Economic Mechanisms for Shortest Path Cooperative Games with Incomplete Information

T.S. Chandrashekar and Yadati Narahari*

Department of Computer Science and Automation, Indian Institute of Science, Bangalore - 560 012, India
Tel: +91-80-22932773; Fax: +91-80-23602911
{chandra, hari}@csa.iisc.ernet.in

Abstract. In this paper we present a cooperative game theoretic interpretation of the shortest path problem. We consider a *buying* agent who has a budget to go from a specified source node s to a specified target node t in a directed acyclic network. The budget may reflect the level of utility that he associates in going from node s to node t. The edges in the network are owned by individual utility maximizing agents each of whom incurs some cost in allowing its use. We investigate the design of economic mechanisms to obtain a least cost path from s to t and to share the surplus (difference between the budget and the cost of the shortest path) generated among the participating agents in a *fair* manner. Previous work related to this problem assumes that cost and budget information is common knowledge. This assumption can be severely restrictive in many common applications. We relax this assumption and allow both budget and cost information to be private, hence known only to the respective agents. We first develop the structure of the shortest path cooperative game with incomplete information. We then show the non-emptiness of the incentive compatible core of this game and the existence of a surplus sharing mechanism that is incentive efficient and individually rational in virtual utilities, and strongly budget balanced.

1 Introduction

The shortest path problem and its many variants such as all pairs shortest paths and stochastic shortest paths occur in a wide variety of contexts and have been studied extensively. More recently, motivated by applications in grid computing, mobile ad-hoc networks, and electronic commerce, game theoretic interpretations of the shortest path problem including both the non cooperative interpretation [1, 2, 3, 4] and the cooperative interpretation [5, 6] have emerged.

In these game theoretic interpretations, economic agents control and provide access to different resources - edges and/or nodes, of the network for a price. In addition, one other agent, hereafter called the buying agent, associates a certain level of utility, in traversing the network between two specified nodes - the source s and the target t. In the remaining part of this section we first summarize the existing state-of-the-art in analyzing shortest path games and then motivate the context for this paper.

* Corresponding author.

X. Deng and Y. Ye (Eds.): WINE 2005, LNCS 3828, pp. 70–79, 2005.

1.1 Extant Work and Motivation

In the noncooperative game theoretic interpretation introduced in [2] it is assumed that the cost of edges in the network are only privately known. In order to find the least cost path, a Vickrey-Clarke-Groves (VCG) mechanism (see chapter 23 of [7]) is employed to elicit truthful cost information. When this approach is analyzed from the perspective of the buying agent, it has been shown that the incentives provided through this mechanism to elicit truthful cost information may be arbitrarily high [8] and cannot be avoided [3].

In the cooperative game theoretic interpretation introduced in [5, 6], the buying agent indicates his budget limitation (which may be a proxy for the utility he obtains) in going from node s to node t. The N agents owning the edges/nodes in the network then agree to cooperate among themselves to identify a shortest (least cost) path and share the surplus revenue, if any, among themselves in some *fair* way as enunciated by solutions concepts such as the core or the Shapley value. It should be noted here that this is a complete information cooperative game and no information is privately held. However, if the agents (both the buying agent and resource owners) have access to private budget and cost information, as is often the case in real world scenarios, then there is reason to expect strategic misrepresentation of either the budget and/or cost information by them. We give below two real world scenarios where such misrepresentation can be commonly expected.

- *QoS routing in communication networks*: Consider a scenario where high quality video is to be transmitted over the public Internet with a Quality of Service (QoS) guarantee. The agent requiring such a service is willing to compensate the owners of the intermediate interconnecting networks. Clearly (1) the agent is going to act strategically so that he obtains the service at the least possible cost and (2) the resource owners will bid strategically so as to maximize their share of the surplus in the transaction.
- *Supply chain procurement*: Consider a procurement scenario with an automotive assembler where a subassembly is to be procured. The way in which a subassembly maybe put together may involve many alternatives in terms of which suppliers participate in supplying components and services. All this may be captured as a network with edges representing value added by each preceding node and the network itself converging into the node representing the assembler. All nodes with no incoming edge may be grouped into the category of source nodes and the objective now is to find a shortest path from any one of the source nodes to the terminal assemblers node. Clearly the automotive assembler as well as the suppliers can be expected to behave strategically and misrepresent budget and cost information.

To model these scenarios, we need to extend the analysis of the class of shortest path cooperative games in two directions: First, by including the buying agent as a game theoretic agent and second, by treating costs and budgets as private information. These extensions prompt us to formulate the scenarios as cooperative games with incomplete information.

1.2 Contributions and Outline

Our contributions in this paper are two-fold.

1. As far as we know this is the first time that a cooperative game with *incomplete information* arising out of the shortest path problem has been addressed. We develop the structure for this game.
2. For this class of games, we investigate the design of surplus sharing mechanisms. We invoke results in cooperative game theory to show the non-emptiness of the incentive compatible core for this class of games and then prove the existence of mechanisms that are incentive efficient and individually rational in virtual utilities, and strongly budget balanced.

The structure of the paper is as follows. In Section 2, we provide the basic notation to develop the structure of the shortest path cooperative game with incomplete information (SPCG-II). In Section 3, we show the non-emptiness of the Incentive Compatible Core for this game. In Section 4, we adapt the bargaining solution based on a generalization of the Shapley Value for the SPCG-II. Finally in Section 5, we summarize the contributions of the paper.

2 A Shortest Path Cooperative Game with Incomplete Information (SPCG-II)

To begin with, in this section we set out the basic notation required and then present the structure of the SPCG-II. We consider a directed graph $N = (V, E)$ with V, the set of vertices, E the set of edges and two special nodes s (source) and t(target). We let $n = |E|$ be the number of edges. Each edge in the network is assumed to be a commodity that is owned by an agent. For expositional clarity we assume that the number of agents is equal to the number of edges in the network. However, any analysis that follows can be extended to cases where agents own multiple edges. We therefore let $I = \{1, 2, \ldots, n, n+1\}$ be the set of all agents where $1, 2, \ldots, n$ are edge owning agents and agent $(n+1)$ is the buying agent.

Each agent $i \in I$ has an initial endowment vector $e_i \in \Re_+^{n+1}$ where $e_{i,j} \in \{0,1\}, \forall j \in \{1, 2, \ldots, n\}$ and $e_{i,(n+1)} \in \Re$. This implies that when agent i owns the edge $j \in E$ then $e_{ij} = 1$ and is otherwise 0. Having assumed that there is a one-to-one correspondence between the edges and the agents, we have $e_{i,i} = 1$ and $e_{i,j} = 0, \forall j \neq i$. In addition the $(n+1)^{th}$ entry in the endowment vector e_i indicates the amount of money that agent i has. For any agent $i \in I$ we let T_i denote the set of possible types. The type $t_i \in T_i$ for all edge owning agents $i \in I \backslash \{(n+1)\}$ is a description of the cost that he incurs when his edge is used. The type $t_{(n+1)} \in T_{(n+1)}$ however describes the budget of the buying agent.

Let $\mathcal{C}$ denote the set of all possible coalitions or non empty subsets of I, that is, $\mathcal{C} = \{S | S \subseteq I, S \neq \emptyset\}$. For any coalition $S \in \mathcal{C}$, we let $T_S = \times_{i \in S} T_i$ so that any $t_S \in T_S$ denotes a combination of types of all agents i in S. For the grand coalition I, we let $T = T_I = \times_{i \in I} T_i$. Now, for any subset $S \in \mathcal{C}$, which includes the agent $(n+1)$, we define a set of market transactions. This follows from a shortest path computation that is carried out after the agents i in S declare their types t_i. We call this set of market transactions as the set of possible outcomes $X_S(t_S)$, such that $X_S(t_S) = \{(\tilde{e}_i)_{i \in S} | \tilde{e}_i \in \Re_+^{n+1}$ and $\sum_{i \in S} \tilde{e}_{ij} \leq \sum_{i \in S} e_{ij}, \quad \forall j \in \{1, 2, \ldots, n, n+1\}$, where $\tilde{e}_i$ is the outcome vector of agent i after the transaction is carried out. The outcome set specifies that the reallocation of resources and money is such that there is no infusion of additional resources into the system. We also define the set X_S and X as the sets that include the outcomes for all possible type declarations $t_S \in T_S$ and all possible coalitions $S \in \mathcal{C}$. So, $X_S = \bigcup_{t_S \in T_S} X_S(t_S)$ and $X = \bigcup_{S \in \mathcal{C}} X_S$.

The reallocation of resources, i.e., the edges and the money, is carried out as follows: Given the set of edges owned by the agents in S and the edge costs declared by them, a shortest path computation identifies the set of edges whose ownership is to be transferred to the buying agent. Following this, each edge agent whose edge is transferred to the buying agent is compensated according to the declared cost. The entire surplus, defined as the difference between the budget and the cost of the shortest path, that results from the transaction is then given to either the buying agent or to one of the agents who plays an active role in providing the shortest path. Note here that from the way in which we define the outcomes, there are only a finite number of outcomes, which is one greater than the number of edge agents who provide the shortest path. In reality however, we would expect the surplus to be shared among the participating agents. This sharing of the surplus is achieved by using randomized mechanisms as state contingent contracts.

Now, for any outcome $x \in X$ and any $t \in T$, we let the utility for an agent $i \in I$ be $u_i(x, t)$. For any agent i and outcome x, the final outcome vector $\tilde{e}_i$ reflects the edges that it currently owns and the money that it has after the transfers have been carried out. That is $\tilde{e}_{i,i}$ can be either 0 or 1 and $\tilde{e}_{i,(n+1)} \in \Re$. So, the payoff that the agent receives from outcome x when its type is t_i is given by $u_i(x, t) = \tilde{e}_{i,(n+1)} + (\tilde{e}_{i,i} - 1)t_i$.

We use the notation $I - i$ to denote $I \backslash \{i\}$ and we write $t = (t_{-i}, t_i)$. Similarly, (t_{-i}, s_i) denotes the vector t where the i^{th} component is changed to s_i. Now, for any $t \in T$, we let $p_i(t_{-i}|t_i)$ denote the conditional probability that t_{-i} is the combination of types for players other than i as would be assessed by player i if t_i were his type. We will assume that these probabilities are consistent as in [9]. We are now in a position to define the structure of the SPCG-II. In line with the structure for cooperative games with incomplete information in [10], the shortest path cooperative game with incomplete information can be described by the structure below.

$$\Gamma = (X, x^*, (T_i)_{i \in I}, (u_i)_{i \in I}, (p_i)_{i \in I}) \tag{1}$$

Here, X refers to the set of all outcomes for all coalitions $S \in \mathcal{C}$ that could be formed; x^* is a default outcome that results when the agents are unable to come to an agreement over the solution. In the context of the SPCG-II, the default outcome is a null transaction whose utility for all types of all agents is 0. T_i,

u_i, and p_i are as defined earlier. This structure Γ of the game is assumed to be known to all agents. In addition we assume that each agent knows his own type before the start of negotiations. Our concern now is to develop a solution to this cooperative game and interpret the results in the context of the SPCG-II and the applications introduced in Section 1.

As opposed to solution concepts for cooperative games with complete information, where the focus is on finding an allocation of the surplus value to the participating agents, in cooperative games with incomplete information the focus is on finding mechanisms or state contingent contracts that a grand coalition of all agents agree to [10, 11]. Here, in the context of the SPCG-II, we use the same conceptual apparatus as presented in [10, 12, 13, 14, 15] and find that the solution approaches for this class of games are analogous to the *Core* and the *Shapley Value*. It appears, as we shall see below, that extensions to these concepts based on important additional insights lead to the *Incentive Compatible Core* [15] and *Myerson's generalization of the Shapley Value* [10] which we adapt as solutions to SPCG-II.

3 The Incentive Compatible Core for SPCG-II

The core as a solution concept for cooperative games with complete information is based on the premise that a group of agents can cooperate and agree upon a coordinated set of actions which can then be enforced; and the resulting feasible allocations of surplus value cannot be improved upon by any other coalition. In the context of incomplete information games, however, since we are concerned with state contingent contracts rather than allocations, the meaning of the two terms - "feasible" and "improve upon" needs a precise clarification. Feasibility in this context refers to contracts that satisfy not only physical resource constraints in each information state but also the incentive constraints that arise when information is private and is inherently unverifiable. And secondly "improving upon" a state contingent contract implies that agents need to examine what information they use at the time of evaluating contracts. The evaluations may be carried out either at the ex-ante stage when none of the agents has any type information or at the interim stage when individuals know their own type information but not that of other agents. For the class of games that we are concerned with here, agents already possess private information when they enter into negotiations. So, the evaluation of contracts should be done at the interim stage and the measure of an agent's well being is based on conditionally expected utilities (conditional on private information). Our further analysis is based on this measure of evaluation. Before that we introduce some additional notation to enable the analysis that follows.

Let $\triangle X$ and $\triangle X_S$ be the sets of all possible probability distributions on X and X_S respectively. Now, we define μ and μ_S as mappings from T and T_S to $\triangle X$ and $\triangle X_S$ respectively. i.e., $\mu : T \rightarrow \triangle X; \quad \mu_S : T_S \rightarrow \triangle X_S$. Now, μ and μ_S may be viewed as direct random mechanisms. Note also that a state contingent contract can be written as function from T to $\triangle X$. So, while strictly a mechanism should be seen as a means to implement a state contingent allocation, here we interpret it as a state contingent contract. Having defined a mechanism, we can now define the conditionally expected utilities of the agents. We let $U_i(\mu, s_i|t_i)$ be the conditionally expected utility of agent i from the mechanism μ, if i's true type is t_i but he reports s_i while all other agents report their types truthfully. So we have,

$$U_i(\mu, s_i|t_i) = \sum_{t_{-i}\in T_{-i}} p_i(t_{-i}|t_i) \sum_{x\in X} \mu(x|t_{-i}, s_i)u_i(x,t) \tag{2}$$

$$U_i(\mu, t_i) = U_i(\mu, t_i|t_i) = \sum_{t_{-i}\in T_{-i}} p_i(t_{-i}|t_i) \sum_{x\in X} \mu(x|t)u_i(x,t) \tag{3}$$

Now, from our definition of outcomes and mechanisms, we can easily verify that the mechanisms always meet the resource feasibility constraints. Since a feasible mechanism is required to satisfy both physical resource constraints and incentive constraints, we now define Π as the set of all mechanisms which are physically feasible and Π^* as the set of all mechanisms which are also incentive compatible. So we have,

$$\Pi = \{\mu : T \to \triangle X \quad | \quad \mu(x|t) \geq 0 \quad \text{and} \quad \sum_{x\in X} \mu(x|t) = 1\} \tag{4}$$

$$\Pi^* = \{\mu \in \Pi \quad | \quad U_i(\mu, t_i) \geq U_i(\mu, s_i|t_i) \quad \forall s_i, t_i \in T_i, \forall i \in I\} \tag{5}$$

These two equations taken together imply that the set of Bayesian incentive compatible mechanisms Π^*, is a subset of Π, the set of resource feasible mechanisms that obey probability constraints. In a similar vein, we define Π_S and Π_S^* as the set of resource feasible mechanisms and incentive compatible mechanisms respectively for a coalition $S \subseteq \mathcal{C}$.

$$\Pi_S = \{\mu_S : T_S \to \triangle X_S \quad | \quad \mu(x_S|t_S) \geq 0 \quad \text{and} \quad \sum_{x_S\in X_S} \mu(x_S|t_S) = 1\} \tag{6}$$

$$\Pi_S^* = \{\mu_S \in \Pi_S \quad | \quad U_i(\mu_S, t_i) \geq U_i(\mu_S, s_i|t_i) \quad \forall s_i, t_i \in T_i, \forall i \in S\} \tag{7}$$

Now, in the spirit of the core defined for cooperative games with complete information, we can say that a mechanism or a state contingent contract is in the core of the SPCG-II if no subset of agents stands to gain by breaking away and negotiating a separate contract which gives them all a better expected utility. But in order to break away and negotiate a more beneficial contract, the agents in the breakaway coalition, say S, must be able to gain in an event that is discernible by all of them. We define an event Λ as $\Lambda = \times_{i\in I}\Lambda_i$, where $\Lambda_i \subseteq T_i$. This event Λ is discernible by a coalition S (or is common knowledge within S) if $p_i(\hat{t}_{-i}|t_i) = 0, \forall i \in S, t \in \Lambda, \text{and } \hat{t} \notin \Lambda$.

Now, coalition S has an objection to a mechanism $\mu \in \Pi^*$ if there exists a contract $\mu_S \in \Pi_S^*$ and an event Λ that is discernible by S such that the following inequality holds for all agents $i \in S$ with strict inequality holding for at least one of them.

$$U_i(\mu_S|t_i) \geq U_i(\mu|t_i), \quad \forall t_i \in \Lambda_i, \forall i \in S. \tag{8}$$

The incentive compatible core consists of all mechanisms $\mu \in \Pi^*$ to which there is no such objection. In other words if a mechanism μ has to belong to the incentive compatible core then there should not exist a coalition S, an incentive compatible mechanism μ_S and an event $\Lambda \in T$ such that:

1. Λ is discernible by S,
2. $U_i(\mu_S|t_i) \geq U_i(\mu|t_i), \quad \forall t_i \in \Lambda_i, \forall i \in S,$
3. $\sum_{i \in S} \tilde{e}_{ij} \leq \sum_{i \in S} e_{ij}, \quad \forall j \in \{1, 2, \ldots, n, n+1\}; \forall t \in T,$

The question that now arises is whether the core of such a game is non-empty. The answer to this lies in recognizing the fact that the utility functions $u_i(x, t)$ are all affine linear in t_i and from Remark 3.1 in [15], it can be deduced that the incentive compatible core is indeed non-empty. This immediately gives us the following theorem.

Theorem 1. *The shortest path cooperative game with incomplete information, $\Gamma = (X, x^*, (T_i)_{i \in I}, (u_i)_{i \in I}, (p_i)_{i \in I})$, has a non-empty incentive compatible core.*

The incentive compatible core is an important solution concept, whose non-emptiness provides strong guarantees for the stability of a coalition. Many times we are also interested in finding a single solution to the SPCG-II as opposed to a set of solutions like in the core. We address this issue next.

4 A Generalization of the Shapley Value for SPCG-II

In the case of cooperative games with incomplete information, we are concerned with finding an incentive efficient mechanism or a state contingent contract agreeable to the grand coalition I so that it is both implementable, and pareto dominates all other incentive compatible mechanisms. In addition, the mechanism chosen should fairly capture the power structure of the agents in the game and must also be an adequate compromise between the types of the agents if type information is to be protected in the bargaining process. A solution for this class of games was proposed in [10] which we adapt here to the SPCG-II.

A mechanism μ is incentive efficient iff $\mu \in \Pi^*$ and there is no other mechanism $\nu \in \Pi^*$ such that $U_i(\nu|t_i) \geq U_i(\mu|t_i), \forall i \in I, \forall t_i \in T_i$ with strict inequality holding at least for one type t_i of some agent i. Since, X and T are finite sets, the set of incentive compatible mechanisms Π^* is a closed convex polyhedron defined by incentive compatibility constraints and probability constraints. These constraints are linear and hence by the supporting hyperplane theorem if there is a set of utility transfer weights $\lambda \in \times_{i \in I} \Re^{T_i}$ then the incentive efficient mechanism μ is a solution to an appropriate linear programming problem (LP1). This is given by:

$$\max_{\mu \in \Pi^*} \sum_{i \in I} \sum_{t_i \in T_i} \lambda_i(t_i) U_i(\mu|t_i) \tag{9}$$

subject to:

$$U_i(\mu|t_i) \geq U_i(\mu, s_i|t_i), \quad \forall i \in I, \forall t_i \in T_i, \forall s_i \in T_i. \tag{10}$$

$$\mu(x|t) \geq 0, \forall x \in X; \quad \text{and} \quad \sum_{x \in X} \mu(x|t) = 1, \forall t \in T \tag{11}$$

The dual of LP1 can be constructed by letting $\alpha_i(s_i|t_i)$ be the dual variable corresponding to the incentive compatible constraint that requires agent i to not gain by claiming that his type is s_i when it is actually t_i. So, $\alpha \in \times_{i \in I} \Re^{T_i \times T_i}$ is a vector of dual variables. Now, a Lagrangian function can be written by multiplying each of the incentive compatibility constraints with its corresponding dual variable and adding it to the objective function of LP1. We then define $v_i(x, t, \lambda, \alpha)$ as the virtual utility of agent i with respect to λ and α for an outcome x when the type profile is t [10]. This is given by:

$$\begin{aligned} v_i(d, t, \lambda, \alpha) = \{\{\lambda_i(t_i) + \sum_{s_i \in T_i} \alpha_i(s_i|t_i)\} p_i(t_{-i}|t_i) u_i(x, t) \\ - \sum_{s_i \in T_i} \alpha_i(t_i|s_i) p_i(t_{-i}|s_i) u_i(x, (t_{-i}, s_i))\}/p(t) \end{aligned} \tag{12}$$

With this definition of virtual utilities, the Lagrangian of LP1 may be written in terms of the virtual utilities as:

$$\max_{\mu \in \Pi^*} \sum_{t \in T} p(t) \sum_{x \in X} \mu(x|t) \sum_{i \in I} v_i(x, t, \lambda, \alpha) \tag{13}$$

Notice that we now seek a mechanism that maximizes the sum of virtual utilities. It is also shown in [10], that when agents face binding incentive constraints, they appear to act according to the preferences of their virtual utilities and not their actual utilities. So, for cooperative games with incomplete information, the bargaining solution is based on conditional transfers of virtual utility rather than transfers of actual utility.

In the computation of the Shapley Value for cooperative games with complete information, the worth of each of the smaller coalitions serves only as a countervailing force to influence the final allocations of the surplus. Analogously, in the model of bargaining for incomplete information games, every coalition selects a threat mechanism against a complementary coalition. We note that the SPCG-II is a game with orthogonal coalitions, in the sense that the threat mechanisms only affect the payoffs of the agents in the coalition. In addition, only a coalition S which includes the buying agent $(n+1)$ can select a threat which has some positive utility for the coalition. All other coalitions can only select threats whose utility for the coalition is zero (empty threats!). We let $\Omega = \times_{S \in \mathcal{C}} \Pi_S$. That is any vector $\omega = (\mu_S)_{S \in \mathcal{C}} \in \Omega$ includes a specification of the threats μ_S that each coalition $S \in \mathcal{C}$ threatens to use in case of a breakdown in negotiations of the grand coalition. Since the significance of these threat mechanisms is only to influence the mechanism $\mu = \mu_I$ chosen by the grand coalition, we do not require them to be incentive compatible nor equitable. So, in our choice of threat mechanisms involving all coalitions S where S includes $(n+1)$ and $S \subset I$, we can restrict ourselves to those mechanisms which place the complete probability weight on the outcome which gives the maximum possible payoff to agent $(n+1)$.

We do this because one of the motivations in our application scenarios was to reduce the high payments that are seen when VCG mechanisms are used.

Now, we can define the warranted claims $W_S(\omega, t, \lambda, \alpha)$ of virtual utility of each coalition S with respect to λ and α given the type profile t and threat profile ω by considering only the parameters relevant to the coalition S and neglecting those of $I \backslash \{S\}$.

$$W_S(\omega, t, \lambda, \alpha) = \sum_{x_S \in X_S} \mu_S(x_S|t_S) \sum_{i \in S} v_i(x_S, t_s, \lambda, \alpha) \tag{14}$$

With these warranted claims of the coalitions we can build a characteristic function form game $W(\omega, t, \lambda, \alpha)$. That is, $W(\omega, t, \lambda, \alpha) = (W_S(\omega, t, \lambda, \alpha))_{S \in \mathcal{C}}$. The Shapley value of such a game is given by:

$$\phi_i(W(\omega, t, \lambda, \alpha)) = \sum_{S \in \mathcal{C}, S \supseteq \{i, (n+1)\}} \frac{(|S|-1)!((n+1)-|S|)!}{(n+1)!}(W_S(\omega, t, \lambda, \alpha)) \tag{15}$$

From the Shapley value, the expected virtual-utility payoff of agent i with type t_i is given by $\sum_{t_{-i} \in T_{-i}} P_i(t_{-i}|t_i) \Phi_i(W(\omega, t, \lambda, \alpha))$. We note here that the mechanism $\mu = \mu_I$ in the warranted claim $W_I(\omega, t, \lambda, \alpha)$ of the grand coalition I is the one that maximizes the sum of virtual utilities of all the agents in I and hence is the one which maximizes the Lagrangian of LP1 expressed in virtual utilities. That such a mechanism exists and is also individually rational is shown in [10]. And from our construction of the outcome set $X = \bigcup_{S \in \mathcal{C}} X_S$, where we ensure that there are no transfers of money into or out of the system, we can infer that the mechanism is strongly budget balanced. This discussion can be summarized as the following theorem:

Theorem 2. *The shortest path cooperative game with incomplete information, $\Gamma = (X, x^*, (T_i)_{i \in I}, (u_i)_{i \in I}, (p_i)_{i \in I})$, has a Shapley value mechanism that is incentive efficient and individually rational in virtual utilities, and strongly budget balanced.*

5 Summary

In this paper we have extended the analysis of shortest path cooperative games to scenarios with incomplete information where a buying agent is also a participant in the game. Such scenarios are routinely encountered in many real life applications such as supply chain procurement, Internet routing, etc. We have developed the structure for the shortest path cooperative game with incomplete information. We have then, using previous results in cooperative game theory, shown the following:

– the non-emptiness of the incentive compatible core of such a game.
– the existence of a Shapley Value mechanism that is incentive efficient and individually rational in virtual utilities and also strongly budget balanced.

We believe that this analysis can be extended to multi commodity network flow scenarios that capture more complex features of the motivating problems.

References

1. Noam Nisan. Algorithms for selfish agents. In *Proceedings of the 16th Symposium on Theoretical Aspects of Computer Science, Trier, Germany*, 1999.
2. Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166–196, 2001.
3. A. Archer and E. Tardos. Frugal path mechanisms. Technical report, Unpublished manuscript. Previously appeared as an extended abstract in Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms, 2002.
4. John Hershberger and Subhash Suri. Vickrey prices and shortest paths: What is an edge worth? In *IEEE Symposium on Foundations of Computer Science*, pages 252–259, 2001.
5. Fragnelli V, Garcia-Jurado I, and Mendez-Naya L. On shortest path games. *Mathematical Methods of Operations Research*, 52, 2000.
6. Mark Voorneveld and Sofia Grahn. Cost allocation in shortest path games. Technical report, Working Paper, Department of Mathematics, Utrecht University, 2000.
7. Andreu Mas-Collel, Michael D Whinston, and Jerry. R Green. *Micoreconomic Theory*. Oxford University Press, 1995.
8. Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, and Scott Shenker. A BGP-based mechanism for lowest-cost routing. In *In 21st ACM Symposium on Principles of Distributed Computing (PODC '02)*, pages 173–182, 2002.
9. John C Harsanyi. Games with incomplete information played by Bayesian players I-III. *Management Science*, 14(7):486–502, March 1968.
10. Roger B Myerson. Cooperative games with incomplete information. *International Journal of Game Theory*, 13:69–96, 1984.
11. B Holmstorm and Roger B Myerson. Efficient and durable decision rules with incomplete information. *Econometrica*, 51:1799–1819, 1983.
12. John C Harsanyi and Reinhard Selten. A generalized Nash solution for two person bargaining games with incomplete information. *Management Science*, 18(5):80–106, January, Part 2 1972.
13. Robert Wilson. Information, efficiency and the core of an economy. *Econometrica*, 46(4):807–816, July 1978.
14. Roger B Myerson. Two-person bargaining problems with incomplete information. *Econometrica*, 52:461–487, 1984.
15. Rajiv Vohra. Incomplete information, incentive compatibility and the core. *Journal of Economic Theory*, 86:123–147, 1999.