

# On the Core of the Multicommodity Flow Game

Evangelos Markakis<sup>a,\*</sup>

<sup>a</sup>*Georgia Tech, College of Computing, 801 Atlantic Drive, Atlanta, GA, 30332*

Amin Saberi<sup>b</sup>

<sup>b</sup>*Georgia Tech, College of Computing, 801 Atlantic Drive, Atlanta, GA, 30332*

---

## Abstract

In [26], Papadimitriou proposed a game theoretic framework for analyzing incentive issues in Internet routing. In particular he defined the following coalitional game: Given a network with a multicommodity flow satisfying node capacity and demand constraints, the payoff of a node is the total flow originated or terminated at it. A payoff allocation is in the *core* of the game if and only if there is no subset of nodes that can increase their payoff by seceding from the network. We answer one of the open problems in [26] by proving that for any network, the core is non-empty in both the transferable (where the nodes can compensate each other with side payments) and the non-transferable case. In the transferable case we show that such an allocation can be computed in polynomial time. We also generalize this result to the case where a strictly concave utility function is associated with each commodity.

*Key words:* Internet, Cooperative Games, Core, Duality

---

## 1 Introduction

The Internet, which has become a common playground for a large number of entities with selfish motives and varying degrees of collaboration, naturally gives rise to new game theoretic issues [26]. Problems that stem from Internet

---

\* Corresponding author

*Email addresses:* [vangelis@cc.gatech.edu](mailto:vangelis@cc.gatech.edu) (Evangelos Markakis),  
[saberi@cc.gatech.edu](mailto:saberi@cc.gatech.edu) (Amin Saberi).

applications are very different from traditional algorithmic problems as the behavior of the participants is determined by their own goals and not by the instructions of the designer. It seems that such problems would require techniques and ideas from both computer science and game theory. In this paper we focus on one specific problem, namely incentive issues in Internet routing.

The Internet is composed of many administrative domains or *Autonomous Systems* (ASes). Each AS is usually administered by a single entity. For example, a corporation or a university campus often defines an autonomous system.

The connectivity of the Internet is determined by agreements between ASes for routing each other's traffic. The current protocol for routing between ASes is the Border Gateway Protocol (BGP). BGP works without a centralized authority by allowing ASes to constantly announce and exchange routing paths.

ASes can be considered independent self-interested agents, following routing policies that serve their own interests. In particular, ASes would like to satisfy their own and their customers' traffic demands and at the same time they would prefer to avoid carrying transit traffic, i.e., traffic that is neither originated nor destined to them or their customers. Avoiding transit traffic though, might result in sub-optimal efficiency and instability in the network. It might even affect the network connectivity.

A question that arises naturally is whether it is possible to have a routing scheme which maximizes network efficiency and is stable in the sense that no AS or subset of ASes has an incentive to secede. In [26], Papadimitriou proposed a game theoretic formulation of this question by defining the following coalitional game: Given a network with a multicommodity flow, satisfying node capacity and demand constraints, the payoff of a node is the total flow originated or terminated at it (flow passing through a node is not included in its payoff). One of the open problems in [26] was to find sufficient conditions under which the *core* of the game is non-empty. An outcome of a game is in the core if no subset of players can collude and obtain a better payoff for its members, either viewed as a set (transferable payoff), or for each player in the coalition individually (non-transferable payoff).

We show that the core of this game is always non-empty. In the transferable case, an allocation in the core can be computed in polynomial time by solving the dual program of the multicommodity flow problem. For the same game with non-transferable payoff our proof of non-emptiness of the core is non-constructive. It is still an open question whether a core allocation can be computed efficiently for this case.

We also generalize this result to the case where a strictly concave utility function is associated with each commodity. In [17], Kelly proposed such a model

for analyzing charging, rate control and routing in communication networks. An optimal outcome in his model is expressed as the solution of a non-linear program. The dual variables of that program (shadow prices) can be interpreted as actual payments of the nodes for their traffic. Using a similar argument as before, we show that if ASes compensate each other according to these shadow prices, the resulting payoff allocation is in the core.

The use of dual variables for producing an allocation in the core is not new. In [4,6,13–16,25,29] classes of games are defined in which a core allocation is obtained as a function of the dual variables. In fact if the demand constraints are dropped and all the nodes have unit capacity then the non-emptiness of the core in the multicommodity flow game with transferable payoff follows from Theorem 1 in [6]. For facility-location games [4,13] show that the dual of the facility location problem is equivalent to the problem of finding core allocations if there is no integrality gap. In some games, e.g. [29] every allocation in the core is obtained via a dual solution. However this is not the case in our game. Several complexity results have also been obtained (e.g. for testing membership or non-emptiness of the core) among others by [6,9,7,5].

Incentive issues in routing have also been addressed in [23] and [10] from a mechanism design point of view. In their models each link [23] or node [10] incurs a cost for routing a packet. VCG-type payment mechanisms are obtained to make the links or the nodes behave truthfully regarding the cost of routing.

In the next section we give some definitions and results from coalitional game theory which will be used later on. In Section 3 we focus on the linear multicommodity flow game and prove that the core is always non-empty in both the transferable and the non-transferable case. In Section 4 we give a game-theoretic formulation of Kelly’s non-linear model [17] and prove that again the core is non-empty. We conclude in Section 5 with open problems and directions for further research.

## 2 Definitions and Notation

A coalitional (or cooperative) game is determined by a set of players  $N = \{1, \dots, n\}$ , a set of possible outcomes  $O(S)$  for every coalition  $S \subseteq N$  and a set of payoff vectors  $V(S)$  corresponding to the outcomes. A payoff vector  $x \in V(S)$  corresponding to the outcome  $o \in O(S)$  determines the payoff of each player if the outcome  $o$  is realized. The set  $N$  is sometimes referred to as the *grand* coalition.

A solution concept in coalitional game theory is usually defined as a set of

payoff allocations that are *stable* in some certain sense. Among all the solution concepts that have been proposed over the years, the *core* is probably the most intuitive one. The core consists of all payoff allocations for which no subset of players (coalition) can improve upon by cooperating only among themselves. This means that once an agreement in the core has been reached, no coalition has an incentive to secede.

We will define the core for two scenarios of coalitional games. In games with *transferable* payoff, players can compensate each other with side payments. In such games a coalition  $S$  can be completely characterized by the maximum total payoff that it can achieve in  $O(S)$ . We will denote this number by  $v(S)$ . The coalition is allowed to split the payoff  $v(S)$  among its members in any possible way. The *core* of the game will be the set of payoff allocations for which no coalition can gain more.

More formally, with each such game we associate a *characteristic function*  $v : P(N) \rightarrow R^+$ , where  $P(N)$  is the powerset of  $N$ . Following standard assumptions in the literature we require that:

- (i)  $v(\emptyset) = 0$ .
- (ii)  $v(S \cup T) \geq v(S) + v(T)$ , if  $S \cap T = \emptyset$ .

We will denote a payoff allocation by a vector  $x = (x_1, \dots, x_n)$ ,  $x_i \geq 0$ , where  $x_i$  is the payoff allocated to player  $i$ . Given an allocation  $x$ , we will denote by  $x(S)$  the payoff that is allocated to a coalition  $S$ , i.e.  $x(S) = \sum_{i \in S} x_i$ . A payoff allocation is an *imputation* if  $x(N) = v(N)$ . The core is the set of stable imputations:

$$core = \{x : x(N) = v(N) \text{ and } x(S) \geq v(S) \quad \forall S \subset N\}.$$

In games with non-transferable payoff, compensations among different players are not possible. In this case a coalition  $S$  is characterized by the set  $V(S)$  of payoff vectors. The interpretation of  $V(S)$  is that it contains all the possible payoff allocations that can be obtained by  $S$ . A coalition  $S$  can improve upon a payoff vector  $x$  if there exists an allocation  $y \in V(S)$  such that  $x_i < y_i$  for all  $i \in S$ . Hence the core will be:

$$core = \{x \in V(N) : \forall S \nexists y \in V(S) \text{ s.t. } y_i > x_i \quad \forall i \in S\}.$$

Necessary and sufficient conditions for the non-emptiness of the core in games with transferable payoff were given by Bondareva and Shapley [1,28]. In [27], Scarf generalized their result and provided a sufficient condition in games with non-transferable payoff.

**Definition 1** *Let  $T$  be a collection of coalitions.  $T$  is said to be a balanced*

collection if and only if we can find nonnegative weights  $\delta_S$  for all  $S \in T$  such that for every  $i \in N$ ,  $\sum_{S \in T: i \in S} \delta_S = 1$ .

Given a coalition  $S$ , we will call a vector  $u$  *attainable* by  $S$  if  $u \in V(S)$ . We will also denote by  $u_S$  the vector whose entries are the entries of  $u$  that correspond to the players of  $S$  (i.e., the projection of  $u$  to  $S$ ).

**Definition 2** *A game is balanced if and only if for every balanced collection  $T$ , if  $u$  is such that  $u_S$  is attainable by  $S$ , for all  $S \in T$ , then  $u$  is attainable by  $N$ .*

**Theorem [Scarf]** : Every balanced game has a non-empty core.

### 3 The Multicommodity Flow Game

As an attempt to address incentive issues in the Internet, Papadimitriou [26] defined the following coalitional game: let  $G$  be an undirected graph on a set of nodes  $N$  with a capacity  $c_i$  on each node and a symmetric demand matrix  $D$  (where  $d_{ij}$  is the demand between nodes  $i$  and  $j$ ). Each node represents an AS and the capacity of node  $i$  is a simplification attempting to capture the capability of the corresponding subnetwork. An outcome of the game is a feasible multicommodity flow subject to demand and capacity constraints, i.e., a vector  $\{f_p\}$  where for a path  $p$  from  $i$  to  $j$ ,  $f_p$  is the flow exchanged between these nodes along path  $p$ . The total flow exchanged between  $i$  and  $j$  will then be equal to  $f_{ij} = f_{ji} = \sum f_p$ , where the sum is taken over all paths connecting  $i$  and  $j$ . Therefore the matrix  $F = (f_{ij})$  will satisfy  $F \leq D$ . In the game with transferable payoff the value  $v(S)$  for a coalition  $S \subseteq N$  is the maximum flow subject to demand and capacity constraints in the graph induced by  $S$ . In other words  $v(S) = \max \sum_i \sum_j f_{ij}/2$ , where the maximum is taken over feasible flows. In the non-transferable case the set  $V(S)$  consists of the vectors  $u = (u_1, \dots, u_{|S|})$  such that there exists a feasible flow  $F$  in the graph induced by  $S$  for which  $u_i = \sum_j f_{ij}/2$ . Note that for a vector  $u \in V(S)$ , it is not necessarily true that the sum  $\sum u_i$  is equal to the maximum flow in the graph induced by  $S$ .

Finding sufficient conditions for the non-emptiness of the core was posed as an open problem in [26]. In the following subsections we will show that the core is always non-empty in both cases.

### 3.1 The Coalitional Game with Transferable Payoff

For each  $i, j \in N$  let  $P_{ij}$  denote the set of all paths between  $i$  and  $j$  and let  $P = \cup P_{ij}$ . A maximum flow satisfying as much of the demands as possible is the solution of the following linear program:

$$\begin{aligned}
& \text{maximize} && \sum_{p \in P} f_p \\
& \text{subject to} && \sum_{p: i \in p} f_p \leq c_i \quad \forall i \in N \\
& && \sum_{p \in P_{ij}} f_p \leq d_{ij} \quad \forall i, j \in N \\
& && f_p \geq 0 \quad \forall p \in P
\end{aligned} \tag{1}$$

The dual program is:

$$\begin{aligned}
& \text{minimize} && \sum_{i \in N} c_i x_i + \sum_{i, j \in N} d_{ij} y_{ij} \\
& \text{subject to} && y_{ij} + \sum_{i \in p} x_i \geq 1 \quad \forall p \in P_{ij} \\
& && x_i \geq 0 \quad \forall i \in N \\
& && y_{ij} \geq 0 \quad \forall i, j \in N
\end{aligned} \tag{2}$$

Here the dual variable  $x_i$  corresponds to node  $i$  in the graph and the variable  $y_{ij}$  corresponds to the unordered pair of nodes  $(i, j)$ .

The first part of the following theorem can also be proved by directly applying the Bondareva-Shapley theorem. However we present the proof that constructs the payoff allocation by using the dual program to establish polynomial running time.

**Theorem 3** *The core of the multicommodity flow game with transferable payoff is non-empty. Furthermore, a payoff allocation in the core can be computed in polynomial time.*

**Proof:** Consider an optimal dual solution  $\{x_i\}, \{y_{ij}\}$ . For each node  $i$  define its payoff to be:

$$p_i = c_i x_i + \frac{\sum_j d_{ij} y_{ij}}{2}$$

To show that the payoff vector  $\{p_i\}$  belongs to the core we need to show that:

- (i)  $\sum_{i \in N} p_i = OPT(N)$ .
- (ii) For every subset  $S$ ,  $\sum_{i \in S} p_i \geq OPT(S)$ .

where for  $S \subseteq N$ ,  $OPT(S)$  is the optimal value of (1) when restricted to the subgraph induced by  $S$ .

For the first part note that:

$$\sum_{i \in N} p_i = \sum_{i \in N} c_i x_i + \sum_{i, j \in N} d_{ij} y_{ij} = OPT(N)$$

by the strong duality theorem [2].

For the second part, consider a coalition  $S$  and the network that is induced by  $S$ . Let  $i \in S, j \in S$  and  $p \in P_{ij}$  such that  $p$  is entirely in the induced graph. Since  $\{x_i : i \in N\}, \{y_{ij} : i, j \in N\}$  is a dual optimal (and hence feasible) solution to the original problem it holds that :

$$y_{ij} + \sum_{i \in p} x_i \geq 1$$

Therefore,  $(\{x_i : i \in S\}, \{y_{ij} : i, j \in S\})$  is a dual feasible solution for the induced linear program on  $S$ . Thus:

$$\sum_{i \in S} c_i x_i + \sum_{i, j \in S} d_{ij} y_{ij} \geq OPT(S)$$

But now the following holds:

$$\begin{aligned} \sum_{i \in S} p_i &= \sum_{i \in S} c_i x_i + \sum_{i \in S} \left( \frac{\sum_{j \in N} d_{ij} y_{ij}}{2} \right) \\ &\geq \sum_{i \in S} c_i x_i + \frac{1}{2} \sum_{i \in S} \sum_{j \in S} d_{ij} y_{ij} \\ &= \sum_{i \in S} c_i x_i + \sum_{i, j \in S} d_{ij} y_{ij} \\ &\geq OPT(S) \end{aligned}$$

Hence  $\{p_i\}$  is in the core.

The above argument directly yields a polynomial time algorithm for computing an allocation that lies in the core by solving the dual program. It should be

noted here that even though the dual program in general has an exponential number of constraints, it is known that it can be solved in polynomial time [31].

□

In the payoff allocation that we constructed, each of the nodes  $i, j$  receives exactly half of the payoff term  $y_{ij}d_{ij}$ . It is easily seen that if we arbitrarily allocate  $\alpha_{ij}y_{ij}d_{ij}$  to node  $i$  and  $(1 - \alpha_{ij})y_{ij}d_{ij}$  to node  $j$  for  $0 \leq \alpha_{ij} \leq 1$  the resulting allocation is also in the core. We should note however that these are not the only core allocations of the game.

### 3.2 The Coalitional Game with Non-Transferable Payoff

In a coalitional game with transferable payoff, we assume that players can compensate each other with a side payment. This assumption is not justified in many cases [24].

We will show that the core of the multicommodity flow game without transferable payoff is not empty using Scarf's Theorem (Section 2). Thus, we need only show that the game is balanced.

**Theorem 4** *The multicommodity flow game with non-transferable payoff is balanced and hence has a non-empty core.*

**Proof :** Consider a balanced collection of coalitions  $T$ . Let  $\delta_S$  be the corresponding weight to each coalition such that for every  $i \in N$ ,  $\sum_{S \in T, i \in S} \delta_S = 1$ . Consider a payoff vector  $u$  which is attainable by every coalition  $S \in T$ . We need to show that  $u$  is attainable by  $N$ . For a coalition  $S \in T$ , since  $u$  is attainable by  $S$ , there exists a feasible flow  $f^S$  subject to demand and capacity constraints such that for every player  $i$ :  $\sum_j f_{ij}^S/2 = u_i$  ( $f_{ij}^S$  is the flow routed for the commodity  $(i, j)$  in the subgraph induced by  $S$ ). We construct the flow  $f = \sum_{S \in T} \delta_S f^S$ . For a node  $i$  the total flow that we route for the commodities containing  $i$  (divided by 2) is:

$$\frac{1}{2} \sum_{S \in T, i \in S} \delta_S \sum_j f_{ij}^S = \sum_{S \in T, i \in S} \delta_S u_i = u_i$$

It is also easy to see that this flow satisfies capacity and demand constraints. Hence  $u$  is attainable by  $N$  and the game is balanced, which implies that the core is non-empty. □



## 4 The Game with Non-Linear Utility Functions

In [17], Kelly defines a mathematical model for analyzing issues of pricing, rate control and routing in communication networks. Similar models have also been used among others by [18,20,21]. The model consists of a network with a set of nodes  $N$ , a capacity for each node  $c_i$  and a set of commodities  $K$ . We will denote by  $P_s$  the set of paths that commodity  $s$  is using to send flow from its source to its sink and  $P = \cup P_s$ . If a commodity  $s$  is sending flow at a rate of  $x_s$  then its source and sink derive a utility of  $U_s(x_s)$  where  $U_s$  is an increasing, strictly concave and continuously differentiable function (according to Shenker [30] traffic that leads to such utility functions is called *elastic* traffic). We further assume that the aggregate utility of the network for flow rates  $\{x_s\}$  is  $\sum_s U_s(x_s)$ .

In this setting, if flow  $f_p$  is sent along each path  $p \in P$  then the total flow rate for commodity  $s$  is  $\sum_{p \in P_s} f_p$ . To find the system's optimal rates we need to solve the following non-linear optimization problem:

$$\begin{aligned}
 & \text{maximize} && \sum_{s \in K} U_s(x_s) \\
 & \text{subject to} && \sum_{i \in p} f_p \leq c_i \quad \forall i \in N \\
 & && \sum_{p \in P_s} f_p = x_s \quad \forall s \in K \\
 & && f_p \geq 0 \quad \forall p \in P \\
 & && x_s \geq 0 \quad \forall s \in K
 \end{aligned} \tag{3}$$

Note that unlike Section 3 we do not have any demand constraints. This is purely for ease of exposition and our results hold even when a demand matrix is specified.

We construct the dual of (3). Consider the Lagrangian form:

$$\begin{aligned}
 L(x, f, \lambda, \mu) &= \sum_{s \in K} U_s(x_s) + \sum_i \mu_i (c_i - \sum_{i \in p} f_p) - \sum_{s \in K} \lambda_s (x_s - \sum_{p \in P_s} f_p) \\
 &= \sum_{s \in K} (U_s(x_s) - \lambda_s x_s) + \sum_{p \in P} f_p (\lambda_{s(p)} - \sum_{i \in p} \mu_i) + \sum_{i \in N} c_i \mu_i
 \end{aligned}$$

where  $\lambda = \{\lambda_s : s \in K\}$ ,  $\mu = \{\mu_i : i \in N, \mu_i \geq 0\}$  are vectors of Lagrange multipliers and for a path  $p$ ,  $s(p)$  denotes the commodity that the path serves. Define the function

$$D(\lambda, \mu) = \max_{x \geq 0, f \geq 0} L(x, f, \lambda, \mu)$$

We can simplify the function  $D(\lambda, \mu)$  by noting that:

$$\frac{\partial L}{\partial f_p} = \lambda_{s(p)} - \sum_{i \in p} \mu_i$$

This means that at a maximum of  $L$  over the orthant  $x \geq 0, f \geq 0$  the following should be true:

$$\text{If } f_p > 0 \text{ then } \lambda_{s(p)} = \sum_{i \in p} \mu_i.$$

Thus

$$\begin{aligned} D(\lambda, \mu) &= \max_{x \geq 0} \sum_{s \in K} (U_s(x_s) - \lambda_s x_s) + \sum_{i \in N} c_i \mu_i \\ &= \sum_{s \in K} \max_{x_s} (U_s(x_s) - \lambda_s x_s) + \sum_{i \in N} c_i \mu_i \end{aligned}$$

The dual program of (3) is:

$$\begin{aligned} &\text{minimize} && D(\lambda, \mu) \\ &\text{subject to} && \mu_i \geq 0 \quad \forall i \in N \end{aligned} \tag{4}$$

The objective function of (3) is differentiable and strictly concave and the feasible region is compact. Hence (3) has an optimal solution. By the duality theorem [22], there exists a dual optimal solution for (4).

As in [17,20] the dual variables of an optimal solution (shadow prices) can be interpreted as congestion control signals. Furthermore they can also indicate actual payments to the nodes for routing traffic. In this case we show that payments defined by an optimal dual solution result in a payoff allocation which lies in the core.

As in Section 3 we can view the nodes of the network as players in a coalitional game with transferable payoff. The outcome of the game is again a multicommodity flow satisfying the constraints in (3) and for a coalition  $S \subset N$  we define its payoff  $v(S)$  to be:  $v(S) = 2OPT(S)$  where  $OPT(S)$  is the optimal value of (3) when restricted to the subgraph induced by the nodes in  $S$ . This is a natural generalization of the game that we studied in Section 3 where now the total payoff of a coalition is not the maximum flow it can send but a concave function of the maximum flow.

The question that arises of course is whether this game has a non-empty core. We will answer this question in the affirmative.

**Theorem 5** Any optimal solution  $(\lambda, \mu)$  to the dual program (4) gives rise to a payoff allocation which is in the core.

**Proof :** The argument is essentially the same as in the proof of Theorem 3. For a node  $i \in N$ , let  $K(i)$  be the set of commodities in which  $i$  is either a source or a sink. We can define the following payoff allocation to the nodes:

$$p_i = \sum_{s \in K(i)} \max_{x_s} (U_s(x_s) - \lambda_s x_s) + 2c_i \mu_i$$

To show that  $p = \{p_i\}$  is in the core note first that:

$$\begin{aligned} \sum_{i \in N} p_i &= \sum_{i \in N} \left( \sum_{s \in K(i)} \max_{x_s} (U_s(x_s) - \lambda_s x_s) + 2c_i \mu_i \right) \\ &= 2 \sum_{s \in K} \max_{x_s} (U_s(x_s) - \lambda_s x_s) + 2 \sum_{i \in N} c_i \mu_i = v(N) \end{aligned}$$

Therefore it remains to show that for every coalition  $S$ ,  $\sum_{i \in S} p_i \geq 2OPT(S)$ . Consider a coalition  $S$  and the dual variables that correspond to commodities and nodes in the subgraph induced by  $S$ . These variables form a feasible solution to the dual of (3) when restricted to this subnetwork. Therefore we have:

$$\sum_{i \in S} p_i = 2D(\lambda, \mu; S) \geq 2OPT(S) = v(S)$$

where by  $D(\lambda, \mu; S)$  we denote the dual objective function restricted to the subnetwork of  $S$ . Hence the allocation  $\{p_i\}$  lies in the core.

□

We should also note that our proof for the non-transferable case in Section 3.2 also holds when the utilities are concave functions of the flow, which is the case here.

## 5 Discussion and Open Questions

In [19], Kelly and Vazirani showed that the problem of charging and rate control as defined in Kelly [17] can be seen as a generalization of Fisher's market equilibrium problem [3,8]. The optimum dual variables in that model correspond to market clearing prices. Hence, Theorem 5 on core allocations in section 4 is along the same lines of the classic result in coalitional game theory that allocations corresponding to an equilibrium in the market lies in the core [24].

The core of a game is a useful concept in a cooperative setting where all the information regarding preferences, capacities and demands is known to all agents. Clearly this is not the case in the Internet. Moreover, in the core allocation that we constructed in sections 3.1 and 4, the payoff that a node receives depends on its capacity. It can be seen by using complementary slackness conditions that a node might receive a bigger payoff if it announces a smaller capacity. It is an interesting problem to design a distributed strategy-proof mechanism such that no node has an incentive to lie about its capacity. For related results on algorithmic mechanism design see [23,10,11].

The proof of non-emptiness of the core in Section 3.2 is based on Scarf's Theorem which is non-constructive. An open problem is to find an algorithm for computing a solution in the core efficiently.

In the non-linear model, if each commodity uses only one path, it is shown in [12,20] that shadow prices can be computed by a distributed algorithm where the local computation is done on each link (on each AS in our case). We are not aware of any result for the general case.

## 6 Acknowledgments

We would like to thank Milena Mihail for introducing the problem to us and for motivating discussions. We would also like to thank Dick Lipton, Christos Papadimitriou and Vijay Vazirani for useful comments and discussions and Tejas Iyer, Amir Ronen and Neal Young for pointers to the literature. Finally we would like to thank the anonymous referees for their comments and suggestions.

## References

- [1] O. Bondareva. The core of an  $n$  person game. *Vestnik Leningrad University*, 13:141–142, 1962.
- [2] V. Chvatal. *Linear Programming*. New York: W. H. Freeman, 1983.
- [3] W. C. Brainard and H. E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*, 1270, 2000.
- [4] P. Chardaire. *Facility Location Optimization and Cooperative Games*. PhD Thesis, University of East-Anglia, UK, 1998.
- [5] V. Conitzer and T. Sandholm. Complexity of determining non-emptiness of the core. Technical Report CS-02-137, CMU, 2002.

- [6] X. Deng, T. Ibaraki, and H. Nagamochi. Algorithms and complexity in combinatorial optimization games. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 720–729, 1997.
- [7] X. Deng and C. Papadimitriou. On the complexity of cooperative game solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994.
- [8] N. Devanur, A. Saberi, C. Papadimitriou, and V. V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *Annual IEEE Symposium on the Foundations of Computer Science*, pages 389–395, 2002.
- [9] U. Faigle, S. Fekete, W. Hochstättler, and W. Kern. On the complexity of testing membership in the core of min-cost spanning tree games. *International Journal of Game Theory*, 26:361–366, 1997.
- [10] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In *Symposium on Principles of Distributed Computing*, pages 173–182, 2002.
- [11] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 1–13. ACM Press, 2002.
- [12] N. Garg and N. Young. On-line end-to-end congestion control. In *Annual IEEE Symposium on the Foundations of Computer Science*, pages 303–312, 2002.
- [13] M. Goemans and M. Skutella. Cooperative facility location games. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 76–85, 2000.
- [14] D. Granot. A generalized linear production model: A unified model. *Mathematical Programming*, 34:212–222, 1986.
- [15] K. Jain and V. V. Vazirani. Application of approximation algorithms to cooperative games. In *Annual Symposium on Theory of Computing*, pages 364–372, 2001.
- [16] E. Kalai and E. Zemel. Generalized network problems yielding totally balanced games. *Operations Research*, 30:998–1008, 1982.
- [17] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- [18] F. P. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [19] F. P. Kelly and V. V. Vazirani. Rate control as a market equilibrium. In preparation, 2003.
- [20] S. Low and D. Lapsley. Optimization, flow control 1: Basic algorithm and convergence. *IEEE Transactions on Networking*, 7(6):861–874, 1999.

- [21] S. Low, F. Paganini, and J. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, 22(1):28–43, Feb. 2002.
- [22] A. Nemirovskii, and A. Ben-Tal. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. MPS-SIAM Series on Optimization, 2001.
- [23] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Annual ACM Symposium on the Theory of Computing*, pages 129–140, 1999.
- [24] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts, 1994.
- [25] G. Owen. On the core of linear production games. *Mathematical Programming*, 9:358–370, 1975.
- [26] C. H. Papadimitriou. Algorithms, games, and the internet. In *Annual ACM Symposium on the Theory of Computing*, pages 749–753, 2001.
- [27] H. Scarf. The core of an n person game. *Econometrica*, 35(1):50–69, 1967.
- [28] L. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14:453–460, 1967.
- [29] L. Shapley and M. Shubik. The assignment game 1: The core. *International Journal of Game Theory*, 1:111–130, 1972.
- [30] S. Shenker. Fundamental design issues for the future internet. *IEEE Journal Selected Areas Communication*, 13:1176–1188, 1995.
- [31] V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.