} c(f_k) + \sum_{i \in \mathcal{A}} (w(i, s_i) + \sum_{j|s_i \neq s_j} w(i, j)) \\ &= \sum_{f_k \in \mathcal{F}, f_k \text{ is open}} c(f_k) + \sum_{i \in \mathcal{A}} w(i, s_i) + \sum_{i \in \mathcal{A}} \sum_{j|s_i \neq s_j} w(i, j) \\ &= \sum_{f_k \in \mathcal{F}, f_k \text{ is open}} c(f_k) + \sum_{i \in \mathcal{A}} w(i, s_i) + 2 \sum_{(i, j)|s_i \neq s_j} w(i, j) \end{aligned}$$

We consider (i, j) as an unordered pair, therefore in $\sum_{i \in \mathcal{A}} \sum_{j|s_i \neq s_j} w(i, j)$, each unordered pair (i, j) that $s_i \neq s_j$ is counted twice.

In this paper, we study the game in which each agent's goal is to minimize its total social cost, and the central coordinator's goal is to find a budget balanced and stable state (s, γ) that (approximately) minimizes the total social cost. A state is budget balanced if each facility f_k is fully paid with the facility cost $c(f_k)$, formally:

Definition 1. A state (s, γ) is **budget balanced** if for each facility f_k , $\sum_{i|s_i=f_k} \gamma_i(f_k) = c(f_k)$.

Before defining the stability of a state, we first define an agent's best response. Consider an agent i with current strategy $s_i = f_k$, and price $\gamma_i(f_k)$ for using this facility. The agent may consider switching to a different facility f_ℓ , but to correctly evaluate their cost after this switch, the agent needs to know exactly how much they will pay after such a switch. We assume that the agents know their connection costs $w(i, f_\ell)$ and their disconnection costs from other agents, as well as which agents are using each facility. What price $\gamma_i(f_\ell)$, however, should they anticipate after switching

to their new facility? If the prices depend on the set of agents (or the number of agents) at the facility, then the price might change from the current one being offered. But how reasonable is it for agents to know the exact details of the pricing schemes used by the facilities (which are modeling IXP's or other private enterprises which do not want to reveal their pricing structures)?

To address these issues, in this paper, every agent assumes it will be charged 0 facility cost for joining a new facility. This allows us to not worry about what an agent may know and what price they may anticipate after switching a facility. At the same time, this assumption does not limit our results on stable solutions. This is because no matter what price $\gamma_i(f_\ell)$ an agent may anticipate after switching to facility f_ℓ , anticipating a price of 0 instead will make it only more likely to switch. Thus, no matter what the agents' beliefs for prices after switching make sense for a particular setting, *a stable solution in our model will still be stable no matter what beliefs about prices $\gamma_i(f_\ell)$ the agents hold, or what price they will actually be charged after switching.*

Thus our results about stable solutions are stronger: they state that even if the agents are extremely optimistic and believe they can switch to any facility without paying facility cost, then there still exist good stable solutions. If they assumed costs higher than 0, then the set of stable solutions would only increase. In other words, if an agent is stable when assuming it will be charged 0 for joining other facilities, then it would also be stable with a higher cost as well.

Definition 2. Given a state (s, γ) , s'_i is agent i 's **best response** if $\forall s''_i \neq s'_i, \tilde{c}_i(s'_i, s_{-i}) + \hat{\gamma}_i(s'_i) \leq \tilde{c}_i(s''_i, s_{-i})$, where $\hat{\gamma}_i(s'_i) = \gamma_i(s_i)$ if $s'_i = s_i$, and $\hat{\gamma}_i(s'_i) = 0$ otherwise. We denote i 's best response at state (s, γ) as $BR_i(s, \gamma)$.

In the definition above, $\hat{\gamma}_i$ is the pricing strategy that agent i assumes would happen after its deviation. If agent i stays at its current facility, then its share of the facility cost does not change. But if i leaves its current facility and joins another one, then it believes that it will be charged 0 facility cost for joining the new facility.

Definition 3. Agent i is stable at state (s, γ) if for any strategy $s'_i \neq s_i$:

$$c_i(s, \gamma) \leq \tilde{c}_i(s'_i, s_{-i}).$$

In other words, agent i is stable at (s, γ) if s_i is i 's best response at (s, γ) .

We define a state (s, γ) to be stable if it is budget balanced, and every agent is stable. Intuitively, if a state is not budget balanced, then a facility would not cover its operating cost $c(f_k)$, and thus would not choose to remain open.

Definition 4. A state (s, γ) is **stable** if it is budget balanced, and for each agent i , for any strategy $s'_i \neq s_i$:

$$c_i(s, \gamma) \leq \tilde{c}_i(s'_i, s_{-i}).$$

Denote an assignment with the minimum total social cost as s^* . Our goal is to find stable states (s, γ) to approximate the minimum total social cost. We use Price of Stability (PoS) to quantify the quality of a stable state. Given an instance, suppose (\hat{s}, γ) is the stable state with the smallest total social cost, then PoS is the worst case ratio between $c(\hat{s})$

and $c(s^*)$ for any instance. A related concept, Price of Anarchy (PoA) is defined as: suppose (\hat{s}, γ) is the stable state with the largest total social cost, then PoA is the worst case ratio between $c(\hat{s})$ and $c(s^*)$ for any instance. So PoS shows the quality of the best stable state, while PoA shows the quality of the worst stable state.

Pricing Strategies and Stability

Recall a state (s, γ) is stable if it is budget balanced, and every agent i is stable. Suppose there is no other constraint on the pricing strategy, then we ask the following question in order to find a budget balanced state: in an assignment s , how much facility cost can we charge an agent while keeping it stable? To answer this question, we first define a special type of best response: with an assignment s , let $\overline{BR}_i(s)$ denote i 's best response, given i is forced to stop using s_i . In other words, it is the strategy $s'_i \neq s_i$ with the smallest $\tilde{c}_i(s'_i, s_{-i})$. Note that if $BR_i(s, \gamma) \neq s_i$, i.e., if i wants to switch from the state (s, γ) , then $BR_i(s, \gamma) = \overline{BR}_i(s)$. But in the case when i 's best response is to stay at its current strategy, $\overline{BR}_i(s)$ would denote the ‘‘next best choice’’ if i is forced to stop using its current facility. Intuitively, the ‘‘value’’ of facility s_i to agent i is how much i 's cost would increase if i is forced to leave s_i and join the next best choice $\overline{BR}_i(s)$. If there are multiple strategies that all satisfy the definition of $\overline{BR}_i(s)$, then we choose an arbitrary one except in one case: we never choose a facility that is closed in s as $\overline{BR}_i(s)$. We can always do this because if there exists such strategy s'_i , such that s'_i is a closed facility in s , then compare $\tilde{c}_i(\emptyset, s_{-i})$ with $\tilde{c}_i(s'_i, s_{-i})$. The connection cost in $\tilde{c}_i(\emptyset, s_{-i})$ is 0, and the disconnection cost is the same as in $\tilde{c}_i(s'_i, s_{-i})$, because i would be the only agent using s'_i . So it must be $\tilde{c}_i(\emptyset, s_{-i}) \leq \tilde{c}_i(s'_i, s_{-i})$, and we define $\overline{BR}_i(s) = \emptyset$ in this case. For every agent i in assignment s , define $Q_i(s) = \tilde{c}_i(\overline{BR}_i(s), s_{-i}) - \tilde{c}_i(s)$; it is not hard to see that agents would be willing to pay this price in order to use facility s_i .

Note that some agents might be unstable even with 0 facility cost, so we also consider the case that agents need to receive payments to be stable. Let Δ_i denote a payment that agent i receives if it does not deviate at state (s, γ) , and denote the total payments as $\Delta = \sum_i \Delta_i$. In this paper, the default setting is that agents do not receive payments ($\Delta_i = 0$), but we do consider the cases that agents are allowed to be paid by a central coordinator in Section or paid by their neighbors in Section . Then we define the stability with payments as follows: a state (s, γ) with payments Δ is stable if it is budget balanced, and for each agent i , for any strategy $s'_i \neq s_i$:

$$c_i(s, \gamma) - \Delta_i \leq \tilde{c}_i(s'_i, s_{-i})$$

The following lemma shows that agent i is stable if $\gamma_i(s_i) - \Delta_i \leq Q_i(s)$. Proofs for all our results can be found in the full version of this paper at <https://arxiv.org/abs/2008.12235>.

Lemma 1. *Given any assignment s , pricing strategy γ , and payments to agents Δ , agent i is stable if $\gamma_i(s_i) - \Delta_i \leq Q_i(s)$.*

Single Facility per Agent: Price of Stability

In the first part of this paper, we show our results in the setting that each agent uses at most one facility.

Facility Cost $c(f_k) = 0$ for Every f_k

In this section, we provide simple baseline results for the case that there is no facility cost. Set the pricing strategy to be $\gamma_i(f_k) = 0$ for any agent i and facility f_k , so all solutions are budget balanced. In this special case, for any agent i in assignment s , we have $c_i(s, \gamma) = \tilde{c}_i(s)$. A state s is stable if for each agent i and strategy s'_i , $\tilde{c}_i(s) \leq \tilde{c}_i(s'_i, s_{-i})$.

Define potential function $\tilde{\Phi}(s)$ as:

$$\tilde{\Phi}(s) = \sum_{i \in \mathcal{A}} w(i, s_i) + \sum_{(i,j) | s_i \neq s_j} w(i, j)$$

When an agent i switches its strategy from s_i to s'_i , it is easy to see that the change of i 's cost is captured exactly by the change of $\tilde{\Phi}(s)$, so $\tilde{\Phi}(s)$ is an exact potential function. We can prove the following theorem using standard potential methods (Tardos and Wexler 2007).

Theorem 1. *If $\forall k, c(f_k) = 0$, then price of stability is at most 2 and this bound is tight.*

Theorem 2. *The price of anarchy is unbounded in our setting.*

See the full version on arXiv for the proofs.

Price of Stability for Arbitrary Facility Costs

In this section, we consider the case that for each f_k , the facility cost $c(f_k)$ is a fixed constant when f_k is open, regardless of how many agents/connections are using f_k . We suppose there is a central coordinator to determine the pricing strategy γ that is budget balanced, with no other constraint on γ .

Note that $\tilde{\Phi}(s)$ is not a potential function in this setting anymore, because agent i also considers the facility cost $\gamma_i(s_i)$ when it deviates to decrease $c_i(s)$. Thus, $\tilde{\Phi}(s)$ does not always decrease when i deviates. We define another potential function $\Phi(s)$:

$$\begin{aligned} \Phi(s) &= \sum_{f_k | f_k \text{ is open}} c(f_k) + \sum_{i \in \mathcal{A}} w(i, s_i) + \sum_{(i,j) | s_i \neq s_j} w(i, j) \\ &= \sum_{f_k | f_k \text{ is open}} c(f_k) + \tilde{\Phi}(s) \end{aligned}$$

We cannot use the potential method (Tardos and Wexler 2007) to analyze the price of stability in our game directly; in fact our game is not a potential game. For this new potential function $\Phi(s)$, a player could still deviate to lower its cost, while the potential increases. The following lemma, however, shows that when a player deviates and decreases its cost $\tilde{c}_i(s)$ (but not necessarily decreases cost $c_i(s, \gamma)$), then $\Phi(s)$ does in fact decrease.

Lemma 2. *In an assignment s , if any agent i switches its strategy to s'_i such that $\tilde{c}_i(s'_i, s_{-i}) < \tilde{c}_i(s)$ and s'_i does not contain any closed facility in s , then $\Phi(s'_i, s_{-i}) < \Phi(s)$.*

Now we use the above potential to prove bounds on the price of stability. While a single player changing its strategy to decrease its cost might actually increase the value of the potential $\Phi(s)$, we give a series of coalitional deviations (i.e., groups of players switching strategies simultaneously) so that the potential is guaranteed to decrease after each such deviation, and so that the cost of the resulting stable solution is not too large.

Theorem 3. *The price of stability is at most 2, and this bound is tight. In other words, there exists a stable state (s, γ) with cost at most twice that of optimum.*

Proof Sketch. We define a coalitional deviation process that converges to a stable state, with $\Phi(s)$ decreasing in each step of the process. Start with the optimal assignment s^* . If there exists an agent i such that when i switches to a strategy s'_i , in which s'_i does not contain any closed facility in s^* , then $\tilde{c}_i(s'_i, s^*_{-i}) < \tilde{c}_i(s^*, s^*_{-i})$, then let agent i switch to s'_i . Repeat this process until no such agent exists. In other words, each agent is now stable if they assume they are not charged any facility cost. By Lemma 2, $\Phi(s)$ decreases during each step in this process. Let s be the current state.

For each facility f_k , consider the following two cases:

Case 1, $c(f_k) > \sum_{i|s_i=f_k} Q_i(s)$. In this case, we close f_k and let each agent i using f_k in s switch its strategy $\overline{BR}_i(s)$. We can show that $\Phi(s)$ decreases after f_k is closed (see the full version).

Repeat the above two steps: let agents switch strategies to reach a “stable” state s . Then if there exist a facility f_k that satisfies the condition in **Case 1**, we close f_k and let every agent i using it switch its strategy to $\overline{BR}_i(s)$. $\Phi(s)$ decreases in each step, so this process always converges to such an assignment s . Then each open facility must satisfy the following **Case 2**:

Case 2, $c(f_k) \leq \sum_{i|s_i=f_k} Q_i(s)$. In this case, we set $\gamma_i(f_k) = Q_i(s)$ for each agent i . We show that this state (s, γ) is stable in the full version.

Because $\Phi(s)$ decreases in each step of the process, so we can directly get the conclusion that PoS is at most 2 by comparing $\Phi(s)$ with the total social cost. \square

The above price of stability result can be easily generalized to approximately stable solutions as well. We say a state (s, γ) is α -approximate stable if it is budget balanced, and no agent could deviate to lower its cost to $\frac{1}{\alpha}$ of its current cost:

Definition 5. *A state (s, γ) is α -approximate stable if it is budget balanced, and for each agent i , for any strategy $s'_i \neq s_i$:*

$$c_i(s, \gamma) \leq \alpha \cdot \tilde{c}_i(s'_i, s_{-i})$$

Theorem 4. *There always exists an α -approximate stable state (\hat{s}, γ) such that $\frac{c(\hat{s})}{c(s^*)} \leq \frac{2}{\alpha}$.*

This theorem implies that, in particular, the optimum solution s^* is a 2-approximate stable state, i.e., no player can improve their cost by more than a factor of 2 by switching its facility.

Payments to Form Good Stable Solutions

Agents Paying Each Other: “Paid Peering”

In this section, we consider the case that agents can pay each other to stabilize the optimal assignment. Formally, for a pair (i, j) such that $s_i = s_j = f_k$, i can pay j up to $w(i, j)$ in order to discourage j from leaving facility f_k and thus disconnecting from i . Given the asymmetry of agent costs (due to connection costs), it may make sense for agents to give their “friends” extra incentives to connect with them using a particular facility. Of course, agent i would never voluntarily pay agent j more than j ’s value to i , i.e., more than $w(i, j)$. Such payments make sense in general settings of group formation, and make sense in our motivating IXP setting as well: when two ISP’s decide to make a peering arrangement to exchange traffic after joining a common IXP, it is often the case that they make a *paid* peering contract (Shrimali and Kumar 2006), in which one ISP pays the other for the traffic exchange, thus giving it extra incentive to remain connected to their joint facility.

Let p_{ij} denote the payments that agent i pays its neighbor j to discourage it from leaving the facility they share. $p_{ij} \geq 0$ means i pays j , and $p_{ij} < 0$ means i receives payment from j . For any pair of agents (i, j) , we have $p_{ji} = -p_{ij}$. In this section, Δ_i denotes the total payments that agent i receives from its neighbors minus the total payments i pays its neighbors. In other words, $\Delta_i = \sum_{j|s_i=s_j} p_{ji}$. We abuse the notation to allow Δ_i to be negative, in which case i pays more than receives from its neighbors. We consider stability with payments defined in Section . It is easy to see that Lemma 1 still holds with this modified definition of Δ_i .

In the optimal assignment s^* with a pricing strategy γ , consider the stability of every agent using f_k : by Lemma 1, we know every agent i would be stable if $\gamma_i(s_i^*) - \Delta_i \leq Q_i(s^*)$. For a pair of agents (i, j) using f_k in s^* , suppose $Q_i(s^*) \geq 0$, and $Q_j(s^*) < 0$, which means we can get some payments from i (to pay the facility or its neighbors) while keeping it stable, but j needs to be paid to become stable at s^* . Thus, it makes sense for i to pay j to stop it from deviating, but i would not pay more than $w(i, j)$, which is the maximum increase of i ’s cost as a result of j ’s deviation.

Theorem 5. *If we allow agents to pay their neighbors, and i pays j no more than $w(i, j)$, then there exist γ and payments of players to each other so that the resulting solution (s^*, γ) is stable, with s^* being the solution minimizing social cost. In other words, the price of stability becomes 1.*

Paying Agents Directly to Stabilize s^*

In this section, we take on the role of a central coordinator, who is paying the agents in order to stabilize the optimum solution s^* . We study the relationship between the Price of Stability and the minimum total payments required to stabilize s^* . We use the notation of Δ_i and stability with payments defined in Section . Δ_i represent the payment each agent i receives from the central coordinator, and the total payments are $\Delta = \sum_i \Delta_i$.

In our main technical result below, we show that the total payment Δ required to stabilize s^* is only a fraction of the social cost of the optimal solution. In fact, there is a tradeoff

between Δ and PoS : when PoS is large, e.g., $PoS = 2$, we only need to pay $\frac{1}{5}c(s^*)$ to stabilize s^* , which is a reasonably small fraction of $c(s^*)$. Thus when PoS is small, there already exist good stable solutions by definition of PoS , and when PoS is large, only a relatively small amount of payments are necessary to stabilize s^* .

Theorem 6. *For any instance, $\frac{\Delta}{c(s^*)} \leq 1 - \frac{2}{5}Pos$, where Δ is the payment needed to stabilize s^* .*

Proof Sketch. Let b_i denote the strategy of agent i such that $\tilde{c}_i(b_i, s_{-i}^*)$ is minimized. We first show that if there exists a state s , such that $\sum_{i \in \mathcal{A}} \tilde{c}_i(b_i, s_{-i}^*) \geq \frac{4}{5}\tilde{\Phi}(s)$, then $\frac{\Delta}{c(s^*)} \leq 1 - \frac{1}{5}Pos$. Therefore, we only need to find an assignment s that satisfies this condition. We define several assignments s^0, s^1, s^2 as candidates that may satisfy this condition, and then prove that at least one of them must do so for every instance. s^0 is the assignment such that every agent i switches its strategy to b_i . Consider all the connection and disconnection costs in the potential function $\tilde{\Phi}(s^0)$. All the connection costs are in the sum of $\tilde{c}_i(b_i, s_{-i}^*)$. As for the disconnection cost between any pair of agents (i, j) , if it is in $\tilde{\Phi}(s^0)$, then it must be in $\tilde{c}_i(b_i, s_{-i}^*)$ or $\tilde{c}_j(b_j, s_{-j}^*)$ except for one situation: if $b_i = s_j^*, b_j = s_i^*$, but $b_i \neq b_j$. This is the only case in which the disconnection cost $w(i, j)$ occurs in $\tilde{\Phi}(s^0)$, but not in either of $\tilde{c}_i(b_i, s_{-i}^*)$ or $\tilde{c}_j(b_j, s_{-j}^*)$. To handle this case, we carefully decompose all the agents into two sets \mathcal{A}_1 and \mathcal{A}_2 , so that this special type of disconnection costs only occur between agents belonging to different sets. Then we define s^1 and s^2 , each representing the assignment such that only one set of agents switches their strategies to b_i , and the other set stays at s^* . Finally, we show by contradiction that at least one of s^0, s^1 , and s^2 must satisfy the condition above, which finishes the proof of the theorem.

Multiple Facilities per Agent

We also consider the case that each agent is allowed to use multiple facilities. All of our results, with the exception of Theorem 6, still hold in this setting, but with a constraint on possible deviations: when agents switch their strategies, they are only allowed to drop from at most one facility each time, although they can join as many new facilities as they want to. The proofs for agents being allowed to connect to multiple facilities are almost the same as for the setting where they can only connect to a single facility. We include the results and proofs in the full version for completeness.

Note that, as is common when studying PoS and PoA, in this paper we consider stable solutions with respect to unilateral deviations, i.e., a single agent changing their strategy. Looking at coalitional stability concepts is an interesting direction of future research, although it is not difficult to prove that Strong equilibrium (SE) may not exist for our setting, while an approximate SE always does.

Computation of Optimum Solutions

In this section, we discuss approximation algorithms to calculate the optimal assignment in polynomial time.

Theorem 7. *If $\forall k, c(f_k) = 0$ and each agent is only allowed to use one facility, then computing the optimum solution s^* is NP-Hard, but there exists a poly-time 2-approximation algorithm.*

When some facility costs are not 0, it is easy to show that this problem is inapproximable to better than $\Omega(\log n)$ unless $P=NP$ (Hochbaum 1982). We now show that even if each agent is allowed to use multiple facilities with non-zero costs, there exist polynomial time algorithms that give a $\min\{O(\log n), m+1\}$ -approximation with high probability. First note that the following is an LP-relaxation for our problem (see the full version):

$$\begin{aligned} \min \quad & \sum_{i,k} w(i, f_k)x_{ik} + 2 \sum_{(i,j)} w(i, j)x_{ij} + \sum_k c(f_k)x_k \\ \text{subject to} \quad & x_{ijk} \leq x_{ik} & \forall (i, j), k \\ & x_{ijk} \leq x_{jk} & \forall (i, j), k \\ & 1 - x_{ij} \leq \sum_k x_{ijk} & \forall (i, j) \\ & x_k \geq x_{ik} & \forall i, k \\ & 0 \leq x_{ij} \leq 1 & \forall (i, j) \\ & 0 \leq x_{ijk} \leq 1 & \forall (i, j), k \\ & 0 \leq x_k \leq 1 & \forall k \end{aligned} \quad (1)$$

Algorithm 1. *Let $x_{ik}^*, x_{ij}^*, x_k^*, x_{ijk}^*$ denote the optimal solution to LP 1. For all $x_{ik}^* \geq \frac{1}{m+1}$, assign agent i to facility k .*

Theorem 8. *Algorithm 1 gives a $(m+1)$ -approximation to the optimal solution s^* .*

Algorithm 2. *Let $x_{ik}^*, x_{ij}^*, x_k^*, x_{ijk}^*$ denote the optimal solution to LP 1. For each facility f_k , we apply correlated randomized rounding on all x_{ik}^* as follows: first order all agents i by increasing order of x_{ik}^* . Without loss of generality, suppose $x_{1k}^* \leq x_{2k}^* \leq \dots \leq x_{nk}^*$. With probability x_{1k}^* , assign $x_{ik} = 1$ for all i . With probability $x_{jk}^* - x_{(j-1)k}^*$, assign $x_{ik} = 1$ for all $i \geq j$, and $x_{ik} = 0$ for all $i < j$. Finally, with probability $1 - x_{nk}^*$, assign $x_{ik} = 0$ for all i . Repeat this randomized rounding process $4 \ln 10n$ times, then assign agent i to facility k iff x_{ik} is assigned to 1 in at least one of the $4 \ln 10n$ runs.*

Theorem 9. *With high probability, the social cost of the solution given by Algorithm 2 is no more than $O(\ln n) \cdot c(s^*)$.*

Acknowledgements

This work was partially supported by NSF awards CNS-1816396 and CCF-1527497. We also thank Koushik Kar and Murat Yuksel for many useful discussions on the incentives of Internet Exchange Points.

References

- Ager, B.; Chatzis, N.; Feldmann, A.; Sarrar, N.; Uhlig, S.; and Willinger, W. 2012. Anatomy of a large European IXP. In *ACM SIGCOMM*.
- Anshelevich, E.; and Sekar, S. 2014. Approximate equilibrium and incentivizing social coordination. In *AAAI*.

- Apt, K. R.; Rahn, M.; Schäfer, G.; and Simon, S. 2014. Coordination games on graphs. In *WINE 2014*.
- Augustine, J.; Chen, N.; Elkind, E.; Fanelli, A.; Gravin, N.; and Shiryaev, D. 2011. Dynamics of profit-sharing games. In *IJCAI 2011*.
- Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2016a. Discrete preference games in heterogeneous social networks: Subverted majorities and the swing player. *CoRR*, abs/1603.02971 .
- Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2016b. Generalized Discrete Preference Games. In *IJCAI 2016*.
- Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2017. Robustness in discrete preference games. In *AAMAS 2017*.
- Aziz, H.; Brandl, F.; Brandt, F.; Harrenstein, P.; Olsen, M.; and Peters, D. 2019. Fractional hedonic games. *ACM Transactions on Economics and Computation (TEAC)* 7(2): 1–29.
- Aziz, H.; Brandt, F.; and Seidig, H. G. 2013. Computing desirable partitions in additively separable hedonic games. *Artificial Intelligence* 195: 316–334.
- Aziz, H.; and Savani, R. 2016. *Hedonic Games*. In Handbook of Computational Social Choice, Chapter 15, Cambridge University Press.
- Bachrach, Y.; Elkind, E.; Meir, R.; Pasechnik, D.; Zuckerman, M.; Rothe, J.; and Rosenschein, J. S. 2009. The cost of stability in coalitional games. In *SAGT 2009*.
- Bhalgat, A.; Chakraborty, T.; and Khanna, S. 2010. Approximating pure nash equilibrium in cut, party affiliation, and satisfiability games. In *EC 2010*.
- Bilò, V.; Fanelli, A.; Flammini, M.; Monaco, G.; and Moscardelli, L. 2014. Nash stability in fractional hedonic games. In *WINE 2014*.
- Bilò, V.; Fanelli, A.; Flammini, M.; Monaco, G.; and Moscardelli, L. 2015. On the price of stability of fractional hedonic games. In *AAMAS 2015*.
- Brânzei, S.; and Larson, K. 2009. Coalitional affinity games and the stability gap. In *IJCAI 2009*.
- Brânzei, S.; and Larson, K. 2011. Social distance games. In *IJCAI 2011*.
- Carosi, R.; and Monaco, G. 2019. Generalized graph k-coloring games. *Theory of Computing Systems* 1–14.
- Chierichetti, F.; Kleinberg, J.; and Oren, S. 2013. On discrete preferences and coordination. In *EC 2013*.
- Christodoulou, G.; and Koutsoupias, E. 2005. The price of anarchy of finite congestion games. In *STOC 2005*.
- Dreze, J. H.; and Greenberg, J. 1980. Hedonic coalitions: Optimality and stability. *Econometrica: Journal of the Econometric Society* 987–1003.
- Feldman, M.; and Friedler, O. 2015. A unified framework for strong price of anarchy in clustering games. In *ICALP 2015*.
- Feldman, M.; Lewin-Eytan, L.; and Naor, J. 2015. Hedonic clustering games. *ACM Transactions on Parallel Computing (TOPC)* 2(1): 1–48.
- Gairing, M.; and Savani, R. 2010. Computing stable outcomes in hedonic games. In *SAGT 2010*.
- Gourvès, L.; and Monnot, J. 2010. The max k-cut game and its strong equilibria. In *TAMC 2010*.
- Hochbaum, D. S. 1982. Heuristics for the fixed cost median problem. *Mathematical programming* 22(1): 148–162.
- Hofer, M. 2007. *Cost sharing and clustering under distributed competition*. Ph.D. thesis, Ph. D. Thesis, University of Konstanz.
- Monderer, D.; and Shapley, L. S. 1996. Potential games. *Games and economic behavior* 14(1): 124–143.
- Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2(1): 65–67.
- Shrimali, G.; and Kumar, S. 2006. Paid peering among internet service providers. In *GameNets 2006*, 11–es.
- Tardos, E.; and Wexler, T. 2007. Network formation games and the potential function method. *Chapter 19, Algorithmic Game Theory* 487–516.