Egerváry Research Group
on Combinatorial Optimization

Technical reportS

TR-2012-18. Published by the Egerváry Research Group, Pázmány P. sétány 1/C, H–1117, Budapest, Hungary. Web site: www.cs.elte.hu/egres. ISSN 1587–4451.

Equilibria and fairness in linear service-providing markets

Tamás Király and Júlia Pap

November 2012
Revised May 2013

Equilibria and fairness in linear service-providing markets

Tamás Király* and Júlia Pap**

Abstract

We investigate equilibria in multi-agent scenarios where agents interact by providing services to each other. Possible strategies and costs of each agent are assumed to be representable by a linear program, while the connection between strategies of different agents is established by the constraint that the total amount provided of a given service must be equal to the total demand of customers. Payments between agents are determined by service prices. Special cases of this model include the service network alliance model of Agarwal and Ergun and the multiplayer multicommodity flow model introduced by the present authors with co-authors.

We study the existence and computational complexity of fair prices and equilibria. We show that it is always possible to find fair service prices guaranteeing that socially optimal solutions are in equilibrium. For the case of fixed service prices, we prove that equilibria always exist under some natural assumptions, but are PPAD-complete to find. We also give a polynomial algorithm for the case when the digraph containing arcs from service providers to customers has the property that every strong component is a simple cycle. The proof is based on a new algorithmic result on approximating a fixed point of a multidimensional function having a cyclic structure.

1 Introduction

The model introduced in this paper originates in previous results on multicommodity flow games. The study of flow problems with selfish agents dates back to the papers of Kalai and Zemel [9, 10], and the multicommodity version was studied by Derks and Tijs [5]. Network formation problems involving bilateral service agreements were studied in [7, 2]. Of particular relevance to our model is the work of Agarwal and Ergun [1] who studied the mechanism design aspects of multicommodity flow games. In their model, each agent owns a given fraction of each arc of the network. An agent also has a set of demands, and revenue is generated by satisfying a fraction of these

*MTA-ELTE Egerváry Research Group, Eötvös Loránd University, Budapest, Hungary. Email: tkiraly@cs.elte.hu

**Department of Operations Research, Eötvös Loránd University, Budapest, Hungary. Email: papjuli@cs.elte.hu

demands. Instead of the usual concept of transferable utilites, only a restricted form of transfer is allowed: agents pay specified capacity exchange prices for the use of arc capacities, divided proportionally among the owners of the arc. Agarwal and Ergun proposed a method for determining capacity exchange prices which provide incentives for agents to route according to the optimal flow. The results were further generalized in [6].

Another precedent to our model is the multiplayer multicommodity flow problem introduced by the present authors with co-authors [3]. As in the aforementioned model, each agent controls a subnetwork. Here agents have hard demands that must be satisfied, and instead of the exchange of arc capacities, there are bilateral agreements specifying source–destination pairs between which one agent undertakes to route the traffic of the other in exchange for a specified per-unit payment. The agents who undertake the routing can freely choose the route in their own network.

We give a common generalization of these two approaches by introducing an abstract linear model of service-providing markets, where subsets of agents provide services to each other. Payments between agents depend linearly on service prices and the amount of service provided. We show that it is always possible to determine service prices that are fair and guarantee that socially optimal solutions are in equilibrium. If the prices are fixed arbitrarily, then we show that the existence of an equilibrium is guaranteed under some natural conditions, but finding an equilibrium is PPAD-complete. Finally, we give a polynomial-time algorithm for finding an equilibrium in a restricted class of instances, using a new algorithmic result on approximating a fixed point of a multidimensional function having a cyclic structure.

1.1 Description of the model

Agents are numbered from 1 to n, so the set of agents is $[n] = \{1, \ldots, n\}$, and S denotes the set of all services. The set of services where agent i is a possible provider (resp. customer) is denoted by S^i (resp. T^i). We do not require S^i and T^i to be disjoint.

The description of the possible strategies of agent i uses three types of variables:

- x^i_s, for $s \in S^i$, represents the amount of service s provided,
- y^i_s, for $s \in T^i$, is the amount bought of service s,
- z^i is a vector of additional variables that may take negative values.

To agent i we associate a bounded polyhedron P^i of the form

$$A^i_1 x^i + A^i_2 y^i + A^i_3 z^i \le b^i, \tag{1}$$
$$x^i, y^i \ge 0, \tag{2}$$

where the matrices A^i_1, A^i_2, A^i_3 can be arbitrary except that A^i_1 is assumed to be non-negative. This bounded polyhedron describes the set of possible strategies of agent i. Let $P_\times$ denote the direct product $P^1 \times \cdots \times P^n$.

There is a cost vector c^i of the same dimension as (x^i, y^i, z^i). The *social cost* of a strategy $(x, y, z) \in P_\times$ is $c^\mathsf{T}(x, y, z)$. We use the notation c^i_s for the cost of variable x^i_s.

Service s has a per unit *service price* p_s, sometimes considered as a variable and sometimes as a constant. A customer j with $s \in T^j$ pays $p_s y^j_s$ for service s, and the income of provider i from service $s \in S^i$ is $p_s x^i_s$.

1.2 Feasible solutions

A strategy vector $(x, y, z) \in P_\times$ is *feasible* if

$$\sum_{i:s\in S^i} x^i_s = \sum_{j:s\in T^j} y^j_s \quad \text{for every } s \in S.$$

The above equation is the *service equation* for service s. The polyhedron of feasible solutions is denoted by P_{feas}. A *socially optimal* solution is a feasible solution (x, y, z) that is optimal for the cost vector c.

If the service prices are fixed, then we can define a *modified cost vector* c_p by decreasing the cost of variable x^i_s by p_s for all $s \in S^i$, and increasing the cost of variable y^i_s by p_s for all $s \in T^i$. The personal interest of agent i is to minimize c^i_p on P^i. Of course, if each agent optimizes independently on their polyhedra, then the resulting vector is typically not feasible.

We also consider the setting where each provider of a service must contribute a fixed proportion of the service. Here we are given non-negative rationals r^i_s for every service s and agent i with $s \in S^i$, satisfying $\sum_{i:s\in S^i} r^i_s = 1$. These values are called *service ratios*. A vector $(x, y, z) \in P_\times$ is called *fixed-ratio feasible* if

$$x^i_s = r^i_s \sum_{j:s\in T^j} y^j_s \quad \text{for every } i \text{ and } s \in S^i.$$

This equation is called the *personal service equation* of agent i for service s.

Clearly, not every instance of the problem has a feasible or fixed-ratio feasible solution; however, there is a condition that turns out to be sufficient for feasibility. An instance is called *safe* if for any $(x, y, z) \in P_\times$ there is a vector $(\overline{x}, \overline{y}, \overline{z})$ in $P_\times$ such that

$$\sum_{i:s\in S^i} \overline{x}^i_s = \sum_{j:s\in T^j} y^j_s \quad \text{for every } s \in S.$$

An instance is *fixed-ratio safe* with respect to service ratios r if for any $(x, y, z) \in P_\times$ and any $i \in [n]$ we can replace (x^i, y^i, z^i) by a vector $(\overline{x}^i, \overline{y}^i, \overline{z}^i) \in P^i$ such that the resulting set of vectors satisfies all personal service equations of agent i. We will see that safe instances have solutions in equilibrium, which implies that they have feasible solutions.

1.3 Relation to previous models

1.3.1 Service network alliances

Agarwal and Ergun [1] gave the following model for multicommodity flow games in service network alliances. There are n agents, and a common directed graph $D = (V, A)$. Each arc $a \in A$ has a capacity u_a, and given ownership ratios r_a^i ($i \in [n]$) with $\sum_{i=1}^n r_a^i = 1$. Each agent i has a demand set Q^i, where a demand $q \in Q^i$ is characterized by a source s_q, a sink t_q, a per-unit revenue w_q, and an upper bound u_q. Agent i should have a flow of size $f_q \leq u_q$ from s_q to t_q; the revenue generated by this flow is $w_q f_q$. The total flow traversing an arc a should not exceed the capacity u_a.

The main result of [1] is a mechanism that distributes the benefits of collaboration by assigning a capacity exchange price p_a to each arc $a \in A$. If the total flow of agent i on arc A is f_a^i, then i has to pay an amount of $p_a f_a^i$, which is distributed among the agents according to the ratios r_a^j ($j \in [n]$). Since the payment to herself can be ignored, agent i actually pays an amount of $(1 - r_a^i) p_a f_a^i$ to the other agents. The paper shows that, given a socially optimal multicommodity flow f^*, it is possible to compute capacity exchange prices with the property that no agent is motivated to deviate from f^*.

We can model this problem in the framework of the present paper by assigning a service to each arc $a \in A$, i.e. $S = A$. All agents are potential customers of all services, while the providers of service a are the agents with $r_a^i > 0$. The service ratios are determined by the values r_a^i. In order to define the polyhedron P^i, we consider the vector variable z^i to be composed of vector variables z_q^i for each demand $q \in Q^i$. In P^i, the variables z_q^i ($q \in Q^i$) describe the multicommodity flow polyhedron of agent i, with cost defined as the opposite of revenue. The variable y_a^i represents the total flow of agent i on arc a, upper bounded by u_a, and variable x_a^i has the single constraint $0 \leq x_a^i \leq r_a^i u_a$; these variables have cost 0.

With these definitions, a fixed-ratio feasible solution is a feasible solution of the original problem, and the c_p^i-cost of agent i represents his total cost in the original problem. Note that the instances we obtain this way are typically not safe or fixed-ratio safe.

1.3.2 The Multiplayer multicommodity flow problem (MMFP)

In the MMFP problem defined in [3], the n agents have separate networks $D^i = (V, A^i)$ on a common node set. Each arc a has a cost c_a and a capactity u_a. Each agent i has a set Q^i of hard demands. A demand $q \in Q^i$ is characterized by a source s_q, a sink t_q and a size d_q. The aim of agent i is to satisfy all his demands at minimum cost. A subset of arcs $B^i \subseteq A^i$ are so-called *contractual arcs*, each with a designated agent called the provider. A contractual arc has a price p_a (either variable or fixed) and a multiplier γ_a. If $a = uv \in B^i$ and the total flow of agent i on a is f_a^i, then an additional demand of size $\gamma_a f_a^i$ from u to v appears in the network of the provider of a (a *contractual demand*). In exchange, i pays an amount of $f_a^i p_a$ to the provider. It is proved in [3] that in safe instances of the problem there always exists an equilibrium.

The MMFP problem can be modeled in our framework by assigning a service to each contractual arc, thus each service has one customer and one provider. Let R^i be the set of contractual demands of agent i. We consider the vector variable z^i to be composed of vector variables z^i_q for each demand $q \in Q^i \cup R^i$. In P^i, the variables z^i_q describe the multicommodity flow polyhedron of agent i, with the modification that if $q \in R^i$, then the size of demand q is $\gamma_{a(q)} x^i_{a(q)}$, where $a(q)$ is the contractual arc corresponding to q. The variable y^i_a for a contractual arc $a \in B_i$ is equal to the total flow of agent i on arc a.

1.4 Definition of equilibrium and fairness

If both the service prices and the service ratios are fixed, then a fixed-ratio feasible solution (x, y, z) is said to be a *fixed-ratio equilibrium* if there is no $i \in [n]$ and $(\overline{x}^i, \overline{y}^i, \overline{z}^i) \in P^i$ such that

$$c_p^{i\mathsf{T}}(\overline{x}^i, \overline{y}^i, \overline{z}^i) < c_p^{i\mathsf{T}}(x^i, y^i, z^i)$$

and the set of vectors obtained by replacing (x^i, y^i, z^i) by $(\overline{x}^i, \overline{y}^i, \overline{z}^i)$ satisfies all personal service equations of agent i. This means that if the solutions of other agents are fixed, then agent i cannot increase his profit without violating one of his personal service equations. It is important to note that this kind of equilibrium is stronger than Nash equilibrium, since the violation of the personal service equations of others is permitted.

Let us now consider the setting when the service prices are fixed but the service ratios are not. Defining equilibrium based on individual interests of agents is problematic, because it is not clear whether agents should be allowed to change the amount of service they provide. Instead, we define a kind of collective equilibrium, as follows. A feasible solution $(x, y, z) \in P_{\text{feas}}$ is a *collective equilibrium* if there is no $(\overline{x}, \overline{y}, \overline{z}) \in P_\times$ such that $c_p^\mathsf{T}(\overline{x}, \overline{y}, \overline{z}) < c_p^\mathsf{T}(x, y, z)$ and

$$\sum_{i:s \in S^i} \overline{x}^i_s = \sum_{i:s \in S^i} x^i_s \quad \text{for every } s \in S.$$

In a collective equilibrium, if a subset of agents consider the amount of service they have to provide as given, then they cannot change their strategies to earn more profit, not even by violating service equations.

Finally, if the service prices are not fixed, then a service price vector p is called *fair* for a feasible solution $(x, y, z) \in P_{\text{feas}}$ if the vector (x, y, z) is optimal for the cost vector c_p in the polyhedron $P_\times$. The motivation for this definition is that the following hold if p is fair for (x, y, z):

- (x, y, z) is a collective equilibrium at prices p,
- (x, y, z) is a fixed-ratio equilibrium at prices p with respect to the service ratios determined by (x, y, z) (in fact, agents cannot improve even by violating their own personal service equations),
- no agent can increase its profit by terminating or restricting a service,

- no agent can increase its profit by persuading others to purchase more of a service.

In other words, if prices are fair for a solution, then the solution is in equilibrium, while the prices are high enough so that agents are not motivated to restrict a service, and also low enough so that no one is motivated to decrease prices in order to get more customers.

An easy observation is that fair prices are possible only for socially optimal solutions; indeed, the c_p-cost of a *feasible* solution is the same as its c-cost, so a socially suboptimal feasible solution cannot be optimal in $P_\times$ with respect to c_p.

1.5 Summary of the results

In Section 2 we show that there are service prices p that are fair for every socially optimal solution. This price vector p can be computed in polynomial time using linear programming.

Theorem 1.1. *It is possible to compute service prices p in polynomial time that are fair for every socially optimal solution (x, y, z). If $c^i_s \geq 0$ for every i and every $s \in S^i$, then the price p_s can be negative only if $x_s = \mathbf{0}$ in every socially optimal solution (x, y, z).*

For fixed service ratios, this result implies that if there exists a fixed-ratio feasible solution that is socially optimal in P_{feas}, then there is a price vector p such that any socially optimal fixed-ratio feasible solution is a fixed-ratio equilibrium.

In Section 3 we prove the following existence results for fixed service prices.

Theorem 1.2. *In a safe instance there always exists a collective equilibrium.*

Theorem 1.3. *In a fixed-ratio safe instance there is always a fixed-ratio equilibrium.*

However, finding an equilibrium is PPAD-complete even in the MMFP model described in Section 1.3.2 (here the notions of collective equilibrium and fixed-ratio equilibrium coincide, since every service has only one provider and one customer). This is proved in Section 4.

Theorem 1.4. *It is PPAD-complete to find an equilibrium in safe instances of MMFP.*

An interpretation of this result is that in the fixed price case it is unlikely that any mechanism can steer agents quickly to an equilibrium, since the computation of an equilibrium is intractable unless all problems in PPAD can be solved in polynomial time.

In Section 5 we present a polynomial time algorithm for finding a fixed-ratio equilibrium in the special case when each strong component of the digraph representing the provider-customer relationships is a simple cycle. To be more precise, this auxiliary directed graph $D^* = ([n], A^*)$ has arcs from the providers of each service to the customers, with possible parallel arcs but excluding loops.

Theorem 1.5. *A fixed-ratio equilibrium can be found in polynomial time in fixed-ratio safe instances where every strong component of D^* is a simple directed cycle.*

In view of this result, the study of other special cases might offer new insights on the borderline between polynomially solvable and PPAD-complete problems. For example, it is open whether there is a polynomial algorithm in case of 3 services, with 1 provider and 1 customer for each.

Our polynomial algorithm can be seen as a more general result on the approximation of fixed points in a certain class of fixed-point problems. Given m interval-valued mappings $\varphi_1, \dots, \varphi_m$ on the unit interval, all with the closed graph property, Kakutani's fixed point theorem implies that there is a vector x such that $x_{i+1} \in \varphi_i(x_i)$ $(i = 1, \dots, m-1)$ and $x_1 \in \varphi_m(x_m)$ (a *cyclically fixed vector*). We show an algorithm for finding m arbitrarily small intervals such that their direct product contains a cyclically fixed vector.

Theorem 1.6. *Let $\varphi_i : [0,1] \to \mathcal{P}([0,1])$ $(i \in [m])$ be given as above with a function evaluation oracle, and let $0 < \epsilon < 1$. In $O(m^2 \log(\frac{1}{\epsilon}))$ steps we can find intervals $I_1, \dots, I_m \subseteq [0,1]$ of length at most ϵ such that there is a cyclically fixed vector x with $x_i \in I_i$ $(i \in [m])$.*

2 Existence of fair prices

In this section we show that there is a price vector p that is fair for any socially optimal solution, and such prices can be found in polynomial time. We start by explaining the polyhedral tools used in the proof.

For a polyhedron P and a face F of it, let opt.cone(F, P) denote the set of objective vectors c for which every point of F is optimal in P, that is, the *optimal cone* of F in P. The *tangent cone* of a point in P is the set of feasible directions from the point. The relative interior of a set $X \subseteq \mathbb{R}^n$ is denoted by relint(X), while lin$X$ is the linear translation of the affine hull of X. We will need the following lemma on optimal cones.

Lemma 2.1. *Let P_1 be a polyhedron, Π an affine subspace and $P_2 = P_1 \cap \Pi$. Let F_2 be a face of P_2 and let F_1 be the smallest face of P_1 that contains F_2. Then*

(i) $\text{opt.cone}(F_2, P_2) = \text{opt.cone}(F_1, P_1) + (\text{lin}\Pi)^\perp$,

(ii) $\text{relint}(\text{opt.cone}(F_2, P_2)) = \text{relint}(\text{opt.cone}(F_1, P_1)) + (\text{lin}\Pi)^\perp$.

Proof. To prove the "$\supseteq$" containment in part (i), let $w \in \text{opt.cone}(F_1, P_1)$ and $a \in (\text{lin}\Pi)^\perp$. Then for any $x \in F_2$ and $x' \in P_2$, $(w+a)x = wx + ax \geq wx' + ax = (w+a)x'$, so $w + a \in \text{opt.cone}(F_2, P_2)$.

For the "$\subseteq$" containment suppose that $w \in \text{opt.cone}(F_2, P_2)$ but w is not in the cone $\text{opt.cone}(F_1, P_1) + (\text{lin}\Pi)^\perp$. By Farkas' Lemma the latter implies that there is a vector y for which $wy > 0$ but $(w' + a)y \leq 0$ for every $w' \in \text{opt.cone}(F_1, P_1)$ and $a \in (\text{lin}\Pi)^\perp$. Clearly $y \in \text{lin}\Pi$. Let x^* be a vector in relint(F_2). Then opt.cone(F_1, P_1) is generated by the normal vectors of the facets of P_1 that x^* satisfies with equality.

This means that y is in the tangent cone of P_1 in x^*, thus, since $y \in \text{lin}\Pi$, y is also in the tangent cone of P_2 in x^*. This contradicts $wy > 0$.

Part (ii) follows from part (i). □

Proof of Theorem 1.1. We assume that P_{feas} is non-empty. Let F_{socopt} be the set of the socially optimal solutions, that is, the optimal face in P_{feas} minimizing the cost c. Let $F_\times$ be the minimal face of $P_\times$ which contains F_{socopt}. Let furthermore H be the subspace determined by the service equations.

We apply Lemma 2.1 with $P_1 = P_\times$, $\Pi = H$, $P_2 = P_{\text{feas}}$, $F_1 = F_\times$, and $F_2 = F_{\text{socopt}}$. By consequence, there is a vector $h \in (\text{lin}H)^\perp$ for which $c + h$ is a member of $\text{relint}(\text{opt.cone}(F_\times, P_\times))$. Note that h can be computed in polynomial time by linear programming. By the definition of H, there is a vector $p \in \mathbb{R}^S$ such that

- The component of h corresponding to x^i_s is $-p_s$,
- the component of h corresponding to y^i_s is p_s,
- the components of h corresponding to z are 0.

Let the price of service s be p_s. Since $c + h = c_p$, it follows that any socially optimal solution $(x, y, z) \in F_{\text{socopt}} \subseteq F_\times$ is optimal for objective function c_p in the polyhedron $P_\times$. That is, p is fair for any socially optimal solution $(x, y, z) \in F_{\text{socopt}}$.

To prove the second part of the theorem, assume that c^i_s is nonnegative for every i and $s \in S^i$. If p_s is negative and $s \in S^i$, then $c^i_s - p_s$ is positive, so by decreasing x^i_s we decrease the c_p-cost. Since the describing matrices A^i_1 are nonnegative, the modified vector is also in $P_\times$ if $x^i_s \geq 0$. Therefore $p_s < 0$ implies that x^i_s must be 0 in a socially optimal solution. □

3 Existence of equilibrium in safe instances

In general, the set of equilibrium solutions can be empty even if P_{feas} is non-empty. In order to show the existence of equilibria in safe instances (Theorems 1.2 and 1.3) we resort to the following fundamental fixed-point theorem of Kakutani.

Theorem 3.1 (Kakutani [8]). *Let C be a compact convex set in $\mathbb{R}^d$ and let $\varphi : C \to \mathcal{P}(C)$ be a set-valued function with the following properties:*

- *$\varphi(x)$ is a nonempty convex subset of C for every $x \in C$,*
- *the graph of φ is closed.*

Then there exists a fixed point of φ, that is, an element $x \in C$ for which $x \in \varphi(x)$.

First we consider collective equilibria. For $\xi \in \mathbb{R}^S_+$, we define the affine subspace

$$\Phi(\xi) = \{(x, y, z) : \sum_{i:s \in S^i} x^i_s = \xi_s \quad \text{for every } s \in S\}.$$

Given $(x, y, z) \in P_\times$, let $x^\Sigma \in \mathbb{R}^S_+$ be the vector defined by $x^\Sigma_s = \sum_{i:s\in S^i} x^i_s$, and similarly let $y^\Sigma \in \mathbb{R}^S_+$ be defined by $y^\Sigma_s = \sum_{i:s\in T^i} y^i_s$. We need the following observation.

Proposition 3.2. *The set of vectors $(x,y,z) \in P_\times$ that minimize c_p in $P_\times \cap \Phi(x^\Sigma)$ is the union of some faces of $P_\times$.*

Proof. Let $(x,y,z) \in P_\times$ be a vector with that property, and let F_2 be the minimal face of $P_\times \cap \Phi(x^\Sigma)$ containing it. If we apply Lemma 2.1 with $P_1 = P_\times$, $\Pi = \Phi(x^\Sigma)$, and F_2, we obtain that $c_p \in \text{opt.cone}(F_2, P_\times \cap \Phi(x^\Sigma)) = \text{opt.cone}(F_1, P_\times) + (\text{lin}\Phi(x^\Sigma))^\perp$, where F_1 is the minimal face of $P_\times$ containing (x,y,z). Observe that for any $(\overline{x}, \overline{y}, \overline{z}) \in F_1$ we have $(\text{lin}\Phi(\overline{x}^\Sigma))^\perp = (\text{lin}\Phi(x^\Sigma))^\perp$, and therefore $c_p \in \text{opt.cone}(F_1, P_\times) + (\text{lin}\Phi(\overline{x}^\Sigma))^\perp \subseteq \text{opt.cone}((\overline{x}, \overline{y}, \overline{z}), P_\times \cap \Phi(\overline{x}^\Sigma))$. □

The social optimum can be obtained by minimizing the cost over the polyhedron P_{feas}, thus the socially optimal solutions form a face of it. On the other hand, the set of equilibrium solutions for fixed prices is not necessarily convex (not even connected), but the following holds.

Proposition 3.3. *The set of collective equilibrium solutions is the (perhaps empty) union of some faces of P_{feas}.*

Proof. A feasible solution $(x,y,z) \in P_{\text{feas}}$ is a collective equilibrium if and only if it is optimal in $P_\times \cap \Phi(x^\Sigma)$ for the objective function c_p. By Proposition 3.2, the set of not necessarily feasible vectors $(x,y,z) \in P_\times$ with this property form the union of some faces of $P_\times$. Since the intersection of a face of $P_\times$ and H (the subspace determined by the service equations) is a face of P_{feas}, the proposition follows. □

Now we are ready to prove the existence of a collective equilibrium in a safe instance.

Proof of Theorem 1.2. Let $C = \{\xi \in \mathbb{R}^S_+ : \exists (x,y,z) \in P_\times \text{ s.t. } y^\Sigma = \xi\}$ and let

$$\varphi(\xi) = \{\zeta \in \mathbb{R}^S_+ : \exists (x,y,z) \text{ optimal for } c_p \text{ in } P_\times \cap \Phi(\xi) \text{ s.t. } y^\Sigma = \zeta\}.$$

Since the instance is safe, $\varphi(\xi)$ is nonempty if $\xi \in C$. The set $\varphi(\xi)$ is the projection of a face of $P_\times \cap \Phi(\xi)$, so it is convex. The graph of φ is closed since it is the projection of $\{(x,y,z) : \ (x,y,z) \text{ is optimal in } P_\times \cap \Phi(x^\Sigma) \text{ for } c_p\}$, which is the union of some faces of $P_\times$ by Proposition 3.2. So by Theorem 3.1 there exists a fixed point ξ^*, for which – by the definition of φ – there exists $(x,y,z) \in P_\times$ such that $x^\Sigma = y^\Sigma = \xi^*$ and (x,y,z) is optimal for c_p in $P_\times \cap \Phi(\xi^*)$. The former implies that (x,y,z) is feasible, so by the latter it is a collective equilibrium. □

In order to prove the analogous result for fixed-ratio equilibrium, we have to consider the personal service equations of the agents.

Proof of Theorem 1.3. Let $C = \{y : \ \exists x, z \text{ s.t. } (x,y,z) \in P_\times\}$ and let

$$\begin{aligned} \varphi(y) = \{\overline{y} : \forall i \in [n] \ \exists \overline{x}^i, \overline{z}^i & \text{ such that} \\ & (\overline{x}^i, \overline{y}^i, \overline{z}^i) \text{ is optimal for } c_p \text{ in } P^i \text{ among vectors that satisfy} \\ & \overline{x}^i_s = r^i_s \ (\ \overline{y}^i_s + \sum_{j \neq i : s \in T^j} y^j_s \ ) \quad \text{if } s \in S^i \cap T^i \quad \text{and} \qquad (3) \\ & \overline{x}^i_s = r^i_s \sum_{j : s \in T^j} y^j_s \quad \text{if } s \in S^i \setminus T^i\}, \qquad (4) \end{aligned}$$

in other words, $(\overline{x}^i, \overline{y}^i, \overline{z}^i)$ is optimal for c_p among the vectors in P^i that satisfy the personal service equations of i with respect to $(y^j)_{j\in[n]\setminus\{i\}}$. For $y \in C$, let

$$\Phi_i(y) = \{(\overline{x}^i, \overline{y}^i, \overline{z}^i) : \overline{x}^i, \overline{y}^i \text{ satisfy (3) and (4)}\}.$$

Note that the dimension of this affine subspace does not depend on the parameters. The set $\varphi(y)$ is nonempty because the instance is fixed-ratio safe, and it is convex because it is the projection of the direct product of faces of the polyhedra $P^i \cap \Phi_i(y)$ $(i \in [n])$. It remains to show that the graph of φ is closed. If a_k is a convergent sequence in C with $\lim_{k\to\infty} a_k = a$, then $\lim_{k\to\infty} \Phi_i(a_k) = \Phi_i(a)$, and thus $\lim_{k\to\infty}(\text{lin}\Phi_i(a_k))^\perp = (\text{lin}\Phi_i(a))^\perp$. It follows from Lemma 2.1 that if $b_k \in \varphi(a_k)$ and $\lim_{k\to\infty} b_k = b$, then $b \in \varphi(a)$, and thus the graph of φ is closed.

By Theorem 3.1 there exists a fixed point, that is, a vector y with $y \in \varphi(y)$. By the definition of φ, there exist x^i, z^i for every agent i such that (x^i, y^i, z^i) is optimal for c_p^i in P^i among the vectors that satisfy the personal service equations of i with respect to $(y^j)_{j\in[n]\setminus\{i\}}$. This also means that (x^i, y^i, z^i) satisfies all personal service equations of agent i, so (x, y, z) is fixed-ratio feasible. Because of the above optimality property, it is a fixed-ratio equilibrium. □

4 PPAD-completeness

In this section we show that if we fix the prices of contractual arcs, then the problem of finding an equilibrium in a safe instance of MMFP (see Section 1.3.2) is PPAD-complete. An interesting aspect of this problem is that it seems easier than many-player Nash equilibrium, because the set of equilibria is the union of some faces of a polyhedron by Proposition 3.3. Membership in PPAD follows from the fact that the computational version of Kakutani's fixed point theorem is in PPAD, see [12].

Completeness is proved by reducing two-player Nash equilibrium to our problem. To be more precise, we reduce approximate 2-Nash, so we use the following breakthrough result of Chen, Deng, and Teng [4].

Theorem 4.1 ([4])**.** *For any $\alpha > 0$, the problem of computing an $m^{-\alpha}$-approximate Nash equilibrium of a two-player game is PPAD-complete, where m is the number of strategies.*

We need a couple of remarks in order to use this theorem. First, it is well known that the problem of finding two-player Nash equilibria can be reduced to finding symmetric Nash equilibria in symmetric games, so we will assume that the game is symmetric, with utility matrix $A \in \mathbb{R}^{m\times m}$. Second, the above theorem is valid if the matrix A is normalized in the sense that its entries are bounded. For our purposes it is convenient to say that a symmetric two-player game is *normalized* if the elements of the matrix A are rationals in the interval $[1, 2]$. Third, approximate equilibria can be defined in several ways; we use a definition in [4]: x^* is an *ϵ-well supported approximate symmetric Nash equilibrium* if $x_k^* > 0$ implies that $\sum_{j=1}^m a_{kj}x_j^* > \max_{i\in[m]} \sum_{j=1}^m a_{ij}x_j^* - \epsilon$. Finally, it is convenient to set $\alpha = 1$. To sum up, we use the following form of the theorem.

Corollary 4.2 ([4]). *The problem of computing a $\frac{1}{m}$-well supported approximate symmetric Nash equilibrium of a symmetric normalized two-player game is PPAD-complete.*

The above problem will be called $\frac{1}{m}$-APPROXIMATE 2-NASH in this paper.

Proof of Theorem 1.4. We have to reduce $\frac{1}{m}$-APPROXIMATE 2-NASH to finding an equilibrium in a safe instance of MMFP. Given a symmetric normalized game defined by a matrix $A \in [1,2]^{m\times m}$, we construct a safe instance of MMFP featuring 4 agents. In order to make the construction more understandable, the agents are named Decision Maker, Combiner, Maximizer and Inverter. The high level view of the role of the agents is that the Decision Maker decides the values of x satisfying $x \geq 0$ and $\sum_{j=1}^m x_j = 1$, while the other agents are "gadgets" that compute $m(M - \sum_{j=1}^n a_{kj}x_j)$ for every k, where $M = \max_{i\in[m]} \sum_{j=1}^m a_{ij}x_j$. These values then appear in the network of the Decision Maker as contractual demands, in such a way that in an equilibrium solution x is guaranteed to be a $\frac{1}{m}$-well supported approximate Nash equilibrium. The prices will be 0 on all contractual arcs. The details of the construction, illustrated on Figure 1, are the following.

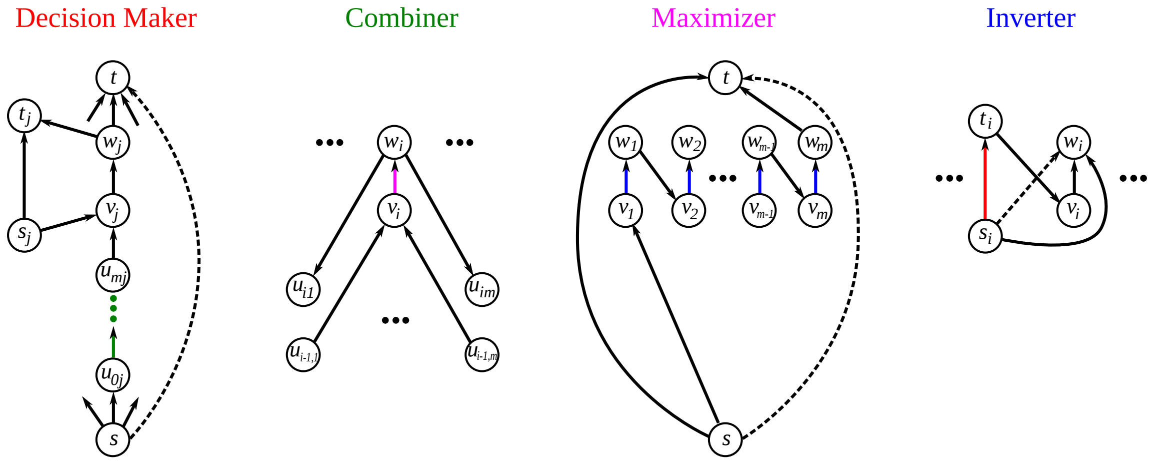

Figure 1: Illustration of the networks of various agents. Dashed arrows are normal demands, coloured arrows are contractual arcs.

The Decision Maker has one normal demand: a unit commodity from s to t. His network contains m internally disjoint st-paths: $s, u_{0j}, u_{1j}, \ldots, u_{mj}, v_j, w_j, t$ ($j = 1, \ldots, m$), each arc having cost 0 and capacity 1. The arcs $u_{i-1,j}u_{ij}$ ($i,j \in [m]$) are contractual arcs to the Combiner, and the multiplier of $u_{i-1,j}u_{ij}$ is a_{ij}.

The network of the Decision Maker also contains nodes s_j, t_j ($j = 1, \ldots, m$), arcs s_jt_j with cost 1 and capacity $2m$, and arcs s_jv_j, w_jt_j with cost 0 and capacity 1.

For a feasible solution x, we will denote by x_j the flow value on the arc su_{0j} in the network of the Decision Maker.

The network of the Combiner consists of paths $u_{i-1,j}, v_i, w_i, u_{ij}$ for $i,j \in [m]$, all arcs having cost 0 and capacity 2. The arcs v_iw_i are contractual arcs to the Maximizer with multiplier 1. Note that the Combiner has a unique way to route his contractual demands, and in a feasible solution the flow value on arc v_iw_i is $\sum_{j=1}^m a_{ij}x_j$, so this is the contractual demand appearing at the Maximizer.

The network of the Maximizer consists of a path $s, v_1, w_1, v_2, w_2, \ldots, v_m, w_m, t$, all arcs having cost 0 and capacity 2. The arcs v_iw_i are contractual arcs to the Inverter, with multiplier 1. There is also an arc st with cost 1 and capacity 2. The Maximizer has a normal demand of 2 from s to t. The routing that the Maximizer has to choose in an equilibrium solution is the following: he must route his contractual demands on the edges v_iw_i; he must route as much of his normal demand on the long $s-t$ path as possible, and route the rest on the arc st. Let M denote $\max_{i\in[m]}\sum_{j=1}^m a_{ij}x_j$ (note that $M \leq 2$). Then the portion of the normal demand routed on the long path is $2-M$ units, so the contractual demand appearing at the Inverter between v_i and w_i is $2-M+\sum_{j=1}^m a_{ij}x_j$ units.

Inverter has a path s_i, t_i, v_i, w_i for each $i \in [m]$ with arcs of cost 0 and capacity 2. The arcs s_it_i are contractual arcs to the Decision Maker, with multiplier m. In addition he has arcs s_iw_i ($i \in [m]$) with cost 1 and capacity 2, and normal demands of size 2 from s_i to w_i. In an equilibrium solution, he routes his contractual demands on the arcs v_iw_i, and routes as much of his normal $s_i - w_i$ demand on the path of cost zero as possible, the bottleneck being the arc v_iw_i. Thus $M - \sum_{j=1}^m a_{ij}x_j$ is routed on the path of cost zero, and $2-M+\sum_{j=1}^m a_{ij}x_j$ on the arc s_iw_i. This means that the contractual demand appearing at the Decision Maker between s_i and t_i is $m(M-\sum_{j=1}^m a_{ij}x_j)$.

It is easy to check that this construction results in a safe instance: we added arcs of non-zero cost with enough capacity to carry the maximum demands.

Let us prove that an equilibrium solution x^* corresponds to a $\frac{1}{m}$-well supported approximate Nash equilibrium. Clearly, $x^* \geq 0$ and $\sum_{j=1}^m x_j^* = 1$, so we have to prove that $x_i^* > 0$ implies that $\sum_{j=1}^m a_{ij}x_j^* > M - \frac{1}{m}$, where $M = \max_{i\in[m]}\sum_{j=1}^m a_{ij}x_j^*$. In the discussion above we have showed that the Decision Maker has a contractual demand of $m(M-\sum_{j=1}^m a_{ij}x_j^*)$ from s_i to t_i. We know that there is an index k such that $\sum_{j=1}^m a_{kj}x_j^* = M$, so the contractual demand from s_k to t_k is 0. This means that the following is a feasible multicommodity flow in the network of the Decision Maker:

- route the normal demand on the path $s, u_{0k}, u_{1k}, \ldots, u_{mk}, v_k, w_k, t$,
- for $i \neq k$, let $\delta_i = m(M - \sum_{j=1}^m a_{ij}x_j^*)$. Route $\min\{1, \delta_i\}$ units of the $s_i - t_i$ contractual demand on the path s_i, v_i, w_i, t_i, and the rest on the arc s_it_i.

It is easy to check that this is a minimum cost multicommodity flow for the given demands. Since x^* is an equilibrium solution, it must have the same cost as this one in the network of the Decision Maker. This is only possible if $\delta_i \leq 1 - x_i^*$ whenever $x_i^* > 0$. Thus $M - \sum_{j=1}^m a_{ij}x_j^* < \frac{1}{m}$ whenever $x_i^* > 0$, which means that x^* is a $\frac{1}{m}$-well supported approximate Nash equilibrium. □

5 Polynomially solvable cases

In the PPAD-completeness proof in Section 4, we reduced approximate 2-Nash to 4-agent MMCF instances where the provider-customer pairs constituted a directed

4-cycle on the set of agents. One may ask if this leaves room for an interesting class of service configurations where an equilibrium can be found in polynomial time.

A natural candidate is the class of problems where the customer-provider relationships form an acyclic digraph, and indeed it is not hard to show that a fixed-ratio equilibrium can be computed efficiently in that case. The main result of this section is an efficient algorithm for a broader class of problems. Let us consider an auxiliary directed graph $D^* = ([n], A^*)$ on the set of agents, in which there are $|S^i \cap T^j|$ parallel ij arcs if $i \neq j$ (so D^* does not have loops).

With this definition, the directed graph corresponding to the hard instances constructed in the proof of Theorem 1.4 is a 4-cycle *with many parallel arcs*. In contrast to this, Theorem 1.5 states that if the strongly connected components of D^* are *simple* directed cycles, then there is a polynomial time algorithm to find a fixed-ratio equilibrium.

The main tool of the proof is an algorithm for finding approximate fixed points of a special kind of mapping, as described in Theorem 1.6. Let $\varphi_i : [0,1] \to \mathcal{P}([0,1])$ ($i \in [m]$) be mappings such that $\varphi_i(t)$ is a non-empty interval for every $t \in [0,1]$, and the graph of φ_i is closed. We are interested in finding a fixed point of the mapping $(x_1, x_2, \ldots, x_m) \mapsto (\varphi_m(x_m), \varphi_1(x_1), \ldots, \varphi_{m-1}(x_{m-1}))$. In other words, we are looking for a vector $x = (x_1, \ldots, x_m)$ for which $x_{i+1} \in \varphi_i(x_i)$ for every $i \in [m]$. Here and later in this section, unless otherwise stated, we consider the indices modulo m, i.e. $x_{m+1} = x_1$ and $\varphi_m = \varphi_0$.

Definition. A vector $x = (x_1, \ldots, x_m)$ that satisfies $x_{i+1} \in \varphi_i(x_i)$ for every $i \in [m]$ is called a *cyclically fixed vector*.

In order to have a meaningful definition of running time, we use the following oracle model: there is an evaluation oracle which, given $t \in [0,1]$ and $i \in [m]$, returns some $z \in \varphi_i(t)$ in one step. We also consider basic arithmetic operations as one step. Of course it is hopeless to compute a cyclically fixed vector exactly in this oracle model, but, as stated in Theorem 1.6, we can approximate it in polynomial number of steps.

Proof of Theorem 1.6. The algorithm itself is quite simple and its time complexity is straightforward, the more involved part being the proof of its correctness. During the algorithm we always follow the rule that if at some point the oracle returns $z \in \varphi_i(t)$, then the triplet (i, t, z) is stored, and we use z in all subsequent evaluations of $\varphi_i(t)$.

The intervals are determined successively in reverse order. To determine I_m, initially let $a_m = 0$ and $b_m = 1$. Let $\psi_m = \varphi_{m-1} \circ \varphi_{m-2} \circ \cdots \circ \varphi_1 \circ \varphi_m$. The function ψ_m is an interval-valued function with a closed graph, since it is the composition of such functions.

Let $t = (a_m + b_m)/2$. We can compute a value $z \in \psi_m(t)$ by m successive oracle calls. If $z \leq t$, then let $b_m = t$. If $z \geq t$, then let $a_m = t$. These steps are repeated until $b_m - a_m \leq \epsilon$. Let $I_m = [a_m, b_m]$.

Suppose that we have already determined $I_{i+1} = [a_{i+1}, b_{i+1}]$. We modify the function φ_i as follows.

$$\varphi'_i(t) = \begin{cases} \varphi_i(t) \cap I_{i+1} & \text{if } \varphi_i(t) \cap I_{i+1} \neq \emptyset, \\ a_{i+1} & \text{if } z < a_{i+1} \text{ for every } z \in \varphi_i(t), \\ b_{i+1} & \text{if } z > b_{i+1} \text{ for every } z \in \varphi_i(t). \end{cases}$$

Note that this modification can be implemented simply by modifying the value returned by the oracle after each oracle call for φ_i: if the returned value is smaller than a_{i+1}, then we change it to a_{i+1}, and if it is greater than b_{i+1}, then we change it to b_{i+1}.

In order to compute I_i, initially let $a_i = 0$ and $b_i = 1$, and let $\psi_i = \varphi_{i-1} \circ \varphi_{i-2} \circ \cdots \circ \varphi_0 \circ \varphi'_{m-1} \circ \cdots \circ \varphi'_{i+1} \circ \varphi'_i$.

In a general step, let $t = (a_i + b_i)/2$, and let us compute a value $z \in \psi_i(t)$ by m oracle calls.

Definition. The sequence of the returned values of these m oracle calls is called the *itinerary* of the pair (i, t).

Let $b_i = t$ if $z \leq t$, and let $a_i = t$ if $z \geq t$. The above steps are repeated until $b_i - a_i \leq \epsilon$, in which case we fix I_i to be the interval $[a_i, b_i]$. We can observe the following.

Observation 5.1. *The m-th step of the itinerary of (i, a_i) is greater than a_i, and the m-th step of the itinerary of (i, b_i) is smaller than b_i.* □

The algorithm described above computes the interval I_i in $O(m \log(\frac{1}{\epsilon}))$ steps for a given i, so the total number of steps is $O(m^2 \log(\frac{1}{\epsilon}))$. It remains to show that the intervals contain a cyclically fixed vector.

Let us define mappings $\varphi^*_i : I_i \to \mathcal{P}(I_{i+1})$ for each $i \in [m]$ by $\varphi^*_i(t) = \varphi_i(t) \cap I_{i+1}$. Note that the image for a value t is either a closed interval or the empty set. We also define for any positive integer k (this time not taken modulo m) the mapping $\psi^*_k = \varphi^*_k \circ \varphi^*_{k-1} \circ \cdots \circ \varphi^*_1$. We can observe that $\psi^*_{m+k}(I_1) \subseteq \psi^*_k(I_1)$ for any k. We define ψ^*_0 to be the identity function.

Claim 5.2. *The set $\psi^*_k(I_1)$ is non-empty for every k.*

Proof. Indirectly, suppose that $im + k$ (where $1 \leq k \leq m$) is the smallest integer for which $\psi^*_{im+k}(I_1) = \emptyset$. We may assume w.l.o.g. that $\varphi_k \circ \psi^*_{im+k-1}(t) > b_{k+1}$ for every $t \in I_1$. This implies that $b_{k+1} \in \psi^*_{(i-1)m+k}(I_1)$ (provided that $(i-1)m + k \geq 0$), because $\varphi_k \circ \psi^*_{(i-1)m+k-1}(t)$ is an interval that contains b_{k+1}.

Let us examine the itinerary of $(k+1, b_{k+1})$. We claim that for any $j < m$ the j-th step of the itinerary is in $\psi^*_{(i-1)m+k+j}(I_1)$, provided that $(i-1)m+k+j \geq 0$. We show this by induction on j, first for $j < m - k$. In this case $\psi^*_{(i-1)m+k+j}(I_1)$ is non-empty and so the set $\varphi'_{k+j} \circ \psi^*_{(i-1)m+k+j-1}(I_1)$, which contains the j-th step of the itinerary, is equal to $\psi^*_{(i-1)m+k+j}(I_1)$.

Next, we show that for $m-k \leq j < m$ it also holds that the j-th step of the itinerary is in $\psi^*_{(i-1)m+k+j}(I_1)$. The difficulty here is that the itinerary proceeds according to the mapping φ_{k+j}, which, as opposed to φ'_{k+j}, may have values outside of I_{k+j+1}.

Suppose that the j-th step is the first to be outside of I_{k+j+1}; we can assume w.l.o.g. that it is greater than b_{k+j+1}. This means that b_{k+j+1} is in $\psi^*_{(i-1)m+k+j}(I_1)$ (because $\psi^*_{(i-1)m+k+j}(I_1)$ is non-empty), so the itinerary of $(k+j+1, b_{k+j+1})$ leads to b_{k+1}, after which it takes the same steps as the itinerary of $(k+1, b_{k+1})$ — here we use that the evaluation oracle cannot return different values for the same input. Thus the m-th step of the itinerary of $(k+j+1, b_{k+j+1})$ is greater than b_{k+j+1}, contradicting Observation 5.1.

We can conclude that the $(m-1)$-th step of the itinerary of $(k+1, b_{k+1})$ is in $\psi^*_{im+k-1}(I_1)$. Therefore the m-th step of the itinerary is greater than b_{k+1} by our assumption, again contradicting Observation 5.1. □

Since the set $\psi^*_{im}(I_1)$ is a non-empty closed interval for every i, and $\psi^*_{(i+1)m}(I_1) \subseteq \psi^*_{im}(I_1)$, we have that

$$R = \cap_{i=1}^{\infty} \psi^*_{im}(I_1) \text{ is a non-empty interval } [a, b].$$

Claim 5.3. *R contains a fixed point of ψ^*_m.*

Proof. It is easy to see that $\psi^*_m(R) = R$. For $k \in [m]$, let

$$R_k = \{t \in R : \psi^*_k(t) = \emptyset, \text{ but } \psi^*_j(t) \neq \emptyset \text{ for } j < k\},$$

and let $R^* = R \setminus \cup_{j=1}^m R_j$. Our aim is to find an interval $I^* = [a^*, b^*] \subseteq R^*$ such that $\psi^*_m(I^*) = R$. To do this, we show that for every $k \in [m]$, there is an interval $I^*_k = [a^*_k, b^*_k] \subseteq R \setminus \cup_{j=1}^k R_j$ such that $\psi^*_m(I^*_k) = R$. We can start with $I^*_0 = R$; suppose that we have already determined I^*_{k-1}. If $t \in R_k \cap I^*_{k-1}$, then either $\psi^*_k([a^*_{k-1}, t]) \subseteq \psi^*_k([t, b^*_{k-1}])$ or vice versa because both are sub-intervals of I_{k+1} containing the same endpoint of I_{k+1}. Consequently, either $\psi^*_m([a^*_{k-1}, t]) = R$ or $\psi^*_m([t, b^*_{k-1}]) = R$. We may assume w.l.o.g. that there is at least one t for which $\psi^*_m([t, b^*_{k-1}]) = R$. Let $a^*_k = \sup\{t \in I^*_{k-1} : \psi^*_m([t, b^*_{k-1}]) = R\}$.

Because of the closed graph property, $a^*_k \notin R_k$ and $\psi^*_m([a^*_k, b^*_{k-1}]) = R$ holds. If $R_k \cap [a^*_k, b^*_{k-1}] = \emptyset$, then we can set $b^*_k = b^*_{k-1}$. Otherwise by the choice of a^*_k we have that $\psi^*_m([t, b^*_{k-1}]) \neq R$ for every $t \in R_k \cap ([a^*_k, b^*_{k-1}]$. On the other hand, for such a t both $\psi^*_k([a^*_k, t])$ and $\psi^*_k([t, b^*_{k-1}])$ contain the same endpoint of I_{k+1}, so one of them contains the other. Since $\psi^*_m([a^*_k, t]) \cup \psi^*_m([t, b^*_{k-1}]) = R$, it follows that $\psi^*_m([a^*_k, t]) = R$. Let $b^*_k = \inf\{R_k \cap [a^*_k, b^*_{k-1}]\}$; then $\psi^*_m([a^*_k, b^*_k]) = R$, and $[a^*_k, b^*_k] \cap R_k = \emptyset$, as required.

We obtained an interval $I^* = [a^*, b^*] \subseteq R^*$ such that $\psi^*_m(I^*) = R$. Now it follows from the closed graph property of ψ^*_m on I^* that it has a fixed point in I^*. □

By definition, a fixed point of ψ^*_m implies the existence of a cyclically fixed vector x with $x_i \in I_i$ $(i \in [m])$. This concludes the proof of the theorem. □

We now show that if the graph of each mapping φ_i is a 2-dimensional polyhedral complex that can be described by inequalities of bit size M, then an exact cyclically fixed vector can be computed in a number of steps polynomial in M.

Corollary 5.4. *Let $\varphi_i : [0,1] \to \mathcal{P}([0,1])$ $(i \in [m])$ be given as above. Then in $O(m^2M^2)$ steps we compute a cyclically fixed vector.*

Proof. Let $\epsilon = \exp(-M^2)$. Since the graph of the function φ_i is a 2-dimensional polyhedral complex of bit size M, the choice of ϵ guarantees that all vertices of this polyhedral complex in the interior of $I_i \times [0,1]$ (if the exist at all) have the same first coordinate t, which can be computed in polynomial time. We modify the algorithm in the proof of Theorem 1.6 the following way: after achieving $b_i - a_i \le \epsilon$, we make an additional step with the above t in place of $t = (a_i + b_i)/2$. This way we obtain intervals I_i such that the graph of φ_i has no vertices in the interior of $I_i \times [0,1]$ ($i \in [m]$). This means that the mappings φ_i are "linear" on these intervals in the sense that their graph is of the form $\alpha_i x_i \le x_{i+1} \le \beta_i x_i$, and an equilibrium can be found by solving an LP. $\square$

We are now ready to prove that a fixed-ratio equilibrium can be found in polynomial time in a fixed-ratio safe instance if the strongly connected components of D^* are simple directed cycles.

Proof of Theorem 1.5. Let us consider a fixed-ratio safe instance of the problem where every strong component of D^* is a simple directed cycle. Let $C_1, \dots, C_q$ be the family of strong components in reverse topological order, and for $1 \le k \le q$ let V_k denote the set of agents in C_k. We will compute solutions $(x^i, y^i, z^i) \in P^i$ of agents $i \in V_k$ for $k = 1, 2, \dots, q$ (in this order), in such a way that the personal service equations of agents in V_k are satisfied. Because of the reverse topological order, these equations are not modified in later stages of the algorithm, so the solution obtained at the end is fixed-ratio feasible.

Suppose that we have already determined the solutions of agents up to V_{k-1}. Let $i_1, \dots, i_m$ denote the agents in V_k, in reverse order of the cycle C_k. Let s_j be the unique element of $S^{i_{j+1}} \cap T^{i_j}$ and let

$$[\ell_j, u_j] = \{t \in \mathbb{R} : \ P^{i_j} \cap \{(x^{i_j}, y^{i_j}, z^{i_j}) : y^{i_j}_{s_j} = t\} \neq \emptyset\}.$$

The function φ_j assigns to each $t \in [\ell_j, u_j]$ a nonempty subinterval of $[\ell_{j+1}, u_{j+1}]$ the following way. If we fix $y^{i_j}_{s_j}$ at t, then by the definition of fixed-ratio safeness there exists a vector $(\overline{x}^{i_{j+1}}, \overline{y}^{i_{j+1}}, \overline{z}^{i_{j+1}}) \in P^{i_{j+1}}$ such that, together with the values already fixed, it satisfies all personal service equations of agent i. Let us take the set of such vectors that have minimum c_p-cost, and let Let $\varphi_j(t)$ consist of the possible $\overline{y}^{i_{j+1}}_{s_{j+1}}$ values in these vectors. This is a non-empty subinterval of $[\ell_{j+1}, u_{j+1}]$. A value of $\varphi_j(t)$ can be computed using linear programming, and the graph of φ_j is a 2-dimensional polyhedral complex describable by inequalities of bit size $O(M)$, where M is the size of the input.

We can use Corollary 5.4 (by appropriately scaling the functions φ_j) to find a cyclically fixed vector of $\varphi_1, \dots, \varphi_m$. This means that we have found vectors $(x^i, y^i, z^i) \in P^i$ ($i \in V_k$) such that the personal service equations of all agents in V_k are satisfied, and for each $i \in V_k$, if we fix the solutions of other agents, then (x^i, y^i, z^i) has minimum c_p-cost among the vectors in P^i that satisfy the personal service equations of i.

The end result is a fixed-ratio feasible solution because all personal service equations are satisfied, and it is a fixed-ratio equilibrium by the above argument. $\square$

6 Open questions

A notable open question is the computational complexity of finding an equilibrium when the number of services is constant. The answer is unknown even in the following simple setting: there are two agents and only three services, two of them offered by agent 1 to agent 2, and one vice versa. Although one might expect this case to be easily solvable, an algorithm similar to the one described in the proof of Theorem 1.6 does not seem to work.

From the mechanism design point of view, it would be interesting to find bargaining mechanisms that lead to fair prices, or at least to some approximate version of fairness.

Acknowledgements

We would like to thank all authors of [3] for fruitful discussions on the topic. The authors received a grant (no. CK 80124) from the National Development Agency of Hungary, based on a source from the Research and Technology Innovation Fund.

References

[1] R. Agarwal and Ö. Ergun, *Mechanism design for a multicommodity flow game in service network alliances*, Operations Research Letters 36 (2008), 520–524.

[2] E. Anshelevich, B. Shepherd, G. Wilfong, *Strategic Network Formation through Peering and Service Agreements*, Proceedings of 47th Annual IEEE Symposium on Foundations of Computer Science (2006), 77–86.

[3] A. Bernáth, T. Király, E.R. Kovács, G. Mádi-Nagy, Gy. Pap, J. Pap, J. Szabó, L. Végh, *Algorithms for multiplayer multicommodity flow problems*, Central European Journal of Operations Research, available online (2012), DOI: 10.1007/s10100-012-0255-6.

[4] X. Chen, X. Deng, S-H. Teng, *Settling the complexity of computing two-player Nash equilibria*, Journal of the ACM, Volume 56 Issue 3, May 2009, Article No. 14

[5] J.J.M. Derks and S.H. Tijs, *Stable outcomes for multicommodity flow games*, Methods of Operations Research, 50 (1985), 493–504.

[6] L. Guyi and Ö. Ergun, *Dual Payoffs, Core and a Collaboration Mechanism Based on Capacity Exchange Prices in Multicommodity Flow Games*, Proceedings of the 4th International Workshop on Internet and Network Economics (WINE '08), 2008, 61–69.

[7] R. Johari, S. Mannor, J.N. Tsitsiklis, *A contract-based model for directed network formation*, Games and Economic Behavior 56 (2006), 201–224.

[8] S. Kakutani, *A generalization of Brouwer's fixed point theorem*, Duke Math. J. 8 (1941), 457–459.

[9] E. Kalai, E. Zemel, *Generalized network problems yielding totally balanced games*, Operations Research 30 (5) (1981) 998–1008.

[10] E. Kalai, E. Zemel, *Totally balanced games and games of flow*, Mathematics of Operations Research 7 (3) (1982) 476–478.

[11] E. Markakis, A. Saberi, *On the core of the multicommodity flow game*, Proceedings of the 4th ACM conference on Electronic commerce (2003), 93–97.

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

[13] C.H. Papadimitriou, *Algorithms, Games, and the Internet*, Proceedings of the thirty-third annual ACM symposium on Theory of computing (2001), 749–753.