Computational Aspects of Game Theory	**Bertinoro Spring School 2011**

Lecture 16: The Price of Truth: VCG Payments and the Core

Lecturer: Bruno Codenotti

In this lecture we will analyze the shortest path game, and in particular the VCG payment scheme which induces link owners to reveal their true costs. We describe a framework under which the question of the *price of truth* can be studied: What is the price that we must pay in order to induce the players to behave truthfully? We then compare these payments with the core of the game.

16.1 Introduction

We consider a cooperative game (which we call the *shortest path game*) where the players are the owners of the edges in a network, and the goal is to send a message between nodes s and t in the network, minimizing the total cost of the transmission. The players might team up to form $s--t$ paths. The problem is that the players' costs are kept private. We will apply the VCG mechanism to select the team to perform the task. We will show that the VCG mechanism may significantly overpay the link owners, and then introduce the notion of *frugality* to characterize instances where the VCG mechanism does not lead to significant overpayments. We will talk about frugality in connection with two game theoretic concepts, i.e., the core and the property of *agents being substitutes*. We will show that the VCG mechanism is frugal if and only if the players are *substitutes*, and then we will see that agents are substitutes if and only if the core has a lattice structure, in which case the VCG payment is the supremum element of the core.

16.2 VCG Mechanisms for the Shortest Path Game

Recall from lecture 10 that in a VCG mechanism each player has a valuation function, which represents her preferences over the possible outcomes. The mechanism selects the outcome that maximizes the sum of the valuations of all the players, and computes a payment to the players. The VCG mechanism induces truthtelling by paying each player an amount such that the net profit of the player is equal to the value that she contributes to the system.

In the application to the shortest path game, the VCG mechanism chooses the shortest path according to the costs reported by the players (which we call *bids*).

Note that if there is a cut edge separating s and t, then such an edge has to be chosen no matter what, so that there is no way to elicit truthtelling by the owner of that edge. Thus we assume that there exist two edge-disjoint paths from s to t.

The payment to an edge e is chosen so that the profit of its owner is equal to the amount by which the presence of e decreases the length of the shortest path. The mechanism pays zero to the owners of edges not in the chosen path, and $c - d$ to the owners of edges in the chosen path, where c is the cost of the shortest path if e had been deleted, and d is the cost of the shortest path if e's owner had bid zero.

16.3 Vickrey Auction vs VCG Mechanism

Consider the graph in Figure 16.1, where the different paths are just disjoint edges connecting s to t. In this case, we can view the game as a (reverse) auction, where the auctioneer wants to *buy* the cheapest edge, and, according to the Vickrey payment scheme, is going to pay an amount equal to the second cheapest edge. So the game coincides with the second price (reverse) auction. Let b_e be the cheapest (reported) cost, and let b' be the (reported) cost of the *second* shortest path. Edge e is chosen and its owner is given a payment of $b' - b_e$, corresponding to its marginal contribution to the game. This instance of the shortest path game coincides with the second price auction, and has the desirable property that, in the presence of competition, we expect the quantity $b' - b_e$ to be small.

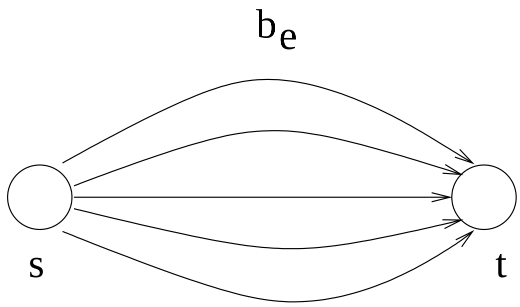

Figure 16.1: A shortest path game for which the VCG mechanism collapses to the Vickrey auction.

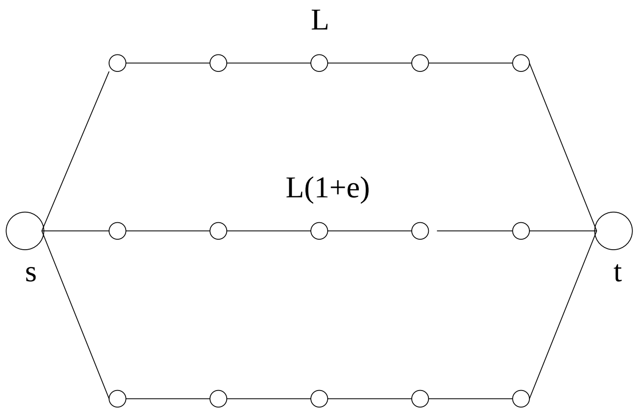

Figure 16.2: Overpayment even in the presence of a competitive path.

Such a nice feature of the Vickrey auction is unfortunately not shared by the more general VCG mechanism which can significantly overpay link owners even in the presence of competitive paths. This happens even in very simple special cases. Consider for example the graph in Figure 16.2. Here we lose the equivalence with the Vickrey auction, and the problem can be described as a game where one special player, the auctioneer, has a certain amount of money K which she can use to hire a team of agents (the link owners) which can provide a path from s to t. Let the cost of the unique shortest path be L and the cost of the next shortest path be $L(1+\epsilon)$. The marginal contribution of each owner of an edge along the shortest path is ϵL, so that ϵL has to be her payoff according to the VCG mechanism. Therefore the total payoff to the team of players in the shortest path is $L(1+\epsilon k)$, where k is the number of links in the shortest path. If ϵ is a fixed constant, then $L(1+\epsilon k) = \theta(kL)$, i.e., the payment exceeds by a factor of k not only the actual cost of the path, but also the cost of the second cheapest path.

This happens despite the presence of a competitive path, just longer by an $(1+\epsilon)$ multiplicative factor. The reason is that the chosen path is owned by many different entities, and each of these must be induced to bid truthfully.

16.4 Core vs VCG Payments: Some Examples

In this section we present some examples which show the interplay between the core of the shortest path game and the VCG payments. In the next section, we will introduce a framework under which this question can be analyzed in full generality.

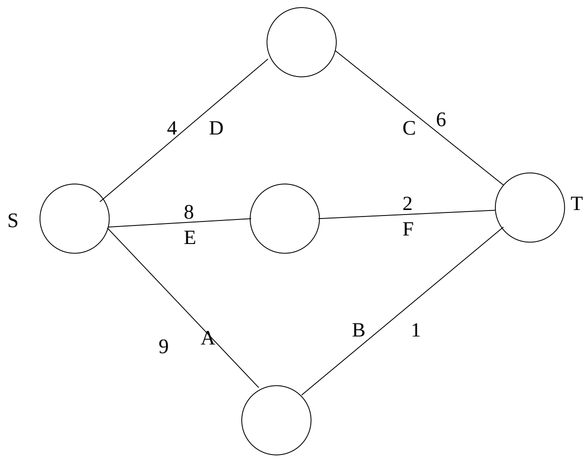

Figure 16.3: A simple graph with three paths of the same length from s to t.

We are given a graph and two distinct nodes s and t. Player 0 (the sender) wants to send a message from s to t. The other players are the owners of the edges (links). We may assume that each player owns one and only one link, although more general scenarios are possible.

Given an edge e, its owner has a cost c_e if edge e is used to route the message, and a cost of 0 is edge e is not used. Let us assume that player 0 has an initial amount of money K which she can use to pay link owners. If an edge e is used, its owner with receive a payment of p_e, which must be at least c_e. Under these assumptions, the payoff to the sender is:

$$p(0) = \begin{cases} K - \sum_{e \in S} p_e & \text{if the message is sent using the links from the set S.} \\ 0 & \text{if the message does not go through} \end{cases}$$

For the owners of the links, the payoffs are:

$$p(e) = \begin{cases} p_e - c_e & \text{if message gets sent using edge e.} \\ 0 & \text{otherwise} \end{cases}$$

The costs c_i are only known to the link owners who might try to misrepresent them in order to get a higher payoff. It is thus reasonable to design the game so that the chosen payment scheme induces the players to reveal their true costs. We consider the VCG payments and contrast them against the core of the game.

Let us define the coalitional game played by the six link owners and the sender for the graph of Figure 16.3. The set of players is $N = \{0, A, B, C, D, E, F\}$.

The characteristic function of the game can be easily defined:

- $v(\{i\}) = 0$, for all players i.
- $v(N) = K - 10$.
- For all $S \subset N$, if $0 \notin S$ then $v(S) = 0$.
- If S contains 0 and a valid path from s to t, then $v(S) = K - 10$.

If x is a core imputation, then we must have $x_A = \cdots = x_F = 0$. Indeed, if any of the link owners gets a nonzero payoff, then other link owners can start bargaining with the sender, and propose an imputation where both them and the sender are better off. Thus the core consists of the single imputation $(K - 10, 0, \ldots, 0)$.

The VCG payment for each player coincides with her marginal contribution to the game. Since this example consists of three disjoint paths of the same length, each edge has a *substitute*, so that its marginal contribution is 0, which makes its VCG payment 0. Thus the VCG payments to the link owners coincide with their payoffs in the unique core imputation.

Let us now consider a different instance of the shortest path game, as shown in Figure 16.4.

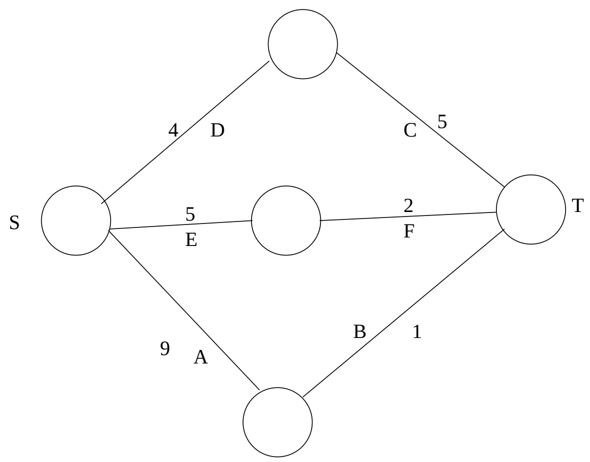

Figure 16.4: A simple graph with a unique shortest path between s and t.

Here the shortest path is unique. The characteristic function is defined by:

- $v(\{i\}) = 0$, for all players i.
- $v(N) = K - 7$.

- For all $S \subset N$, if $0 \notin S$ then $v(S) = 0$.
- $v(\{0, E, F\}) = K - 7$.
- $v(\{0, C, D\}) = K - 9$.
- $v(\{0, A, B\}) = K - 10$.

For any core imputation x, we have $x_A = x_B = x_C = x_D = 0$.

Since $V(\{0, C, D\}) = K - 9$, we have that the payoff to the sender must satisfy $x_0 \geq K - 9$, which in turn implies that $x_E + x_F \leq 2$. The core is thus described by the constraints:

- $x_0 \geq K - 9$,
- $x_E + x_F \leq 2$,
- $x_0 + x_E + x_F = K - 7$.

Consider now the VCG payments. As before, players $A \dots D$ do not provide a positive marginal contribution to the game, and so they get a payoff of 0. However, the marginal contribution of both player E and player F is 2, so that each of them gets a payment of 2, which leaves a payoff of $K - 11$ to the sender.

Therefore the VCG payments lie outside the core.

A third example is given by the graph of Figure 16.5.

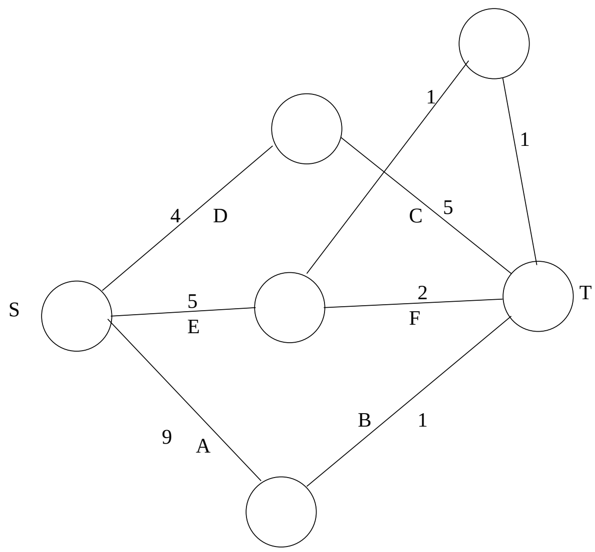

Figure 16.5: A simple graph where only one edge does not have substitutes.

Here link owner F has now substitutes, and thus her payoff in the core is 0. The only nonzero coordinates in the core are for players 0 and E: $x_E \leq 2$ and $x_0 \geq K - 9$.

The VCG payment is 2 for E and $K - 9$ for player 0, so that the VCG payment lies on the boundary of the core.

16.5 Frugality, Substitutes, and the Core

Let $N = \{0, 1, 2, \ldots, n\}$ be the set of players including the auctioneer (player 0). We define $v(S)$ for all $S \subseteq N$ as follows. If a set $S \subseteq N$ does not contain 0 or a feasible solution then $v(S) = 0$. Otherwise $v(S) = K - C(S)$, where $C(S)$ is the cost of the optimal solution entirely contained in S. In the specific case of the shortest path game, $C(S)$ is the cost of the shortest path computed according to the reported costs.

We assume that $v(\cdot)$ satisfies $v(S_1) \leq v(S_2)$ for $S_1 \subseteq S_2$ ("zero cost of disposal" property).

Let $\Theta_1, \ldots, \Theta_m$ be the set of optimal solutions, and let $\Theta = \bigcap_j \Theta_j$ be the set of agents in all the optimal solutions. For any imputation in the core only the agents in Θ receive a nonzero payoff. To see this, consider an agent a such that $a \in \Theta_i$ but $a \notin \Theta_j$. Then $v(N \setminus \{a\}) \geq v(\{0\} \cup \Theta_j) = v(N)$, where the inequality follows from the zero cost of disposal property, since $\{0\} \cup \Theta_j \subseteq N \setminus \{a\}$. Therefore $v(N) - v(N \setminus \{a\})$, the marginal contribution of player a, is zero, which implies that her VCG payoff is zero.

Definition 16.1 *We say* agents are substitutes *if* $\forall S \subseteq N$ *the following condition is satisfied*

$$v(N) - v(N \setminus S) \geq \sum_{i \in S} (v(N) - v(N \setminus \{i\})).$$

Definition 16.2 *Given a payoff vector* $x = (x_1, x_2, \ldots, x_n)$, *we define its frugality ratio as*

$$F(x) = \max_{\mathcal{O} \subseteq \Theta} \frac{\sum_{i \in \mathcal{O}} x_i}{v(N) - v(N \setminus \mathcal{O})}.$$

We say that x *is frugal if its frugality ratio satisfies* $F(x) \leq 1$.

Let π^V be the payoff induced by the VCG mechanism. We have:

$$\pi_i^V = \begin{cases} v(N) - v(N - \{i\}) & \text{if } i \in \Theta \\ 0 & \text{otherwise} \end{cases}$$

Theorem 16.3 π^V *is frugal if and only if the agents are substitutes.*

Proof: From the definitions, it is easy to see that the *agents are substitutes* property is sufficient for the frugality of π^V. Now assume that π^V is frugal. From the definition of VCG payments, we have that, for $A \subseteq \Theta$, $v(N) - v(N \setminus A) \geq \sum_{i \in A} (v(N) - v(N \setminus \{i\}))$. Let $K \subseteq N$ and $A = K \cap \Theta$.

$$\begin{aligned} v(N) - v(N \setminus K) &\geq v(N) - v(N \setminus A) \\ &\geq \sum_{i \in A} (v(N) - v(N \setminus \{i\})) \\ &= \sum_{i \in K} (v(N) - v(N \setminus \{i\})). \end{aligned}$$

Note that the first inequality follows from "zero cost of disposal" property, while the equality depends on the fact that the marginal contribution of players who do not belong to Θ is zero. ∎

In the previous section, we have seen that the VCG payments may or may not lie in the core of the game. Here we give a general answer to this question.

Given a set of payoffs π_j to players $j \in N \setminus \{0\}$, the corresponding payoff to player 0 is $\pi_0 = K - C - \sum_{j \in N \setminus \{0\}} \pi_j$.

For two allocations x^1 and x^2 in the core, we now define their *max* $\bar{x}$ as follows:

- $\forall i \in N \setminus \{0\}$, $\bar{x}_i = \max\{x_i^1, x_i^2\}$;
- $\bar{x}_0 = K - C - \sum_{j \in N \setminus \{0\}} \bar{x}_j$.

It is easy to show that $\bar{x}_0$ is always nonnegative. Furthermore, it is possible to prove that, if agents are substitutes, then if x^1 and x^2 are in the core then $\bar{x}$ is in the core as well (see [4, 5]). Therefore, if agents are substitutes, then the max of two core payoffs is in the core.

We can similarly define the min of two core payoffs (which we denote by $\underline{x}$) as follows:

- $\forall i \in N \setminus \{0\}$, $\underline{x_i} = \min\{x_i^1, x_i^2\}$;
- $\underline{x}_0 = K - C - \sum_{j \in N \setminus \{0\}} \underline{x}_j$.

We now prove that $\underline{x}$ is (unconditionally) in the core.

Theorem 16.4 *If x^1 and x^2 are in the core then $\underline{x}$ is in the core.*

Proof:

By definition of $\underline{x}_0$, we have that $\underline{x}(N) = v(N)$. We now show that $\forall S \subseteq N$, $\underline{x}(S) \geq v(S)$. Let $\Theta \cap S = A$. We have

$$
\begin{aligned}
\underline{x}(S) &= \underline{x}_0 + \sum_{i \in A} \underline{x}_i + \sum_{i \in S-A} \underline{x}_i \\
&= K - C - \sum_{i \in \Theta} \underline{x}_i + \sum_{i \in A} \underline{x}_i \\
&= K - C - \sum_{i \in \Theta - A} \underline{x}_i \ \geq \\
&\geq \ K - C - \sum_{i \in \Theta - A} x_i^1 \ = \ x^1(S) \ \geq \ v(S),
\end{aligned}
$$

where the equality $K - C - \sum_{i \in \Theta - A} x_i^1 \ = \ x^1(S)$ follows from the fact that

$$
\begin{aligned}
x^1(S) &= x_0^1 + x(A) \\
&= K - C - \sum_{i \in \Theta} x_i^1 + \sum_{i \in A} x_i^1 \\
&= K - C - \sum_{i \in \Theta - A} x_i^1.
\end{aligned}
$$

■

Summarizing, we have seen that if x^1 and x^2 are in the core and agents are substitutes, then $\bar{x}$ and $\underline{x}$ are in the core.

The following Lemma shows that the converse also holds.

Lemma 16.5 *If for every x^1 and x^2 in the core, $\bar{x}$ and $\underline{x}$ are also in the core, then agents are substitutes.*

Proof:

For all $i \in N \setminus \{0\}$, let s^i be the payoff vector which is nonzero only in positions 0 and i, i.e., $s_i^i = v(N) - v(N \setminus \{i\})$, and $s_0^i = K - C - s_i^i$.

It is immediate to see that s^i is in the core of the game. We now repeatedly apply the assumption of the Lemma (i.e., that if x^1 and x^2 are in the core then $\bar{x}$ is in the core) to prove that the vector $s^{N\setminus\{0\}}$, whose k-th component is s_k^k, for $k \geq 1$, and whose 0 component is $K - C - \sum_i s_i^i$, is in the core.

Consider now any $S \subseteq N$ such that $\Theta \subseteq S$, and let $A = S \cap \Theta$. By definition of core, we have that

$$\sum_{j \in S} s_j^{N\setminus\{0\}} \geq v(S)\,,$$

from which we get

$$s_0^{N\setminus\{0\}} + \sum_{j \in S\setminus\{0\}} s_j^{N\setminus\{0\}} \geq v(S)\,,$$

which in turn means that

$$K - C - \sum_{j \in N\setminus\{0\}} s_j^{N\setminus\{0\}} + \sum_{j \in S\setminus\{0\}} s_j^{N\setminus\{0\}} \geq v(S)\,.$$

But now

$$v(N) - \sum_{j \in N\setminus S} s_j^{N\setminus\{0\}} \geq v(S),$$

which can be rewritten as

$$v(N) - v(S) \geq \sum_{j \in N\setminus S} \left(v(N) - v(N \setminus \{j\})\right),$$

i.e., the agents are substitutes. ■

We say that $\pi^1 \preceq \pi^2$ if $\pi_i^1 \leq \pi_i^2$ for all $i \in N \setminus \{0\}$.

Definition 16.6 (Lattice) *A lattice is a partially ordered set in which any two elements have a supremum and an infimum.*

Theorem 16.7 Agents are substitutes *if and only if the core is a lattice, w.r.t. the $\preceq$ relation.*

Proof: The claim can be proved by first showing that the core is a lattice if and only if x^1 and x^2 in the core implies that $\bar{x}$ and $\underline{x}$ are in the core. The details of the proof can be found in [6]. ■

It is then immediate to see that, when the core is a lattice, the VCG payment is the supremum element of the lattice.

Bibliographic notes

The application of the VCG mechanism to the shortest path game has been studied by Nisan and Ronen in [12].

The computation of the VCG payments makes it necessary to compute the optimal solution to all the subgames obtained by excluding each of the agents in turn. Thus Nisan and Ronen [11, 12] raised the question of whether or not, for the shortest path game, it is possible to compute all the required shortest paths at the same asymptotic cost required to compute the shortest path for the original problem. This question was answered in the affirmative in [7].

Motivated by the overpayments induced by the VCG mechanism, Archer and Tardos [1] considered the *frugal path mechanism problem*, where they sought the existence of truthful mechanisms not incurring into the above overpayments. They defined a pretty general family of mechanisms, but they obtained negative results which essentially show that overpayments can not be avoided.

The results of the last section have been presented in [6].

References

[1] A. Archer, E. Tardos, Frugal Path mechanisms, to appear in Transaction on Algorithms. Preliminary version appeared in the Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2002, pp. 991-999..

[2] A. Archer, E. Tardos, Truthful mechanisms for one-parameter agents, Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (2001) 482-491.

[3] A. Archer, C. Papadimitriou, K. Talwar, E. Tardos, An approximate truthful mechanism for combinatorial auctions with single parameter agents, Internet Mathematics, Volume 1 Issue 2, 2004. Prleiminary version appeared in the Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2003, pp. 205 - 214.

[4] S.Bikhchandani, J.M.Ostroy, The Package Assignment Model, Journal of Economic Theory, 107(2):377-406, 2002.

[5] S. Bikhchandani, S.E. Vries, J. Schummer, R.V. Vohra, Linear Programming and Vickrey Prices, in Mathematics of the Internet: E-auction and Markets, B. Dietrich and R. Vohra eds., Springer-Verlag, New York, 2002, pp. 75–116.

[6] R. Garg, V. Kumar, A. Rudra, A. Verma, Coalitional games on graphs: core structure, substitutes and frugality, ACM Conference on Electronic Commerce, 2003.

[7] J. Hershberger, S. Suri, Vickrey Prices and Shortest Paths: What is an Edge Worth? Proceedings of the 42nd IEEE symposium on Foundations of Computer Science (2001).

[8] A. Mas-Colell, M.D. Whinston, J.R. Green, Microeconomic Theory, Oxford University Press (1995).

[9] N. Nisan, Algorithms for selfish agents, in Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, 1–15 (1999).

[10] N. Nisan, Introduction to Mechanism Design (for Computer Scientists). In Algorithmic Game Theory by N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani, editors, September 2007.

[11] N. Nisan, A. Ronen, Computationally feasible VCG mechanisms, *ACM Conference on Electronic Commerce*, 2000.

[12] N. Nisan, A. Ronen, Algorithmic Mechanism Design, Games and Economic Behavior 35, pp. 166–196 (2001). Extended abstract appeared at the Thirty-First Annual ACM Symposium on Theory of Computing (STOC 99).

[13] K. Talwar, The price of truth: Frugality in truthful mechanisms, STACS 2003.

[14] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, Journal of Finance 16 (1961), 8–37.