Reducibility Among Fractional Stability Problems

Shiva Kintali*, Laura J. Poplawski†, Rajmohan Rajaraman†, Ravi Sundaram†, and Shang-Hua Teng ‡
*College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, Email: kintali@cc.gatech.edu
†College of Computer and Information Science, Northeastern University, Boston, MA 02115
Email: ljp@ccs.neu.edu, rraj@ccs.neu.edu, koods@ccs.neu.edu
‡Computer Science Department, University of Southern California, Los Angeles, CA 90089, Email: shanghua@usc.edu

Abstract— In a landmark paper [32], Papadimitriou introduced a number of syntactic subclasses of $\mathbf{TFNP}$ based on proof styles that (unlike $\mathbf{TFNP}$) admit complete problems. A recent series of results [11], [16], [5], [6], [7], [8] has shown that finding Nash equilibria is complete for $\mathbf{PPAD}$, a particularly notable subclass of $\mathbf{TFNP}$. A major goal of this work is to expand the universe of known $\mathbf{PPAD}$-complete problems. We resolve the computational complexity of a number of outstanding open problems with practical applications.

Here is the list of problems we show to be $\mathbf{PPAD}$-complete, along with the domains of practical significance: Fractional Stable Paths Problem (FSPP) [18] - Internet routing; Core of Balanced Games [34] - Economics and Game theory; Scarf's Lemma [34] - Combinatorics; Hypergraph Matching [1]- Social Choice and Preference Systems; Fractional Bounded Budget Connection Games (FBBC) [26] - Social networks; and Strong Fractional Kernel [2]- Graph Theory. In fact, we show that no fully polynomial-time approximation schemes exist (unless $\mathbf{PPAD}$ is in $\mathbf{FP}$).

This paper is entirely a series of reductions that build in non-trivial ways on the framework established in previous work. In the course of deriving these reductions, we created two new concepts - preference games and personalized equilibria. The entire set of new reductions can be presented as a lattice with the above problems sandwiched between preference games (at the "easy" end) and personalized equilibria (at the "hard" end). Our completeness results extend to natural approximate versions of most of these problems. On a technical note, we wish to highlight our novel "continuous-to-discrete" reduction from exact personalized equilibria to approximate personalized equilibria using a linear program augmented with an exponential number of "min" constraints of a specific form. In addition to enhancing our repertoire of $\mathbf{PPAD}$-complete problems, we expect the concepts and techniques in this paper to find future use in algorithmic game theory.

Keywords-Complexity theory, Game theory, Stability

1. INTRODUCTION

Intuitively, the notion of stability implies the absence of oscillations over time and encompasses the concepts of fixed points and equilibria. Stability is important in a variety of fields ranging from the practical – the Internet – to the theoretical – combinatorics and game theory. For important practical systems (e.g. Internet), the existence and computational feasibility of stable operating modes is of profound real-world significance. On the more abstract front, the study of stable solutions to combinatorial problems has a distinguished tradition dating back to, at least, the Gale-Shapley algorithm [14]. It is often the case, as with Nash's celebrated theorem [30], that *fractional* stable points are guaranteed to exist even when integral points do not.

In this paper, we focus on fractional stability and introduce two new fractional stability problems. The first, **Preference Games**, is a $\mathbf{PPAD}$-complete problem with a simple definition that can be reduced to a number of pre-existing problems. The other, **Personalized Equilibria** for matrix games, is also $\mathbf{PPAD}$-complete and generalizes several pre-existing problems. We use these new tools to resolve the complexity of four pre-existing problems with applications to a variety of different domains. Below we provide elaborate motivation for two of the pre-existing problems - **Fractional Stable Paths Problem (FSPP)** and **Core of Balanced Games**. The others are: **Scarf's lemma**, a fundamental result in combinatorics with several applications [34], and **Fractional Hypergraph Matching** [1], useful for modeling preferences in social-choice and economic systems. In the full version of this paper ([23]), we use the same tools to also resolve the complexity of **FBBC**, the fractional version of the Bounded Budget Connection (BBC) game [26], which models decentralized overlay network creation and social networks, and **Strong Fractional Kernel** [2], of relevance to structural graph theory.

Fractional Stable Paths Problem. Griffin, Shepherd and Wilfong [17] showed how BGP (Border Gateway Protocol, the routing mechanism of the Internet) can be viewed as a distributed mechanism for solving the Stable Paths Problem (SPP). They showed that there exist SPP instances with no integral stable solutions, a phenomenon that would explain why oscillation has been observed in Internet routes. Route oscillation is viewed as a negative, since it imposes higher system overheads, reorders packets, and creates difficulties for tracing and debugging. Subsequently, Haxell and Wilfong [18] introduced FSPP: a natural fractional relaxation of SPP with the property that a (fractional) stable solution always exists. Intuitively, FSPP can be viewed as a game played between Autonomous Systems that each assign fractional capacities to the different paths leading to a destination in such a way that they maximize their utility without violating the capacity constraints of downstream nodes. Understanding the computational feasibility of finding the

‡work done at Boston University and Microsoft Research New England

equilibria of this game could help to develop techniques for stable routing in the Internet.

Core of balanced games. The notion of *core* in cooperative games is analogous to that of Nash equilibrium in non-cooperative games. Informally, a core is the set of all outcomes in which no coalition of players has an incentive to secede and obtain a better payoff, either viewed as a set (transferable utilities) or individually (non-transferable utilities). Necessary and sufficient conditions for the nonemptiness of the core in games with transferable utilities is given by the classic Bondareva-Shapley theorem [4], [35], which also yields a polynomial-time algorithm for finding an element in a nonempty core. Subsequently, in a celebrated paper, Scarf [34] generalized their result, developed certain sufficient *balance* conditions for the nonemptiness of the core in games with non-transferable utilities, and presented an algorithm for finding a point in the core. As noted by Jain and Mahdian in Chapter 15 of [31], "However, the worst case running time of this algorithm (like the Lemke-Howson algorithm) is exponential." Resolving the computational feasibility of finding the core in balanced games is of considerable significance in the theory of cooperative games.

Personalized equilibria for matrix games - a generalization. Personalized equilibria is a solution concept intended to generalize fractional stability problems. We are given any matrix game – a set of players, a set of actions for each player, and a definition of the utility for each player for any combination of actions. In a personalized equilibrium, each player may choose any distribution over his actions, and (unlike a traditional matrix game) each player may match his actions to those of the other players in such a way as to maximize his individual payoff. To better illustrate the concept, let us imagine a business manufacturing and selling outfits consisting of a pant (solid or striped) and a shirt (cotton or wool). The manager of the location producing pants decides on the ratio of striped pants produced to solid pants while the manager at the location producing shirts decides on the ratio of cotton shirts produced to wool shirts. Each manager is then given half the total number of shirts and pants (in the proportions decided) and has to match them into outfits and sell them at her own location in such a way as to maximize her individual profits.

Preference games - a specialization. In a preference game, the set of actions for each player is just the set of all players. Each player also has an ordinal preference list across the actions. Each player must choose some distribution of weights across the action set such that (a) it does not assign more weight to another player than that player assigns to itself, and (b) no weight may be moved from a lower preference action to a higher preference action. To better illustrate this problem, consider a world of bloggers where each blogger has a choice of actions. They can fill their blogs with original content or they can copy from the original content on others' blogs. Naturally, each blogger has a preference order over the content of the different bloggers (as well as their own). Also, of course, more cannot be copied from another blog than the amount that other blogger has written. The preference game models each blogger's choice of what percentage of his blog is original and what percentages are copied from which other blogs. Preference games are reducible in polynomial-time to all the problems considered in this paper.

1.1. Our Contributions

Hewing to the dictum that a picture is worth a thousand words, we present a diagram (Figure 1) showing the different reductions. The takeaway is that all of these problems are $\mathbf{PPAD}$-complete. To be precise, we show the exact versions of these problems are in $\mathbf{PPAD}$, and our reductions extend to natural approximation versions to show that there are no fully polynomial-time approximation schemes (unless $\mathbf{PPAD}$ is in $\mathbf{FP}$). Our reductions build on prior work in intricate and involved fashion.

From a conceptual standpoint, we believe there is merit in the definitions of preference games and personalized equilibria. Preference games are very simple to describe and model a number of real-world situations, such as the blogger example mentioned earlier. However, we can show that the set of equilibria of preference games can be nonconvex and in fact, are hard even to approximate. Personalized equilibria of matrix games are, we believe, a fascinating solution concept which constitute a natural generalization of a variety of predefined games, such as FSPP and FBBC. Our results on the hardness of approximating personalized equilibria for k-player games apply for $k \geq 4$. We show in the full version [23] that finding personalized equilibria of 2-player games is in $\mathbf{FP}$. The $k = 3$ case is open.

1.2. Related Work

Nash profoundly changed game theory by demonstrating the existence of mixed equilibria [29], [30]. Decades later, on the computational front [31], the complexity class $\mathbf{TFNP}$ was introduced by Megiddo and Papadimitriou [28]. Papadimitriou's seminal work [32] not only defined a number of syntactic subclasses of $\mathbf{TFNP}$ (including $\mathbf{PPAD}$), but also proved that a variety of problems, including discrete versions of Brouwer's fixed point theorem and Sperner's lemma, are $\mathbf{PPAD}$-complete. The problem of finding Nash equilibria was left open. Recently, a series of papers comprising different author combinations of the two teams, Daskalaikis-Goldberg-Papadimitriou [11], [16] and Chen-Deng-Teng [5], [6], [7], [8] culminated in establishing that even *approximating* Nash equilibria with two players, 2-NASH, is hard. The reductions in our work build on the framework established in these papers.

BGP has been the focus of much attention since its inception [33], [36]. As mentioned earlier SPP was introduced by Griffin, Shepherd and Wilfong [17] to explain the nonconvergence of BGP [38]. Haxell and Wilfong [18] defined FSPP and proved the existence of an equilibrium using Scarf's lemma and a compactness-type argument. They left open the complexity of finding an equilibrium. Our reduction from personalized equilibria to End-of-the-Line is a different approach that generalizes the Haxell-Wilfong existence result while preserving computational tractability. Kintali [22] presented a distributed algorithm for finding an ϵ-approximation for FSPP that is guaranteed to converge, although no bounds are given on the time-to-convergence (our results imply a polynomial time bound is unlikely).

Cooperative games, the study of mechanisms to sustain and enforce cooperation among willing agents, has a rich and extensive literature [9], [15], [10], [13], [25]. As mentioned earlier, Scarf [34] generalized the classical Bondareva-Shapley theorem [4], [35] and developed an algorithm for finding a point in the core of balanced games with non-transferable utilities. More recently, Markakis and Saberi [27], Immorlica, Jain and Mahdian [20] studied certain classes of games with non-transferable utilities in the context of the Internet; however, it is unclear that their problems are even in $\mathbf{TFNP}$. Scarf's paper [34] also contains Scarf's lemma, an important result in combinatorics which played a crucial part in the FSPP existence proof of Haxell and Wilfong [18]. Aharoni and Fleiner [1] proved that every hypergraphic preference system has a fractional stable matching. This proof is based on Scarf's lemma. The computational complexity of these problems was left unresolved.

2. The Class PPAD

A major contribution of this paper is to expand the set of problems known to be $\mathbf{PPAD}$-complete. The class $\mathbf{PPAD}$ (***P**olynomial **P**arity **A**rgument in a **D**irected graph*) was introduced by Papadimitriou in [32], which defined a number of syntactic classes in the semantic class $\mathbf{TFNP}$, or the set of all total search problems. A search problem $\mathcal{S}$ consists of a set of inputs $I_{\mathcal{S}} \subseteq \Sigma^*$ such that for each $x \in I_{\mathcal{S}}$ there is an associated set of solutions $\mathcal{S}_x \subseteq \Sigma^{|x|^k}$ for some integer k. For each $x \in I_{\mathcal{S}}$ and $y \in \Sigma^{|x|^k}$, it is decidable in polynomial time whether or not y is in $\mathcal{S}_x$. A search problem is *total* if $\mathcal{S}_x \neq \emptyset$ for all $x \in I_{\mathcal{S}}$. $\mathbf{TFNP}$ is the set of all total search problems [28]. Since every member of $\mathbf{TFNP}$ is equipped with a mathematical proof that it belongs to $\mathbf{TFNP}$, a number of syntactic classes can be defined based on their proof styles. The complexity class $\mathbf{PPAD}$ is the class of all search problems whose totality is proved using a directed parity argument.

Problems in $\mathbf{PPAD}$ are reducible to the End of the Line problem. In End of the Line, we are given a finite directed graph in which each node has at most one outgoing edge and at most one incoming edge. The input to the problem is not a complete list of the nodes and edges; such a list may be exponentially large in the size of the input. Instead, we are given an initial source node and a circuit. The circuit takes a node name as input and in polynomial time returns the *next* node (the other end of the outgoing edge from the input node) and the *previous* node (the other end of the incoming edge into the input node). If the input node is a source (or sink), null is returned as the previous (or next) node. The problem for End of the Line is to find a sink or a source other than the initial source.

Throughout this paper, we use problem A $\leq_P$ problem B to mean "There exists a polynomial time reduction from finding a stable point in problem A to finding a stable point in problem B."

3. Preference Games

In this section, we define a very simple game, the *preference game*. Each player has a preference list across the set of players and must assign a weight to each player. No player may put more weight on another player than that player puts on itself. A set of weight assignments is stable if no player can move weight from a less preferred player to a more preferred player. We precisely define preference games in Section 3.1. In Section 3.2, we show that finding an equilibrium in preference games is $\mathbf{PPAD}$-hard, and in Section 3.3, we define an ϵ-equilibrium for the preference game. Our $\mathbf{PPAD}$-hardness result can be extended to approximate equilibria (see the full version [23]). Our notion of approximation is very general and carries though all of the reductions in later sections, so we prove that there are no fully polynomial-time approximation schemes (unless $\mathbf{PPAD}$ is in **FP**) for computing stable points in any of the $\mathbf{PPAD}$-complete problems discussed in this paper. Finally, in section 3.4, we define the *degree* of a preference game and show that it is $\mathbf{PPAD}$-hard to find an equilibrium even in a preference game with constant degree.

3.1. Definition

In a preference game with a set S of players, each player's strategy set is S. Each player $i \in S$ has a preference relation $\succeq_i$ among the strategies. For strategies j and k, $j \succeq_i k$ indicates that player i prefers j at least as much as k. When it is clear from context that we are talking about the preferences for player i, we write $j \succeq k$ instead of $j \succeq_i k$. Each player i chooses a *weight distribution*, which is an assignment $w_i : S \to [0, 1]$ satisfying two conditions: (a) the weights add up to 1: $\sum_{j \in S} w_i(j) = 1$; and (b) the weight placed by i on j is no more than the weight placed by j on j: $w_i(j) \leq w_j(j)$ for all $i, j \in S$.

Given weight assignments w_i, w'_i, and w_{-i} such that (w_i, w_{-i}) and (w'_i, w_{-i}) are both feasible, we say w_i is *lexicographically at least* w'_i (with respect to w_{-i}) if for all $j \in S$, $\sum_{k \succeq_i j} w_i(k) \geq \sum_{k \succeq_i j} w'_i(k)$. We say that w_i is *lexicographically maximal* if (w_i, w_{-i}) is feasible and w_i is lexicographically at least every assignment w'_i such that (w'_i, w_{-i}) is feasible. An equilibrium in a preference game is an assignment $w = \{w_i : i \in S\}$ such that w_i is lexicographically maximal for all $i \in S$.

Every preference game has an equilibrium, a fact which can be shown using standard fixed-point theorems.

3.2. PPAD Hardness

In this section, we study PREFERENCE GAME, the problem of finding an equilibrium in a preference game. If the preferences exhibit symmetry, then it is easy to find an equilibrium in which all weights are either 0 or 1 (see [23]). With general preferences, however, the set of equilibria may have a more complicated structure. There exist instances for which the set of equilibria is not convex.

We show that PREFERENCE GAME is **PPAD**-hard. Based on the framework in [11], we can reduce finding an exact equilibrium to the **PPAD**-complete problem 3-D BROUWER if we can show how to create a set of gadgets in the preference game to do simple computations and logic operations. In our preference game gadgets, each player will assign to itself an amount of weight $\in [0, 1]$ which equals the result of the computation or logic operation. Descriptions of these gadgets follow. The complete reduction can be found in the full version [23]), including correctness analysis for each of these gadgets.

For the gadget descriptions, assume we are given player X that plays itself with weight v_1 in any equilibrium, and player Y that plays itself with weight v_2 in any equilibrium. For the first three gadgets, $v_1, v_2 \in \{0, 1\}$. For the rest of the gadgets, $v_1, v_2 \in [0, 1]$. $L(X)$ is the preference list for node X. Additional nodes are required for some of these gadgets, as indicated in the description of how to create each gadget.

R = **OR**(X, Y). Plays itself: $max(1, v_1 + v_2) = v_1 \vee v_2$.
To create: $L(R_1) = (X, Y, R_1)$, $L(R) = (R_1, R)$.

N = **NOT**(X). Plays itself: $1 - v_1 = \neg v_1$.
To create: $L(N) = (X, N)$.

A = **AND**(X, Y). Plays itself: $v_1 \wedge v_2$.
To create: $A = \text{NOT}(\text{OR}(\text{NOT}(X), \text{NOT}(Y)))$.

S = **SUM**(X, Y). $S = \text{OR}(X, Y)$.

D = **DIFF**(X, Y). Plays itself: if $v_1 > v_2$, $v_1 - v_2$. else, 0.
To create: $D = \text{NOT}(\text{SUM}(\text{NOT}(X), Y))$

C = **COPY**(X). Plays itself: v_1.
To create: $C = \text{NOT}(\text{NOT}(X))$

M = **DOUBLE**(X). Plays itself: $\min(1, v_1 * 2)$
To create: $M_1 = \text{COPY}(X)$, $M = \text{SUM}(X, M_1)$.

L = **LESS**(X, Y). Plays itself: (given ϵ_l, $0 < \epsilon_l \leq \frac{1}{2}$) $v_1 - v_2 \geq \epsilon_l$ then 1. if $v_1 \leq v_2$, then 0.
To create: $D = \text{DIFF}(X, Y)$. $M_1 = \text{DOUBLE}(D)$. For $i = 1$ to $-\log \epsilon_l$, $M_{i+1} = \text{DOUBLE}(M_i)$. L = the last DOUBLE node

H = **HALF**(X). Plays itself: $v_1/2$.
To create: $L(H_1) = (X, H_1)$. $L(H_2) = (H_1, H_3, H_2)$. $L(H_3) = (H_1, H, H_3)$. $L(H) = (H_1, H_2, H)$.

Based on the above gadgets and the framework from [11], we get the following.

Theorem 3.1. 3-D BROUWER $\leq_P$ PREFERENCE GAME. □

3.3. Approximate equilibria

Given the hardness of finding exact equilibria in preference games, a natural next question is whether it is easier to find approximate equilibria. We define an ϵ-equilibrium of a k-player preference game to be a set of weight distributions $w_1, \ldots, w_k$ that satisfy the following conditions for every player i: (a) $\sum_j w_i(j) = 1$; (b) for each j, $w_i(j) \leq w_j(j) + \epsilon$; and (c) for each j, either $\sum_{\ell:\ell \succeq j} w_i(\ell) \geq 1 - \epsilon$ or $|w_i(j) - w_j(j)| \leq \epsilon$. The problem of finding an ϵ-equilibrium is ϵ-APPROXIMATE PREFERENCE GAME.

BROUWER $\leq_P$ ϵ-APPROXIMATE PREFERENCE GAME. Thus, it is **PPAD**-hard to find an ϵ-equilibrium for preference games for ϵ inverse polynomial in n. This can be shown by adjusting the **PPAD**-hardness reduction to start with a higher dimensional fixed point problem, BROUWER, as in [7], [8]. The gadgets described above can be analysed carefully to show that each only slightly amplifies errors in the input, so that following each gadget with a LESS gadget will reduce the errors back to at most ϵ_l. For a full proof, refer to [23].

3.4. Constant degree preference games

For a given preference game, define in(v) (resp., out(v)) of a player v to be the set $\{u : v \succ_u u\}$ (resp., $\{u : v \prec_v u\}$). We define the in-degree (resp., out-degree) of a player v to be $|\text{in}(v)|$ (resp., $|\text{out}(v)|$). The degree of the player is defined to be the sum of the in-degree and the out-degree of the player. The degree of the preference game is defined to be the maximum degree of any node. Notice that this is the same as the degree in a directed graph in which each player is represented by a node, and an edge from u to v means that u prefers v over itself. DEGREE d PREFERENCE GAME is the problem of finding an equilibrium in a preference game with constant degree d.

The players used to show **PPAD**-hardness in Section 3.2 all have out-degree at most 2, and the in-degree can be reduced using COPY gadgets to ensure in-degree at most 1 for all players except COPY players, and in-degree at most 2 for COPY players (which have out-degree 1). This immediately implies that it is **PPAD**-hard to find an equilibrium even in a preference game with degree 3. We will use this fact in later sections, where we show **PPAD**-hardness of several other problems via reductions from constant degree preference games (and specifically from DEGREE 3 PREFERENCE GAME). In the full version [23], we also give an explicit reduction to show that PREFERENCE GAME $\leq_P$ DEGREE d PREFERENCE GAME.

4. Personalized Equilibria

In this section, we introduce a new notion of equilibrium for matrix games, in which a player may individually match her strategies to her opponents strategies without obeying a product distribution. Since this equilibrium allows different players to simultaneously choose different matchings across the strategies, we call this a *personalized equilibrium*. Section 4.1 contains a formal definition. In Section 4.2, we characterize the set of all personalized equilibria in a k-player game. In Section 4.3, we show that finding a personalized equilibrium is $\mathbf{PPAD}$-complete.

4.1. Definition

Suppose we are given a k-player matrix game between players $1, \ldots, k$. Each player i has strategy set S_i. We are also given a utility function for each i specified by $u_i : E \to \mathbb{R}$, where $E = \prod_j S_j$. Now, given probability distributions $p_j(S_j)$ for each $j \neq i$, a best response for player i (using traditional Nash payoffs) is defined by the $p_i(S_i)$ that satisfies the following, where w is a weight function over $e \in E$.

$$\begin{aligned} \max \sum_{e \in E} & w(e) u_i(e) \\ w(e) &= \prod_{s \in e \cap S_j} p_j(s) \qquad \text{for all } e \in E \\ w(e) &\geq 0 \qquad\qquad\qquad \text{for all } e \in E \end{aligned}$$

The correlator in a correlated equilibrium [3] relaxes the requirement that w be a product distribution; however, w does satisfy, among other conditions, the projection constraint $\sum_{e:s \in e} w(e) = p_j(s)$ for all $s \in S_j, 1 \leq j \leq k$. For a personalized equilibrium, we further relax this by allowing each player to define her own weight function, w_i, so that in the best response of player i, $p_i(s)$ and $w_i(e)$ satisfy the following.

$$\begin{aligned} \max \sum_{e \in E} & w_i(e) u_i(e) \\ \sum_{e:s \in e} w_i(e) &= p_j(s) \quad s \in S_j, 1 \leq j \leq k \\ w_i(e) &\geq 0 \qquad e \in E \end{aligned}$$

We can view a matrix game as a hypergraph with nodes $V = \cup_j S_j$ and edges $E = \prod_j S_j$. Then, if we interpret the $p_j(s)$ values as capacities on the nodes and the utility function for player i as weights on the edges from the perspective of player i, a personalized equilibrium is simultaneously a maximum-weight fractional hypergraph matching for each player.

The description of the game above is exponential in the number of players, since we require that every edge connects one strategy of each player. To allow for more succinct descriptions, we generalize the game as follows. For each player i, we introduce a hypergraph with nodes $V = \cup_j S_j$ and edges E_i. The set E_i is required to satisfy two conditions (that are satisfied by E): (i) for each e in E_i and player j, e contains at most one element of S_j; (ii) there do not exist distinct e and e' in E_i such that $e \subset e'$. In the game, player i places a weight $w_i(e)$ on each edge in E_i. A player must still place a total of weight 1 on all her edges, and all weights must be non-negative. Since the edges of E_i may not connect all players, however, we relax the projection constraint to $\sum_{e:s \in e} w_i(e) \leq p_j(s)$. Thus, the collection of weights $w_i(e)$, $e \in E_i$, and probability distributions $p_i(s)$, $s \in S_i$, over all players i, form a personalized equilibrium if for each i, $w_i(e)$ and $p_i(s)$ maximize $\sum_{e \in E_i} w_i(e) u_i(e)$ subject to the following constraints.

$$\begin{aligned} \sum_{e:s \in e} w_i(e) &\leq p_j(s), \qquad \forall s \in S_j, \forall j \neq i \\ \sum_{e:s \in e} w_i(e) &= p_i(s), \qquad \forall s \in S_i \\ w_i(e) &\geq 0, \qquad\quad \forall e \in E_i \end{aligned}$$

Just as mixed Nash equilibria exist for every matrix game, we show that every game thus defined has a personalized equilibrium. We defer the proof of the following theorem to the full version [23].

Theorem 4.1. *For every multi-player matrix game, a personalized equilibrium always exists.*

We define Personalized Equilibrium as the problem of finding a personalized equilibrium in a given matrix game. k-Personalized Equilibrium is the same problem in a game with k players for constant k. Note that the traditional definition of a graphical game [21] may be used in this setting with smaller edges. In d-Graphical Personalized Equilibrium, each player i has a neighborhood N_i of at most d other players, and all edges defined for player i are in $\prod_{j \in N_i} S_j$. Finally, we define ϵ-Approximate Personalized Equilibrium as the problem of finding a set of weight assignments ($w_i(e) \geq 0$ is the weight assigned by player i to edge e) such that (a) for every player i, $1 - \epsilon \leq \sum_e w_i(e) \leq 1$, (b) for each player pair i and j, and for each strategy s, $\left|\sum_{e:s \in e} w_i(e) - \sum_{e:s \in e} w_j(e)\right| \leq \epsilon$, and (c) for any best response weight assignment w_i^* for any player i, $\sum_e w_i^*(e) u_i(e) - \sum_e w_i(e) u_i(e) \leq \epsilon$.

4.2. Characterizing personalized equilibria in k-player games

The set of all personalized equilibria for a two-player game is just the set of all linear combinations of cycles in an appropriately defined graph (see [23]), which is easy to compute in polynomial time. However, for k player games ($k > 3$), we will give a reduction from finding an equilibrium in a preference game to finding a personalized equilibrium in a k player game, thereby showing that finding personalized equilibria is $\mathbf{PPAD}$-hard. Nevertheless, we are able to give a concise characterization of the set of all personalized equilibria for arbitrary multi-player games. This characterization will be useful for showing $\mathbf{PPAD}$ membership of PERSONALIZED EQUILIBRIUM.

Theorem 4.2 (Personalized Equilibrium Characterization)**.** *The following program represents the set of all exact personalized equilibria. The variables are $w_i(e)$, the weight placed by player i on edge e, $\forall e \in E_i$.*

$$\begin{aligned}
&\sum_{e\in E_i:s\in e} w_i(e) \le \sum_{e\in E_j:s\in e} w_j(e),\ s\in S_j, 1\le j,i\le k\\
&\sum_{e\in E_i} w_i(e) = 1, \qquad 1\le i\le k \qquad (1)\\
&w_i(e)\ge 0, \qquad 1\le i\le k, e\in E_i\\
&\min_{e\in F} w_i(e) = 0, \qquad \text{for all players } i \text{ and subsets. } F\subseteq E_i \text{ such that LP (2) is feasible}
\end{aligned}$$

The following linear program is defined for each player i and $F \subseteq E_i$ (referred to as an improvement set*). The variables are $\delta(e)$ for each edge $e \in E_i$.*

$$\begin{aligned}
&\sum_{e\in E_i} \delta(e)u_i(e) > 0 \qquad (2)\\
&\sum_{e\in E_i:s\in e} \delta(e) = 0 \qquad s\in S_j, 1\le j\le k, j\ne i\\
&\delta(e) < 0,\ (e\in F) \qquad \delta(e)\ge 0,\ (e\notin F)
\end{aligned}$$

A formal proof of the above theorem can be found in the full version ([23]), but here we provide some intuition. The first two constraints of program 1 specify a feasible weight assignment, and the first two constraints of LP 2 specify feasible "weight changes" that would increase the payoff for player i. How do we know that checking this for all subsets of edges is enough to find any possible improvement, and how does the last constraint of program 1 ensure that no improvement is possible? We can think of the δ values found in any solution to LP 2 as an "improvement direction." This is a vector that is orthogonal to the vector of all 1's and has a positive dot product with the utilities of i. In other words, if player i were to move weight in this direction, her payoff would improve. Of course, there may be a continuum of such improvement directions. However, there are most an exponential number of negative supports, or "improvement sets". These are exactly the F values for which LP 2 is feasible. Given an improvement set, the associated player can get a higher payoff by removing weight from all of those edges and adding the weight instead to edges with positive δ value. This improvement will be possible unless the player has zero weight on some edge in this entire improvement set; that is, unless $\min_{e\in F} w_i(e) = 0$. Theorem 4.2 leads to the following corollary, since the above linear program with additional min constraints can be re-written as a union of linear programs, one with each subset of one edge from each improvement set explicitly set to 0. See [23] for more details.

Corollary 4.3. *For any matrix game with all rational payoffs, there exists a personalized equilibrium in which the probability assigned by each player to each strategy is a rational number.*

4.3. Finding personalized equilibria is $\mathbf{PPAD}$-complete

In order to show that PERSONALIZED EQUILIBRIUM is $\mathbf{PPAD}$-complete, we will use two chains of reductions. First, to show $\mathbf{PPAD}$-hardness (with ≥ 4 players): DEGREE 3 PREFERENCE GAME $\le_P$ 3-GRAPHICAL PERSONALIZED EQUILIBRIUM $\le_P$ 4-PERSONALIZED EQUILIBRIUM. It is easy to verify that the reductions in this first chain can also be used to show ϵ-APPROXIMATE PREFERENCE GAME $\le_P$ ϵ-APPROXIMATE PERSONALIZED EQUILIBRIUM. Then, to show $\mathbf{PPAD}$-membership: PERSONALIZED EQUILIBRIUM $\le_P$ ϵ-APPROXIMATE PERSONALIZED EQUILIBRIUM $\le_P$ END OF THE LINE (as long as ϵ is sufficiently small).

Theorem 4.4. DEGREE d PREFERENCE GAME $\le_P$ d-GRAPHICAL PERSONALIZED EQUILIBRIUM

Proof: Given a preference game over player set $[n]$, with the preference lists specified as a set of values Q_{ij} for all $i, j \in [n]$: Q_{ij} = the number of players k such that $j \succeq_i k \succeq_i i$. Define a game as follows, in which we will find a personalized equilibrium. The set of players = $\{p_1, \dots, p_n\}$. S_i (the set of strategies for player p_i) = $\{s_{ij} : Q_{ij} > 0\}$. E_i = the set of edges for player p_i = $\{\{s_{ij}, s_{jj}\} \forall s_{ij} \in S_i, j \ne i\} \cup \{s_{ii}\}$. $u_i(\{s_{ij}, s_{jj}\})$ (the payoff to player i for this edge) = Q_{ij}. $u_i(\{s_{ii}\}) = Q_{ii} \ge 1$. Notice that the degree of the game is preserved, and the number of edges defined is at most n times the degree. Suppose we are given weights x_{ij} for each player i and edge $\{s_{ij}, s_{jj}\}$, and x_{ii} for player i and edge $\{s_{ii}\}$. These weights form a personalized equilibrium if and only if weights $w_{ij} = x_{ij}$ are an equilibrium in the preference game. We defer the correctness proof to the full version [23]. ■

Two of the remaining reductions follow patterns introduced in previous work. Our reduction showing 3-GRAPHICAL PERSONALIZED EQUILIBRIUM $\le_P$ 4-PERSONALIZED EQUILIBRIUM follows the pattern established in [16] for mixed Nash equilibria, and our reduction showing ϵ-APPROXIMATE PERSONALIZED EQUILIBRIUM $\le_P$ END OF THE LINE follows the pattern from [12][Section 3.2] demonstrating $\mathbf{PPAD}$-membership of approximate mixed Nash equilibria. Both of these reductions can be found in the full version of this paper [23].

Finally, we have PERSONALIZED EQUILIBRIUM $\le_P$ ϵ-APPROXIMATE PERSONALIZED EQUILIBRIUM for sufficiently small ϵ. Our reduction is based on an LP Compactness Lemma, Lemma 4.5. The basic idea for this reduction follows. First, we assume that we can find an approximate personalized equilibrium for our matrix game. We use Corollary 4.6 of Lemma 4.5 to show that this approximate equilibrium *almost* obeys every constraint in LP 1. We can adjust the solution by forcing any variable that is very close to 0 down to 0 and get a solution that completely obeys all of the min constraints from LP 1 and still almost obeys each of the other constraints. Now, we again use Lemma 4.5, which says that if there is a point that comes very close to obeying each constraint of a linear program, then the linear program must be feasible. We apply this to LP 1 with the min constraints replaced by "$= 0$" constraints for those values we've already set to 0. Our adjusted approximate solution comes close to obeying each constraint, so the LP has a feasible solution. Since the min constraints have been removed, this is now a polynomially sized LP, so we can solve it in polynomial time to find an exact personalized equilibrium. A full proof as well as the exact bound on ϵ can be found in the full version [23], but the key LP Compactness Lemma and its Corollary are included here.

Lemma 4.5 (LP Compactness)**.** *If an LP with n variables and rational coefficients, each represented by at most β bits, is such that there is a point obeying each constraint to within $\epsilon = \frac{1}{2^{3n\beta}}$, then the LP is feasible.*

Proof: If t, b, t_i, b_i, y_i, z_i, (for $1 \leq i \leq n$), are β-bit integers, then either $\sum_{i=1}^{n} \frac{t_i y_i}{b_i z_i} \geq \frac{t}{b}$ or $\sum_{i=1}^{n} \frac{t_i y_i}{b_i z_i} < \frac{t}{b} - \frac{1}{2^{3n\beta}}$. To show this, suppose we have $\sum_{i=1}^{n} \frac{t_i y_i}{b_i z_i} < \frac{t}{b}$. Then the difference $\frac{t}{b} - \sum_{i=1}^{n} \frac{t_i y_i}{b_i z_i}$ is at least $1/(b \cdot \prod_i b_i \prod_i z_i)$, which is at least $1/2^{\beta+2n\beta} < 2^{-3n\beta}$ since each integer in the product is at most 2^β.

This gives us the following. If x satisfies $\sum_{i=1}^{n} a_i x_i \geq b - \frac{1}{2^{3n\beta}}$, where each a_i and b are rational numbers whose numerators and denominators are representable as β-bit integers, then x satisfies $\sum_{i=1}^{n} a_i x_i \geq b$. This immediately implies the Lemma. ■

Corollary 4.6. *Given a linear program with $\leq n$ variables and coefficients of the form $\frac{a}{b}$ for integers a and b, each represented by at most β bits, each coordinate of a vertex must be representable by $\frac{c}{d}$ for integers c and d, each represented by less than $3n\beta$ bits.*

5. Scarf's Lemma and Fractional Stability Problems

This section discusses the complexity of a number of well-known combinatorial problems that can be categorized as *fractional stability problems*. We begin with Scarf's Lemma, a fundamental result in combinatorics, originally introduced to prove that every balanced cooperative game with non-transferable utilities has a nonempty core (see Section 5.3) [34]. The core (no pun intended) of his argument is an elegant and constructive combinatorial argument which has been applied to diverse combinatorial problems, including fractional stable matchings in hypergraphic preference systems, strong kernels in digraphs, and the fractional stable paths problem [2], [1], [24], [18]. We first show that the computational version of Scarf's lemma is **PPAD**-complete (Section 5.1). We then establish the **PPAD**-completeness of stable matchings in hypergraphic preference systems (Section 5.2), core of balanced games with non-transferable utility (Section 5.3), and the fractional stable paths problem (Section 5.4). We mention two additional problems in Section 5.5 which are also shown to be **PPAD**-complete in the full version of this paper [23].

5.1. Scarf's Lemma

In the computational version of Scarf's lemma (Scarf) we are given matrices B, C and a vector b satisfying the conditions in Theorem 5.1. The goal is to find $\alpha \in \mathbb{R}_+^n$ satisfying the desired properties.

Theorem 5.1. *(Scarf's lemma [34]) Let $I = [\delta_{ij}]$ be an $m \times m$ identity matrix. Let $[n] = \{1, 2, \ldots, n\}$. Let $m < n$ and let B be an $m \times n$ real matrix such that $b_{ij} = \delta_{ij}$ for $1 \leqslant i, j \leqslant m$. Let b be a non-negative vector in $\mathbb{R}^m$, such that the set $\{\alpha \in \mathbb{R}_+^n : B\alpha = b\}$ is bounded. Let C be an $m \times n$ matrix such that $c_{ii} \leqslant c_{ik} \leqslant c_{ij}$ whenever $i, j \leqslant m$, $i \neq j$ and $k > m$. Then there exists a subset $J \subset [n]$ of size m such that* **(P1)** *$B\alpha = b$ for some $\alpha \in \mathbb{R}_+^n$ such that $\alpha_j = 0$ whenever $j \notin J$, and* **(P2)** *For every $k \in [n]$ there exists $i \in [m]$ such that $c_{ik} \leqslant c_{ij}$ for all $j \in J$.*

A subset $J \subset [n]$ of size m is called a *feasible basis* of (B, b) if it satisfies **(P1)**, and *subordinating* if it satisfies **(P2)**. To compute α of Scarf, it suffices to have a $J \subseteq [n]$ that is simultaneously subordinating and a feasible basis. Once such J is computed, α can be computed by solving a system of linear equations. Also, given a solution α, J is easy to compute, since J is α's support. Hence finding α and J are computationally equivalent to within polynomial time. Scarf's original proof [34], together with Todd's orientation technique [37], gives an end of the line argument for the existence of a subordinating and feasible basis, thus showing that Scarf $\leq_P$ End of the Line, so Scarf is in **PPAD**. We refer to [23] for details.

In Section 5.2, we establish the **PPAD**-hardness of Fractional Hypergraph Matching, which reduces to Scarf in polynomial time [1], thus completing the proof that Scarf is **PPAD**-complete.

5.2. Hypergraphic Preference Systems

A hypergraphic preference system is a pair $(H, \mathcal{O})$, where $H = (V, E)$ is a hypergraph, and $\mathcal{O} = \{\preceq_v : v \in V\}$ is a family of linear orders, $\preceq_v$ being an order on the set of edges containing the vertex v. A set M of edges is called a *stable matching* with respect to the preference system if (a) it is a matching and (b) for every edge e there exists a vertex $v \in e$ and an edge $m \in M$ containing v such that

$e \preceq_v m$. A nonnegative function w on the edges in H is called a *fractional matching* if $\sum_{v \in h} w(h) \leq 1$ for every vertex v. A fractional matching w is called *stable* if every edge e contains a vertex v such that $\sum_{v \in h, e \preceq_v h} w(h) = 1$.

Aharoni and Fleiner [1] used Scarf's lemma to prove that every hypergraphic preference system has a fractional stable matching. This naturally leads to a computational problem – FRACTIONAL HYPERGRAPH MATCHING : given a hypergraphic preference system $(H, \mathcal{O})$, find a fractional stable matching. We first observe that the proof of [1] is a polynomial time reduction from FRACTIONAL HYPERGRAPH MATCHING to SCARF, thus placing it in $\mathbf{PPAD}$. We now show that FRACTIONAL HYPERGRAPH MATCHING is $\mathbf{PPAD}$-hard via a reduction from preference games.

Theorem 5.2. DEGREE d PREFERENCE GAME $\leq_P$ FRACTIONAL HYPERGRAPH MATCHING.

Proof: We are given a preference game over players $[n] = \{1, \ldots, n\}$. We construct a hypergraph matching instance $(H, \mathcal{O})$, $H = (V, E)$. The set V of vertices is $[n] \cup \{i^* : i \in [n]\}$; that is, we have two vertices i and i^* for each player i. The set of edges is given by $\{\{i^*\} : i \in [n]\} \bigcup \{\{i, i^*\} \cup J_i : i \in [n], J_i \subseteq \text{in}(i)\}\}$. (Note that J_i is a subset of players that prefer i over themselves.)

We next describe the linear order for a given vertex i. Let e_1 and e_2 be two edges containing i. By our construction of E, there exists a unique i_1 such that $\{i_1, i_1^*\}$ is a subset of e_1. Similarly, there is a unique i_2 such that $\{i_2, i_2^*\}$ is a subset of e_2. If $i_1 \neq i_2$, then we require that $e_1 \succeq_i e_2$ if and only if $i_1 \succeq_i i_2$. If $i_1 = i_2$, then we require $e_1 \succeq_i e_2$ whenever $e_1 \supseteq e_2$. Finally, for any vertex i^*, we select any linear order in which $e_1 \succeq_{i^*} e_2$ whenever $\{i, i^*\}$ is a subset of e_1 and $e \succeq_{i^*} \{i*\}$ for all e.

Since we are given a preference game of constant degree, the above construction is polynomial time.

Suppose f is a stable fractional matching for the hypergraph preference system. We set w_{ij} (an assignment for the preference game) to be the sum of the weights of edges containing the subset $\{j, j^*, i\}$. This is an equilibrium for the preference game.

Similarly, given any equilibrium for the preference game, we can construct a correponding equilibrium in this hypergraph preference system. For a proof of correctness of this reduction, see [23]. ■

5.3. Cooperative Games with Non-Transferable Utilities

Definition 5.3. *A game with non-transferable utilities over n players is specified by a function V that for each subset S of $N = \{1, 2, \ldots, n\}$ returns a set $V(S)$ of outcomes – each outcome being a vector of utility values, one component for each player in S. A collection T of coalitions is* balanced *if there exists an assignment of reals δ_S for each coalition S in T such that for all v, $\sum_{S: v \in S} \delta_S = 1$. We say that u is* attainable *by S if $u \in V(S)$. A game is* balanced *if and only if for any balanced collection T and any u, if u_S is attainable by all S in T, then u is attainable by N.*

As mentioned earlier, Scarf [34] proved that every balanced game has a nonempty core. We define CORE-BALANCED-NTU below, a natural computational version of this claim. Scarf's proof [34], which is a reduction to SCARF, and Theorem 5.4 establish its $\mathbf{PPAD}$-completeness.

CORE-BALANCED-NTU: The game is specified by a set N of players, a collection $\mathcal{S}$ of proper subsets of N (the coalitions), and for each $S \in \mathcal{S}$, vectors $u_1, \ldots, u_{k_S}$ in $\mathbb{R}^{|S|}$ such that $V(S) = \{u \in \mathbb{R}^{|S|} : \exists j \ u \leq u_j\}$. For a coalition $S \notin \mathcal{S}$, $V(S) = \{0\}^{|S|}$ and $V(N)$ is defined as the set of all u for which there exists a balanced collection T such that u_S is attainable by all S in T. The goal is to find an element in the core.

Theorem 5.4. FRACTIONAL HYPERGRAPH MATCHING $\leq_P$ CORE-BALANCED-NTU.

Proof: Suppose we are given a hypergraph H and for each vertex i, a preference ranking among all edges containing i. We first add, for each vertex i in H, a new vertex i^* and edge $\{i, i^*\}$. We set the preference of i for the edge $\{i, i^*\}$ to be the least among all the edges containing i. Let N denote the new set of nodes and E the new set of edges. For $S \in E$ and $i \in N$, let $r_i(S)$ denote the rank of S in i's preference list, with 0 assigned to the least preferred edge (thus for every i, $r_i(\{i, i^*\}) = 0$). We now define a balanced cooperative game with non-transferable utilities. For each node in N, we have a player in the game. For any coalition S, we consider two cases. If $S \in E$, then we have a single vector $r_S = (r_{i_1}(S), r_{i_2}(S), \ldots, r_{i_{|S|}}(S)$, where $S = \{i_1, i_2, \ldots, i_{|S|}\}$. Note that by definition, if $S \notin E$ and $S \neq N$, then $V(S)$ equals $0^{|S|}$.

For N, note that $V(N)$ is precisely the set of all u such that u_S is attainable by all S in some balanced collection T. We first observe that we can determine in polynomial time whether a given u is in $V(N)$. For each S, if $u \leq r_S$, then we have a variable x_S for S. Now we simply solve the linear program: $\sum_{S: i \in S} x_S = 1$. It is easy to see that the linear program is feasible if and only if u is in $V(N)$.

Proof of correctness for this reduction can be found in [23]. ■

5.4. Fractional Stable Paths Problem

The Fractional Stable Paths problem, introduced in [18], is defined as follows. Let G be a graph with a distinguished destination node d. Each node $v \neq d$ has a list $\pi(v)$ of simple paths from v to d and a preference relation $\succeq_v$ among the paths in $\pi(v)$. For a path S, we also define $\pi(v, S)$ to be the set of paths in $\pi(v)$ that have S as a suffix. A *proper suffix* S of P is a suffix of P such that $S \neq P$ and $S \neq \emptyset$.

A *feasible fractional paths solution* is a set $w = \{w_v : v \neq d\}$ of assignments $w_v : \pi(v) \to [0, 1]$ satisfying: (1)

Unity condition: for each node v, $\sum_{P \in \pi(v)} w_v(P) \leq 1$, and (2) **Tree condition**: for each node v, and each path S with start node u, $\sum_{P \in \pi(v,S)} w_v(P) \leq w_u(S)$. In other words, a feasible solution is one in which each node chooses at most 1 unit of flow to d such that no suffix is filled by more than the amount of flow placed on that suffix by its starting node. A feasible solution w is *stable* if for any node v and path Q starting at v, one of the following holds: **(S1)** $\sum_{P \in \pi(v)} w_v(P) = 1$, and for each P in $\pi(v)$ with $w_v(P) > 0$, $P \succeq_v Q$; or **(S2)** There exists a proper suffix S of Q such that $\sum_{P \in \pi(v,S)} w_v(P) = w_u(S)$, where u is the start node of S, and for each $P \in \pi(v, S)$ with $w_v(P) > 0$, $P \succeq_v Q$. In other words, in a stable solution: if node v has not fully chosen paths that it prefers at least as much as Q, then it has completely filled path Q by filling some suffix with paths it prefers at least as much as Q.

We define a computational version, FRACTIONAL SPP: given an instance of the fractional stable paths problem, find a fractional stable solution.

Theorem 5.5. PREFERENCE GAME $\leq_P$ FRACTIONAL SPP $\leq_P$ PERSONALIZED EQUILIBRIUM.

Proof: We list the reductions below, but defer proofs of correctness to the full version ([23]).

PREFERENCE GAME $\leq_P$ FRACTIONAL SPP. Given a preference game over player set $[n]$, including preference relation $\succeq_i$ for all $i \in 1 \dots n$. We will convert this into a fractional stable paths problem. Create a node v_i for each i. Also create a universal destination node d. For all i, define P_{ii} = the path (v_i, d). For all i, j, define P_{ij} = the path (v_i, v_j, d). Let $\pi(v_i)$ (the set of preferred paths for v_i) = $\{P_{ij} : j \succeq_i i\}$. If $k \succeq_i j$, then $P_{ik} \succeq_i P_{ij}$. Let $w_i(j)$ refer to the amount of weight placed by node v_i on path P_{ij} in a fractional SPP solution, and let $w_i(i)$ be the amount of weight placed by i on path P_{ii}. w is a fractional stable paths solution if and only if w defines an equilibrium of the preference game.

FRACTIONAL SPP $\leq_P$ PERSONALIZED EQUILIBRIUM. Suppose we are given an instance of FRACTIONAL SPP, consisting of a set of nodes V, a set of preferred paths $\pi(v)$ for all $v \in V$, and a preference relation $\succeq_v$ for each set $\pi(v)$. We can also find $\pi(v, S)$, the set of all $P \in \pi(v)$ such that S is a subpath of P. Let $q_v(P)$ = the number of paths Q such that $P \succeq_v Q$. We will create the following instance of PERSONALIZED EQUILIBRIUM. The set of players is V. The set of strategies S_v for a node V is $\pi(v) \cup \{N\}$ (N stands for "No path"). For node v, there is exactly one edge defined for each strategy. Edge P' for strategy $P = \{S : P \in \pi(v, S)\}$. The edge for strategy N (N') is a singleton edge, containing only that strategy. The payoffs to player v are: $u_v(P') = q_v(P) + 1$, $u_v(N) = 1$. Suppose w is a set of weights in a personalized equilibrium of the game defined above. $w_v(P')$ represents the weight assigned by v to edge P'. w is a personalized equilibrium if and only if $w' : w'_v(P) = w_v(P')$ is a fractionally stable solution to the FRACTIONAL SPP instance. ■

In the full version ([23]), we extend the proof of $\mathbf{PPAD}$-completeness to cases where path preferences are based on shortest path lengths as well as to approximate FRACTIONAL SPP.

5.5. *Additional Problems*

In the full version ([23]), we also show two additional problems are $\mathbf{PPAD}$-complete using reductions from PREFERENCE GAME and reductions to PERSONALIZED EQUILIBRIUM. The first of these, STRONG KERNEL, is the problem of finding a strong fractional kernel in a clique-acyclic digraph with largest clique of constant size. A kernel of a directed graph is a subset of vertices that is both independent and dominating. Fractional kernels [2] relax this concept to allow nodes to fractionally belong to a kernel. The second additional problem, FRACTIONAL BBC, is the problem of finding an equilibrium in a fractional Bounded Budget Connection game – a network connection game in which nodes in a graph may spend up to a fixed budget to fractionally purchase edges to other nodes with the goal of achieving a small minimum-cost flow to a destination node.

ACKNOWLEDGMENT

We would like to thank H. Venkateswaran for helpful discussions, Gordon Wilfong for sending a preprint of [19], Christos Papadimitriou for pointing us to [37], and David Karger for his insightful motivation for the preference game.

REFERENCES

[1] R. Aharoni and T. Fleiner, "On a lemma of Scarf," *Journal of Combinatorial Theory, Series B*, vol. 87(1), pp. 72–80, 2003.

[2] R. Aharoni and R. Holzman, "Fractional Kernels in Digraphs," *Journal of Combinatorial Theory, Series B*, vol. 73(1), pp. 1–6, 1998.

[3] R. J. Aumann, "Subjectivity and correlation in randomized strategies." *Journal of Mathematical Economics*, vol. 1, pp. 67–96, 1974.

[4] O. Bondareva., "Some applications of linear programming to cooperative games." *Problemy Kibernetiki*, vol. 10, pp. 119–139, 1963.

[5] X. Chen and X. Deng, "3-NASH is PPAD-complete," *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 134, 2005.

[6] ——, "Settling the complexity of two-player Nash equilibrium," *FOCS*, pp. 261–272, 2006.

[7] X. Chen, X. Deng, and S.-H. Teng, "Computing Nash equilibria: Approximation and smoothed complexity," *FOCS*, pp. 603–612, 2006.

[8] ——, "Settling the complexity of computing two-player Nash equilibria." *Journal of the ACM (JACM)*, vol. 56(3), 2009.

[9] V. Conitzer and T. Sandholm, "Complexity of determining nonemptiness of the core." *EC*, pp. 230–231, 2003.

[10] ——, "Complexity of constructing solutions in the core based on synergies among coalitions." *Artificial Intelligence*, vol. 170(6), pp. 607–619, 2006.

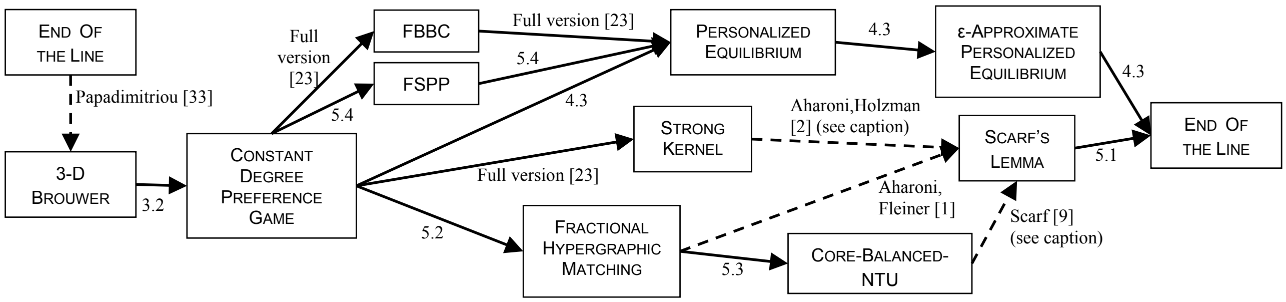

Figure 1. We show these problems to be **PPAD**-complete. Each reduction line is labeled with the Section or Citation where the reduction can be found. Two of these reductions, STRONG KERNEL $\leq_P$ SCARF and CORE-BALANCED-NTU $\leq_P$ SCARF, are only polynomial time reductions for the specific versions of the problems discussed in this paper. In our definition of STRONG KERNEL, we assume that the largest clique in the graph has constant size, since otherwise it is not clear whether the problem is even in **TFNP**. CORE-BALANCED-NTU, as defined in Section 5.3, assumes that the game description explicitly lists the possible coalitions and their Pareto-optimal outcomes.

[11] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, "The complexity of computing a Nash equilibrium," *STOC*, pp. 71–78, 2006.

[12] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou., "The complexity of computing a Nash equilibrium." 2008, invited to *SiCOMP*.

[13] Q. Fang, R. Fleischer, J. Li, and X. Sun, "Algorithms for core stability, core largeness, exactness, and extendability of flow games." *13th Annual International Computing and Combinatorics Conference*, 2007.

[14] D. Gale and L. S. Shapley, "College admissions and the stability of marriage," *American Mathematical Monthly*, vol. 69, pp. 9–14, 1962.

[15] R. Garg, V. Kumar, A. Rudra, and A. Verma, "Coalitional games on graphs: core structure, substitutes and frugality." *EC*, pp. 248–289, 2003.

[16] P. W. Goldberg and C. H. Papadimitriou, "Reducibility among equilibrium problems," *STOC*, pp. 61–70, 2006.

[17] T. G. Griffin, F. B. Shepherd, and G. Wilfong, "The stable paths problem and interdomain routing," *IEEE/ACM Transactions on Networking*, vol. 10(2), pp. 232–243, 2002.

[18] P. E. Haxell and G. T. Wilfong, "A fractional model of the Border Gateway Protocol (BGP)," *SODA*, pp. 193–199, 2008.

[19] ——, "On the Stable Paths Problem," 2008, preprint.

[20] N. Immorlica, K. Jain, and M. Mahdian, "Game-theoretic aspects of designing hyperlink structures." *WINE*, vol. 4286, pp. 150–161, 2006.

[21] M. Kearns, M. L. Littman, and S. Singh., "Graphical models for game theory." *17th Conference in Uncertainty in Artificial Intelligence*, pp. 253–260, 2001.

[22] S. Kintali, "A Distributed Protocol for Fractional Stable Paths Problem," *Georgia Tech, College of Computing Technical Report GT-CS-08-06; Presented at the DIMACS/DyDAn Workshop on Secure Internet Routing, Rutgers University*, March 24-26 2008.

[23] S. Kintali, L. J. Poplawski, R. Rajaraman, R. Sundaram, and S.-H. Teng, "Reducibility among fractional stability problems," *arXiv:0904.1435v1 [cs.CC]*, 2009.

[24] T. Király and J. Pap, "Kernels, stable matchings, and Scarf's lemma," *The Egerváry Research Group Technical Report TR-2008-13, http://www.cs.elte.hu/egres/tr/egres-08-13.pdf*.

[25] M. Kumabe and H. R. Mihara, "Computability of simple games: A characterization and application to the core." *arXiv:0705.3227v2*, 2007.

[26] N. Laoutaris, L. J. Poplawski, R. Rajaraman, R. Sundaram, and S.-H. Teng, "Bounded budget connection (BBC) games or how to make friends and influence people, on a budget." *PODC*, pp. 165–174, 2008.

[27] E. Markakis and A. Saberi., "On the core of the multicommodity flow game." *EC*, pp. 93–97, 2003.

[28] N. Megiddo and C. H. Papadimitriou, "On total functions, existence theorems and computational complexity," *Theoretical Computer Science*, vol. 81(2), pp. 317–324, 1991.

[29] J. Nash., "Equilibrium point in n-person games." *PNAS*, vol. 36(1), pp. 48–49, 1950.

[30] ——, "Noncooperative games." *Annals of Mathematics*, vol. 54, pp. 286–295, 1951.

[31] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[32] C. H. Papadimitriou, "On the complexity of the parity argument and other inefficient proofs of existence," *Journal of Computer and System Sciences*, vol. 48(3), pp. 498–532, 1994.

[33] Y. Rehkter and T. Li, "A Border Gateway Protocol (BGP version 4) RFC 1771," 1995.

[34] H. E. Scarf, "The Core of an N Person Game," *Econometrica*, vol. 69, p. 35:50, 1967.

[35] L. S. Shapley., "On balanced sets and cores." *Naval Research Logistics Quarterly*, vol. 14, pp. 453–460, 1967.

[36] J. W. Stewart, *BGP4: Inter-domain routing in the Internet.* Addison Wesley, 1998.

[37] M. J. Todd., "Orientation in complementary pivot algorithms." *Mathematics of Operations Research*, vol. 1(1), pp. 54–66, 1976.

[38] K. Varadhan, R. Govindan, and D. Estrin, "Persistent route oscillations in inter-domain routing," *Computer Networks*, vol. 32, pp. 1–16, 2000.