



HAL
open science

Self-managed Inter-domain pricing: An overview and discussion of possible approaches

Hoang-Hai Tran, Bruno Tuffin

► **To cite this version:**

Hoang-Hai Tran, Bruno Tuffin. Self-managed Inter-domain pricing: An overview and discussion of possible approaches. RR-6992, 2009. inria-00403953v4

HAL Id: inria-00403953

<https://inria.hal.science/inria-00403953v4>

Submitted on 28 Oct 2009 (v4), last revised 27 Nov 2009 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Self-managed Inter-domain pricing: An overview
and discussion of possible approaches*

Hoang-Hai Tran — Bruno Tuffin

N° 6992

July 2009

Thème COM

*R*apport
de recherche

Self-managed Inter-domain pricing: An overview and discussion of possible approaches

Hoang-Hai Tran*, Bruno Tuffin†

Thème COM — Systèmes communicants
Équipe-Projet DIONYSOS

Rapport de recherche n° 6992 — July 2009 — ?? pages

Abstract: Telecommunication networks are now an interconnection of competitive operators that need to cooperate to ensure end-to-end delivery of traffic. Inter-domain agreements have to be performed, and pricing is seen as a relevant way to reward intermediate domains for forwarding the traffic of others and such that the network can be efficiently self-managed, in a way as decentralized as possible. In the devoted literature, Vickrey-Clark-Groves (VCG) auctions have been highly considered because they provide proper incentives, lead to an efficient use of the network and verify other relevant characteristics. On the other hand, it has been highlighted in this specific context that they are neither budget-balanced nor robust to collusion. After reviewing the list of properties that we ideally would like the Inter-domain pricing to verify and the fact that no mechanism can ensure them all together, we discuss the properties that could be relaxed. Furthermore, we present related works on pricing and resource allocation not especially devoted to inter-domain pricing, and discuss how they can be applied to our specific problem and what are the corresponding results.

Key-words: Internet pricing, game theory, auctions, network management, resource allocation

* thoangha@irisa.fr

† bruno.tuffin@irisa.fr

Tarification inter-domaine auto-gérée : une vue d'ensemble et discussion des approches possibles

Résumé : Les réseaux de télécommunication sont maintenant une interconnection d'opérateurs compétitifs ayant besoin de coopérer pour assurer la livraison point-à-point de leur trafic. Des accords inter-domaines doivent être réalisés, et la tarification semble être une méthode pertinente de récompenser les domaines intermédiaires transférant le trafic des autres et telle que le réseau soit efficacement auto-géré, de manière aussi décentralisée que possible. Dans la littérature, les enchères Vickrey-Clark-Groves (VCG) ont fait l'objet d'une attention particulière car elles fournissent les bonnes incitations, conduisent à une utilisation efficace du réseau, et vérifient également d'autres propriétés intéressantes. Cependant il a été établi pour ce contexte spécifique qu'elles ne sont ni équilibrées sur le plan budgétaire, ni résistantes à la collusion entre opérateurs. Après avoir passé en revue la liste des propriétés qu'idéalement nous souhaiterions voir vérifiées par la tarification inter-domaine, nous établissons qu'elles ne peuvent être obtenues toutes en même temps. Nous discutons alors les propriétés qui peuvent être relâchées. Nous présentons également des travaux proches sur la tarification et l'allocation de ressources non spécifiquement consacrés à la tarification inter-domaine, et discutons leur application à ce problème.

Mots-clés : Tarification, théorie des jeux, enchères, gestion des réseaux, allocation de ressources

1 Introduction

The Internet has evolved from an academic (and cooperative) network into a highly competitive inter-connected one with many service providers trying to optimize their own profit, even if at the expense of the others. The question is therefore to better understand the (economic) relations between the different domains, also called Autonomous Systems (ASs) and their implications from an engineering point of view. The Internet consists of thousands interconnected-AS defined as different Internet Service Providers (ISPs). Up to now, routing between different ASs are managed by the *Border Gateway Protocol* (BGP) or *Exterior Gateway Protocol* (EGP), and *Interior Gateway Protocols* (IGPs) handle traffic within a single AS. Inter-domain has indeed been studied for many years, from an engineering or “protocol design” point of view, but the competition and selfishness issues have more recently shifted it to a more economic or “mechanism design” oriented point of view[?, ?]. The basic tool is *game theory*, a branch of applied mathematics whose motivation is to mathematically investigate agent’s behavior in strategic situations. Game theory has been applied to networking and telecommunication to solve routing and allocation problems in a competitive environment [?]. It can also be used to see if rewarding intermediate domains for forwarding the traffic of others can help to efficiently self-manage inter-domain relationships, in a way as decentralized as possible [?]. Our work is therefore motivated by studying Inter-domain routing and resource (mainly bandwidth) allocation from an economic point of view. The issues that are raised and which require to be solved are summarized as follows:

- The ASes need to exchange traffic on behalf of their customers, and transparently to them, in order to provide end-to-end service. This requires other providers to agree to forward that traffic. However, the providers are strategic (or selfish) agents and they usually behave according to their best interest. Why would they use their own resources for competitors without a proper economic incentives? Similarly, an AS may discriminate the traffic that originate from its own customers and against the traffic from customers of its competitors. The existing routing protocols have not dealt with those situations in which the providers do not follow the cooperative rules.
- Defining economic incentives to participate and forward traffic helps to create a “working” network, but we also expect it to be efficient in the sense that the resulting allocation of resources is optimal with respect to the sum of agents’ true preferences, also called social welfare.
- Another major issue with a general network topology is that determining the optimal resource allocation, calculating the cheapest paths, computing the price paid for each provider, and the maximization of social welfare, are computationally demanding and may even be generally NP-hard problems. A goal would be to provide a mechanism which is self-managed, decentralized, and each non-cooperative AS’s own interest is to behave in the general interest.

This paper aims at reviewing the problem of inter-domain pricing. In the literature, the main family of work is based on Vickrey-Clarke-Groves (VCG)

auctions, which have the properties to be incentive compatible and efficient, meaning that each agent reveals his real valuation for bandwidth, and an optimal resource (bandwidth) allocation is derived, in a decentralized way. This has been successfully applied to inter-domain pricing in the seminal work [?] using an adaptation of BGP. On the other hand, some important issues have been forgotten. First of all the pricing scheme is not budget-balance, meaning that the sum of revenues is larger than the sum of charges. As a consequence, a central authority (a government?) has to permanently inject money in the system to make it work. This is very unlikely. Actually, it can be proved that it is impossible to design a mechanism verifying at the same time incentive compatibility, efficiency, individual rationality (no agent will get a negative utility, that is each agent has interest in participating into the game) and budget balance. For this reason, we list as exhaustively as possible the various properties that an inter-domain mechanism should verify in order to see which set would be the most relevant (as it is impossible to have them all together). We then review all the related works on pricing for multiple objects that can be related to inter-domain pricing, and organize them in families verifying given set of properties. We additionally apply those schemes specifically to our Inter-domain model with its specificities, to investigate their behavior and how the results can then be adapted. From those results, we are able to discuss the most relevant pricing scheme for inter-domain.

2 Mechanism design and desirable properties

In this section, we make a short introduction of mechanism design and review the properties that are the most relevant to Inter-domain pricing.

2.1 Mechanism design

Mechanism design theory provides us a tool for studying and designing a distributed algorithm for Inter-domain pricing. The mechanism is designed in a way that all agents participated in a game should be forced or at least incentive to act according to the rules of the designer. For more detail and understanding of mechanism design, we recommend to the readers [?][?] and the references therein.

First, let us bring some basic notations from mechanism design theory that will be helpful to understand the main results. Assume there is n agents in set $N = \{1, \dots, n\}$ and each agent has a choice within a set of action \mathcal{A}_i . Let $x_i \in \mathcal{A}_i$ be the action of agent i ($1 \leq i \leq n$), and define (x_1, \dots, x_n) as the strategy profile. To this set of actions, the designer maps an outcome vector $\theta = (\theta_1, \dots, \theta_n)$, where θ_i represents the outcome (or allocation) of agent i . Similarly, with each outcome is associated a payment vector $\rho = (\rho_1, \dots, \rho_n)$, with ρ_i the payment of agent i . $\rho_i < 0$ means that agent i is rewarded for his action, while $\rho_i > 0$ means that he has to pay for it. Each agent i has his own preferences over outcomes, reflected by his valuation function V_i . Valuation function v_i assign a non-negative, real number $V_i(\theta_i)$ over each possible outcome θ_i . Each agent is assumed to act selfishly and strives to maximize his own *profit* or *utility* and U_i is the utility function of player i . The utility is assumed to be quasi-linear between *valuation* V_i for the *outcome* θ_i , and the *payment* $\rho_i(\theta_i)$,

i.e.,

$$U_i(\theta_i) = V_i(\theta_i) - \rho_i(\theta_i).$$

Since the outcome depends on the whole action profile, the underlined framework is that of *non-cooperative game theory*, and the equilibrium notion that of *Nash equilibrium*. A Nash equilibrium is a strategy profile (x_1^*, \dots, x_n^*) , to which is associated an outcome $\theta^* = (\theta_1^*, \dots, \theta_n^*)$ and a payment $\rho^* = (\rho_1^*, \dots, \rho_n^*)$ such that no agent can increase his utility by *unilaterally* changing his action. Formally, we have a Nash equilibrium if $\forall 1 \leq i \leq n$, if $x = ((x_1^*, \dots, x_{i-1}^*, x_i, x_{i+1}^*, \dots, x_n^*))$ with corresponding outcome θ (and payment ρ), we have $U_i(\theta^*) \geq U_i(\theta)$. A notion related to Nash equilibrium is the concept of *dominant strategy*. A strategy x_i is called to be a dominant strategy if the strategy earns agent i a larger payoff than any other, whatever the strategies of others. It is obvious that a strategy profile (x_1, \dots, x_n) in which every x_i is dominant for player i is a Nash equilibrium.

It may happen that no pure Nash equilibrium exists, and the concept of Nash equilibrium *in mixed strategies* comes in, where instead of choosing a strategy, each agent plays with a probability distribution over the the set of strategies, from which the eventual one is randomly chosen. A *Nash equilibrium in mixed strategies* is a profile of probability distributions over which no agent can unilaterally increase his expected utility (where the expectation over the whole profile of distributions [?]).

Bayesian games are games for which information about characteristics of the other players (their utilities) is incomplete, a common and relevant assumption in inter-domain. Formally, a Bayesian game with n players is defined by

- A set $N = \{1, \dots, n\}$ of n agents/players;
- A set Ω of states describing all the players' relevant characteristics;
- A space of strategies $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ where $x_i \in \mathcal{A}_i$ is the set of actions available to agent $i \in N$.
- A finite set T_i of *signals, or types*, that player i can observe due to the incomplete information, and a *signal function* $\tau_i : \Omega \rightarrow T_i$ mapping each state to a signal.
- A probability measure p_i of *prior beliefs* for each agent i , defined over Ω and such that $p_i(\tau_i^{-1}(t_i)) > 0$ for all $t_i \in T_i$.
- Utility functions defined as in the above normal form of the game.

Each strategy profile $x = \{x_1, \dots, x_n\}$ leads to an outcome $\theta = \{\theta_1, \dots, \theta_n\}$ and a payment vector $\rho = \{\rho_1, \dots, \rho_n\}$. The payment ρ_i is given to agent i to incentive him to behave according to the rules for the overall goals. A strategy-proof mechanism is one in which truthful report of an agent is always *dominant strategy*. A stronger form of mechanism called strategy-proofness in which no coalition of agents has incentive to misreport its willingness to pay (its true valuation).

Each player knows his type, but needs to hold beliefs about actions of other players, but this belief may depend the player i 's own actions. A Nash equilibrium in a Bayesian is a Nash equilibrium where each expected utilities are used, with *posterior belief* distributions $p_i(\omega)/p_i((\tau_i^{-1}(t_i))$ for each player i .

Those general notions of game and equilibrium will be useful later on to analyze the inter-actions of ASes and design a proper mechanism. Indeed, the goal of mechanism design is to build a mechanism (that is an allocation/outcome mapping from the strategy profile, and a payment rule) with some desired properties among the ones listed below.

2.2 Properties

What is the list of properties that can be of interest for a mechanism designer (and in particular in the case of inter-domain pricing)? We now review them in detail.

1. *Individual Rationality*: a mechanism is said to be individually rational if and only if rational playing by agents does not lead to any agents being worse-off after the game. More specifically, the mechanism is individually rational if and only if for all $1 \leq i \leq n$ and all θ , $U_i(\theta) \geq 0$.
2. *Truthfulness, also called Incentive Compatibility*: Truthfulness is hold when agents truthfully disclose their preferences to the mechanism in equilibrium. A mechanism is said to be truthful if every agent i 's best interest is to play or submit a strategy that corresponds to the true valuation agent i has for the service. A precise and more formal definition is difficult in a general context but will be made clear for several mechanisms. As a short illustration, in auctions, the principle would be to bid the real valuation of the player for the quantity asked, the real cost when buying or sell something.
3. *Budget balance*: Besides allocating the goods, the mechanism collects and in some cases distributes the money from buying agents and to selling agents. Thus, the mechanism is said to be budget-balanced if the mechanism guarantees not to lose money. Formally, the mechanism is budget-balanced if and only if

$$\sum_{1 \leq i \leq n} \rho_i(\theta) = 0.$$

The mechanism will be said to be *weakly budget-balanced* if the sum of money exchanges is non-negative, i.e., if $\sum_{1 \leq i \leq n} \rho_i(\theta) \geq 0$. As we will see

later in the paper, some mechanisms appealing *a priori* actually verify $\sum_{1 \leq i \leq n} \rho_i(\theta) < 0$, meaning that they require an authority injecting money permanently.

4. *Allocative Efficiency*: A mechanism is called efficient the mechanism maximizes the total utility gain or social welfare of all agents involved in the game when they implement at the final Nash equilibrium. In other words, the mechanism enforces the agents to play strategies such that an outcome is the same as the maximization in the cooperative case. Formally, the mechanism is efficient if it leads to an outcome θ^* such that

$$\theta^* = \arg \max_{\theta} \sum_{i \in N} V_i(\theta_i).$$

When the mechanism is budget-balanced, the maximization of social welfare is equivalent to $\max \sum_{1 \leq i \leq n} U_i(x_i)$.

5. *Decentralization*: This property is related to scalability issue, so that the mechanism is robust in terms of computational time to an increasing number of agents. Decentralization basically means that decisions can be spread all over agents and there is no need to concentrate them at a central authority. Decisions can therefore be made locally with a limited information.
6. *Robustness to collusion*: due to the game-theoretic framework, it may happen that agents have can increase their utility if they collude with other agents in a group. Each agent in the group could then gain from this situation, at the expense of global social welfare. Designing a mechanism for which collusion is never profitable is a desirable property.

2.3 Feasibility issue

While designing a mechanism verifying these properties all together would be expected, it has been shown that verifying a subgroup of them, namely individual rationality (IR), incentive-compatibility (IC), efficiency and budget-balance (BB), is already not possible [?, ?].

Theorem 1. [?, ?, ?] *No dominant-strategy incentive-compatible mechanism is always both efficient and weakly budget balanced, even if agents are restricted to the simple exchange setting.*

Theorem 2. [?, ?] *No Bayes-Nash incentive-compatible mechanism is always simultaneously efficient, weakly budget balanced, and interim individually rational, even if agents are restricted to quasi-linear utility functions.*¹

Following the two theorems, it turns out that it is impossible to design a mechanism that simultaneously achieves incentive-compatibility, efficiency, budget-balance and individually rationality. The challenges are now for us

1. To model our inter-domain pricing problem,
2. to investigate the properties verified or not by the existing propositions,
3. to discuss the most relevant properties and those which can be relaxed.

An important remark from the literature is that this impossibility to verify all properties has driven to an *almost-the property* framework where the idea is to relax one of the properties, but to be "close" to it. This is for instance the case for incentive-compatibility in [?, ?], and for budget balance in [?, ?].

¹ At the interim stage, each agent can decide whether or not to participate in the mechanism and the decision has to be made without knowing other private information. The agent will decide to participate in the game if his expected net-payoff is non-negative. However, It is possible in interim stage that his net pay-off can be negative. On the other hand, ex-post IR allow agents to learn the term of the game before deciding whether to participate in the game or not, and his net pay-off is always positive. Hence, no agent will ever regret the choice he made, and it is obvious that ex-post IR is stronger condition than Interim IR.

3 Inter-domain problem and modeling

We now specifically define our inter-domain pricing problem, in order investigate the properties verified by the mechanisms from the literature, and of other propositions not already applied to that problem, but closely related. We begin with the description of Border Gateway Protocol (BGP), the standard protocol currently used for Inter-domain. Afterward, the Inter-domain resource allocation problem and mathematical modeling are presented.

3.1 Inter-domain routing: the Border Gateway Protocol

The Border Gateway Protocol (BGP) is currently the core routing protocol of the Internet [?]. It maintains a table of IP networks which designates network reachable among ASes. It is a path vector protocol and makes routing decisions based on path, next hop policies. The basic idea of BGP is to build a routing tree from a source node to a given destination hop-by-hop. Communication between nodes is made through *update messages* [?]. The *update messages* give nodes a way to construct a consistent view of the network topology. The process is initialized when a destination node d announces itself to its neighbors by sending update messages. For each source node i , the route to d is iteratively established as follows:

- Importing routes to d from its neighbor nodes via update messages.
- Choosing the best route to d through one of its neighbor based on routing metric. BGP only uses single routing metric to determine the best path to a given destination node. This metric is assigned to each link by the network administrator which may be based on delay, cost or speed.
- If there is any change in the route to d , i updates its routing table via new update messages.

Each node therefore maintains its routing table which contains all feasible, shortest paths to a particular AS, transmits or receives update messages and makes routing decisions. In BGP, nodes do not regularly update their routing table, only in case there is an update for an optimal path to a network.

3.2 Inter-domain resource allocation problem modeling

The network is modeled as a graph $G = (V, L)$, where nodes in V are the ASes and links in L are the existing direct links between ASes. We assume for now that there is no capacity constraints. There is a set \mathcal{R} of routes. Each route $r \in \mathcal{R}$ is defined as an ordered list of nodes, each node appearing only once, and such that for two successive nodes in the list there exists a link $\ell \in L$ between those two nodes. Define also for node $v \in V$, $\mathcal{R}^S(v)$ the subset of routes starting at v and $\mathcal{R}^D(v)$ the subset of routes ending at v .

Define y_r the flow for route r and y the vector of all flows $y = (y_r)_{r \in \mathcal{R}}$. The total rate initiated in node v is

$$x_v^I = \sum_{r \in \mathcal{R}^S(v)} y_r,$$

and the flow transferred through v is

$$x_v^T = \sum_{r \in \mathcal{R} \setminus (\mathcal{R}^S(v) \cup \mathcal{R}^D(v))} y_r.$$

From those allocated rates, the utility of an AS v is given by

$$U_v(y) = V_v(x_v^I) - \sum_{r \in \mathcal{R}^S(v)} C_r^I(v) + \sum_{r \in \mathcal{R} \setminus (\mathcal{R}^S(v) \cup \mathcal{R}^D(v))} C_r^T(v) \quad (1)$$

where

- $V_v(x_v^I)$ is the valuation of AS v for sending traffic x_v^I , which deals with the relation its customers. It is defined here as the valuation of the total rate because we assume that the traffic value is the same on all routes. If differentiation needs to be done, this component can be replaced by $\sum_{r \in \mathcal{R}^S(v)} V_{v,r}(y_r)$ to provide different valuation function for the different routes.
- $C_r^I(v)$ is the total charge imposed to v for sending traffic on route r and initiated from v . The charge can be linear in the traffic $C_r^I(v) = c_r^I(v)y_r$, but not necessarily. Similarly, depending on the mechanism applied, the charge can be the total charge from all the intermediate nodes, or the charge from neighbors only, which are themselves charged from the next hops and then passes it on the sender (what will be typically done when using BGP).
- $C_r^T(v)$ is the total reward for traffic forwarded by v on route r . Here again, it can be linear in y_r , i.e., $C_r^T(v) = c_r^T(v)y_r$, and can be directly charged to the initiator, or, for decentralized purposes, it can be the difference between the charge to the predecessor on the route and the charge from the successor (that v is paying).

Note that two different routes can have the same source and destination. Therefore, the model is general here in the sense that (i) multi-homing among several routes is possible (ii) or routing is fixed because only one route is proposed in \mathcal{R} between each possible pair of ASes (iii) or several routes are proposed but at the end of the allocation process, only one potentially gets a strictly positive flow. Constraints can be imposed without loss of difficulty to let those choices possible.

An important remark is that utilities depend here on allocations y_r on all routes r and charges, i.e., on allocation and charging rules. This is therefore dependent of *the outcome of the game* where each AS strives to maximize its utility function. Similarly, we do not explicitly describe how ASes play, i.e., what is their action space, because it is mechanism-dependent and will be described in next sections. Basically though, they can play on the charge for transferring traffic of competitors and/or the quantity they want to send on each route (initiated from them or not) and the price they would accept to pay for it. Several situations can indeed be envisaged, and they will be explicitly defined when describing the different variants/mechanisms in the remainder of the paper.

4 The popular mechanism: Vickrey Clarke Groves (VCG) auction

4.1 VCG auction and Inter-domain

Vickrey original auction mechanism for a single item and multiple identical items provides an efficient allocation mechanism in which the dominant strategy of all participants is to bid truthfully. Vickrey auction for a single item is often referred to *the second-price sealed-bid auction* in which the highest bidder wins the item and pays the second highest bid. Vickrey original auction has been merged with the Clarke-Groves design for heterogeneous goods and does not require the bidders have non-increasing marginal values. In the resulting mechanism, called Vickrey-Clarke-Groves or VCG mechanism [?, ?, ?], the auctioneer allocates the goods to the highest bids. However, rather than paying the price he bids for his units, a winning bidder pays the *opportunity cost* for the units won [?]. In other words, what you pay is the sum of declared valuations when you are in the game minus the sum of declared valuations for the outcome when you are *excluded* from the game.

Theorem 3. [?] *Truthful reporting is a dominant strategy for each bidder in the VCG mechanism. Moreover, when each bidder reports truthfully, the outcome of the mechanism maximizes social welfare. It is the only mechanism [?] which verifies in general incentive compatibility, individual rationality and efficiency at the same time.*

This theorem makes VCG particularly attractive for resource allocation problems with distributed computing. To illustrate this, truthful routing in algorithmic mechanism design has been first investigated from the seminal paper of Nisan and Ronen [?]. They provide a polynomial-time, strategy-proof mechanism for cheapest route selection in a centralized model. They formulate the communication network by a directed graph G , and focus on a route between a source and a destination. Each edge e of the graph is a strategic agent and the algorithm ensures that when agents honestly report their cost c_e (by applying VCG), the cheapest path is selected. The payment given to each edge e on the cheapest path is $p_e(y_r) = \sum_w c_w 1_{(w \in (y_{-e}^*))} - \sum_{w \neq e} c_w 1_{(w \in (y^*))}$ and 0 if not on this path. $\sum_w c_w 1_{(w \in (y_{-e}^*))}$ is the cost of the cheapest path which does not contain e and $\sum_{w \neq e} c_w 1_{(w \in (y^*))}$ is the cost of the cheapest path including e , but not counting its cost. The given mechanism is a pure VCG mechanism. This preliminary suggestion is improved on the two papers of J. Feigenbaum et al. toward incentive-compatible Inter-domain routing [?, ?] described now.

Let us now explicitly define VCG auctions in this inter-domain context. We start by illustrating the problem with the seminal work of Feigenbaum in this area. In this work, the amounts of traffic x_v^T initiated from each node v are fixed. Each node is then seen as an intermediate node and his action is to define the per unit of volume cost c_v^T that he declares for transferring traffic. The goal of the mechanism designer is to minimize the *real total cost of routing*

$$\sum_v \bar{c}_v^T x_v^T, \quad (2)$$

where \bar{c}_v^T designs the real cost (not necessarily the same than the declared one). This fits our general model by saying that the valuation $V_v^I(x_v^I)$ of node v is $-c_v^T y_r 1_{(v \in r \text{ and } y_r > 0)}$ (that is, represents his costs) and then $\sum_v V_v^I(x_v^I) = -\sum_v \bar{c}_v^T x_v^T$ minimizing routing of cost is the same as maximizing the social welfare, i.e., the sum of valuations of all actors in the game. All nodes are assumed here to be *intermediate nodes*, which means that $C_r^I(v) = 0 \forall r, v$. The transfer payments $C_r^T(v) = p_v^T y_r$ (with unit price p_v^T) are defined below. The allocation rule of the mechanism allocates flows to the route that maximize the declared total cost of routing

$$y^* = \operatorname{argmax}_y \sum_v \bar{c}_v^T x_v^T.$$

To represent formally the payment rule, define y_{-v}^* as the allocation if AS v is not part of the game (that is maximizing the social welfare of the same game, but without v , which is equivalent to considering an infinite cost for v). Then the payment of node v is the opportunity cost imposed by his presence, i.e.,

$$p_v^T = c_v^T 1_{(v \in y_r^* \text{ and } y_r^* > 0)} + \sum_w c_w^T 1_{(w \in (y_{-v}^*)_r \text{ and } (y_{-v}^*)_r > 0)} - \sum_w c_w^T 1_{(w \in y_r^* \text{ and } y_r^* > 0)}. \quad (3)$$

VCG actually drives the ASes to declare their real costs $c_v^T = \bar{c}_v^T$ for sending traffic, and therefore yields an allocation that minimizes the total (real) cost of routing, verifying at the same time incentive compatibility [?, ?]. The BGP-based computational model and distribution of price computation is described. In Feigenbaum model, that the (cost-based) routing table of each provider is determined from the transit costs for each neighbor. The source provider uses the declared cost from intermediate providers to find the optimal path to destination. Whenever there are some changes in network topology, all providers will update their routing tables. When dealing with auctions, the computational outcome of the game is often seen as a linear program, or more generally as an optimization under constraints, and is computationally hard. In [?, ?], there is no bandwidth capacities or budget constraints for players, leading to a simpler computation. The complexity in this model comes from the distributed price computation due to BGP-compatible motivation. Complexity issue will be discussed more in next section.

Consider an illustration of this mechanism. Figure ?? represents an Inter-domain topology, in which a source ISP would like to forward its traffic/flow x_{source}^I to the destination node through several intermediate ASes.

On the Figure, transit costs for the source ISP are expressed for each intermediate node above the node. On this topology, using DFS (Deep First Search) [?], there are 6 possible paths that can carry the traffic flow y_r to destination. The loop paths and the paths having more than two hops are not considered. The path minimizing the total (declared) routing cost is *Source-C-E-Destination* having total cost 5 (or total valuation -5) for both C and E , see the Table 1 for more details. To compute the price paid to node C for instance, remark that if C is excluded from the game, the cost minimizing path becomes *Source-B-E-Destination* with total unit cost 7. Therefore, the payment to C is $p_C = 2 + (7 - 5) = 4$. In a similar way, the payment to E is 5 per unit.




Figure 1: Simple Inter-domain Topology

Possible path	Total cost
<i>Source - A - D - Destination</i>	9
<i>Source - A - E - Destination</i>	8
<i>Source - B - E - Destination</i>	7
<i>Source - B - F - Destination</i>	10
<i>Source - C - E - Destination</i>	5
<i>Source - C - F - Destination</i>	8

Table 1: Possible paths between source and destination provider with their associated cost (remember that valuation is minus cost)

Node	declared cost	payment
<i>A</i>	5	0
<i>B</i>	4	0
<i>C</i>	2	4
<i>D</i>	4	0
<i>E</i>	3	5
<i>F</i>	6	0

Table 2: Costs and payments to nodes for the example of Figure ???. Costs presented here are only those really applied, that is for the allocated route.

Table ?? displays the payments associated to all nodes. The reader can check the values.

Remark that with our very general model, the strategy choice can take different forms. Instead of working with costs (only) and playing on the declared cost, we may add to the valuations V_v^I a function due to the valuation from customers, and then play in addition on the amount of traffic the user wishes to send. Opportunity costs can be computed as well. Many other forms of game can be imagined, and some will be presented later on.

4.2 VCG auction applied to other networking problems

As we have emphasized, VCG is a powerful mechanism which gives us a way to construct a truthful, dominant strategy mechanism. It can maximize total social welfare and gives the bidders the incentive to bid their true valuations for their services. As a result, VCG have been widely applied in various distributed applications which require the agents to interact with each others to achieve a global goal.

Another work by J. Feigenbaum et al. is applied VCG mechanism on Multicast cost sharing [?]. In their papers, the authors make use of a previous work of H. Moulin [?] in order to design a distributed algorithm for multi-cast cost sharing, namely Marginal Cost (MC) Pricing and Shapley Value mechanism [?]. MC uses VCG payment rules to charge the cost share. In the case of a symmetric cost function, it charges only marginal cost $C_v^O(y_r) - C_{-v}^O(y_r - v)$, to agent v on selected route r (for flow y_r according to our Inter-domain model), where $C_v^O = y_r (\sum_v c_v^T 1_{(v \in y_r \text{ and } y_r > 0)})$ is the total cost for path r and flow y_r , and C_{-v}^O is for the similar cost for the selected path when v is not part of the network, and charges nothing to agents to which the flow is not allocated. By sub-modularity of costs, the payments collected from the users fall short of the actual cost of service. The mechanism cannot be used by a group of agents in autarcy.

Other works than Inter-domain or multi-cast are the works of S. Eidenbez, L. Anderregg, W. Wang, and others in incentive-compatible wireless ad hoc and Internet routing [?, ?, ?, ?].

4.3 Deficiencies of VCG onto Inter-domain problem

Most works on inter-domain pricing have been done using VCG auction [?, ?]. Indeed, VCG is a very powerful mechanism, since it is the only mechanism [?] which gives us a general way of constructing a dominant-strategy, truthful (incentive-compatible), individually rational, and maximizing the social welfare in quasi-linear settings. However, despite of these attractive properties, VCG also brings out some undesirable characteristics when applied to our context, as described below. Indeed, we have already seen that no mechanism can verify the above properties and at the same time budget balance. The deficiencies of VCG are therefore described now.

1. Probably the most important drawback of VCG is the fact the VCG does not verify budget balance. This has been well discussed when VCG is applied to Inter-domain but also to ad hoc networks in [?], and was already known in general and highlighted among other weaknesses by other authors [?, ?, ?, ?]. The issue is that the optimization problem (??), minimizing total routing cost, only deals with intermediate nodes, and not the traffic sender. If we include the sender in the VCG auction, he would pay the sum of declared costs only. But total VCG payment is always more than the sum of declared cost. This can be verified from (??), where the first term is the declared cost, but the remaining value is refereed as subsidies, which are always more than zero because total cost of cheapest path without v is always more than the cheapest path involved by v . Those subsidies incentive the intermediate nodes, acting selfishly, to truthfully reveal their transfer costs and making routing efficient [?]. But as a result,

the balance of the mechanism is always negative, i.e., the money has to be added from an external actor, for instance a governmental one to make the scheme work. This situation is very unlikely. Some procedures have been designed to partially reduce the total value which has to be injected to the model and will be described in next subsection.

As a numerical illustration, consider again the example of Figure ?? . If we look at the sum of payments Table ?? , it is 9, while cheapest path has a declared valuation of 5. The VCG payment for these intermediate nodes is even more than the second (declared) cheapest cost 7.

2. Vulnerability to collusion: we noticed that truth telling is a dominant strategy for each agent, i.e., no one can gain by declaring something else than his true valuation/cost. However, this is not confirmed in case of groups of agents. Indeed, it turns out that groups of colluding agents can achieve a higher utility by coordinating their declarations rather than honestly reporting their valuations [?, ?]. In a decentralized context without the presence of a central authority, the group of agents can achieve higher utility by coordinating their declared cost rather than honestly bidding its valuations. For example, the cost paid by the source provider can be increased to 6 for C and 7 for E if the two winners simultaneously decrease their declared cost by one. As a result, the solution are still the same for C and E but the total cost charged on source provider is increased because the loss of declared welfare the auction winners impose to other provider is increased (because their declared cost is simultaneously decreased).
3. Computational intractability: in combinatorial auctions (as in our Inter-domain model where in general we have to deal with auctions on individual nodes) or in general multi-unit auctions, if the inter-domain deals with non-elastic demands from intermediate providers, determining the winner is computationally hard [?]. The above model from Feigenbaum et al. simplifies the issue by not considering any bandwidth or capacity constraint. In that case, the only complexity issue comes from the distributed price computation from hop-to-hop computation for each AS in the network. For BGP to converge (when all Lowest Cost Path (LCPs) found), it requires d stages of computation where d is the maximum number of AS hops in an LCP and N is the number of ASes in V . Each stage involves $O(Nd)$ communication on any links [?]. If Inter-domain resource allocation problems involve more constraints (bandwidth, budget, etc), the optimization problem can become computationally intractable. The example of a comprehensive optimization model for Inter-domain allocation will be presented lately.

4.4 Improvements for applying VCG to Inter-domain

VCG mechanism has therefore been recently applied in various network applications. Some authors have applied VCG mechanism in order to achieve an dominant-strategy, incentive-compatible and efficient allocation algorithm. Unfortunately, we have already seen thanks to *Myerson theorem* that such an algorithm cannot verify budget balance. In this section, we consider several algorithms and applications which are variations (mostly) based on VCG and

decompose them into two sub-groups, budget-balanced and budget-imbalanced. In each sub-group, we investigate the properties of mechanism design in cases which property they verify or remove when applied to inter-domain resource allocation problems.

4.4.1 The imbudget-balanced approaches

H. Moulin recently investigated the problem of assigning multiple objects to N agents using a modification of VCG mechanism [?], and to improve the imbalance of the budget. The main idea is to distribute a suitable rebate function τ to each participant, in such a way that the sum of the rebate functions almost exhaust the revenue from the auction. Adapting it to our inter-domain problem, source nodes would like to form a route r to its wished destination. Let's denote by ζ as the number of traffic flows (i.e. couple source-destination) that have to be carried, if possible, and remember that is y_r is the flow on route r . The problem now becomes assigning an object y_r^* to ζ agents. The set of flows y_r^* can be chosen as the solution of the optimization problem:

$$\min_y \quad \sum_{v \in N} c_v^{-T} \cdot x_v^T \quad \text{s.t.} \quad y_r^* = x_v^T, \forall v \in r^*$$

where, again, x_v^T is total amount of traffic forwarded by node v and c_v^{-T} its transfer cost per unit of flow. We therefore again try to minimize the total routing cost. The total cost for the sender of traffic flow y_r^* on route r is determined as

$$\rho_{r^*} = v_{y_r^*} - v_{-y_r^*} + \tau_{-y_r^*} \quad (4)$$

where

$$\tau_{-y_r^*} = \sum_{k=2}^{\zeta-1} (-1)^{k-2} \frac{L^*(\zeta, 1)}{kL^*(\zeta, k)} y_{-r^*}^{*k}, \quad L^*(\zeta, k) = \frac{\binom{\zeta-1}{1}}{\sum_{k=1}^{\zeta-1} \binom{\zeta-1}{k}}$$

$v_{y_r^*}$ is the efficiency surplus with a strategy profile y_r^* ² and $v_{-y_r^*}$ is the worst efficiency surplus in the absence of strategy profile y_r^* ³. L^* is the smallest efficiency loss⁴ of the mechanism and $y_{-r^*}^{*k}$ is the flow on the k th cheapest path for a given element of ζ . The main idea is that the total payment is determined on the route r for y_r^* , instead of to each node as in original VCG mechanism or in J. Feigenbaum et al's work. The total payment for a subset of node formed a route r^* based on Moulin afterward is then distributed to each intermediate node/seller. The assignment of one object to ζ agents using VCG become a second sealed bid auction. The advantage of the payment function is that, by using a rebate function τ from Moulin's approaches, total cost imposed on source provider for r^* is always lesser than the second cheapest path. On the

²Given a profile of valuation x , a vector x^* is its permutation where coordinates are arranged decreasingly $x^{*1} \geq x^{*2} \geq \dots x^{*\zeta}$. The efficient surplus given one object and the profile of valuations y_r^* is $v_{y_r^*} = x^{*1}$.

³Similarly, $v_{-y_r^*} = x_{-y_r^*}^{*1}$

⁴The worst efficiency loss is the maximum ratio of the mechanism's budget imbalance to efficiency surplus over all profiles of non-negative valuations. In our case of assigning one object, the efficiency loss is $L^*(\zeta, 1) = \max \frac{\text{out_flow_money}}{\text{efficient_surplus}} = \max_{x \in \mathbb{R}^\zeta} \frac{x^{*1} - x_{-y_r^*}^{*1}}{v_{y_r^*}}$.

example defined Figure ??, the total payment received by intermediate nodes in r^* is $\rho_{r^*} = -5 + \tau_{-y_{r^*}}$ where

$$\tau_{-y_{r^*}} = \frac{13}{62}y_{-r^*}^{*2} - \frac{8}{93}y_{-r^*}^{*3} + \frac{3}{62}y_{-r^*}^{*4} - \frac{1}{31}y_{-r^*}^{*5} \simeq -1.1.$$

The total payment the source provider should pay the intermediate nodes is 6.1 per unit traffic, almost approach the total declared cost from the sellers: the difference between between payment and rewards is now 1.1 instead of 4 as in the *original* VCG of Equation (?). To illustrate more the budget imbalance between VCG and VCG-Moulin payment rules, we make 100.000 different simulations on *Matlab* using the topology described Figure ??, with for each node a transit cost randomly generated from a uniform distribution over $\{2, \dots, 10\}$. The empirical cumulative distribution of budget imbalance is depicted in Figure ?? for the two algorithms. We can see that VGC generates more monetary loss in the market than VCG-Moulin's payment, but on the other hand yields a larger social welfare as depicted in Figure ??. The Moulin's mechanism is claimed to achieve asymptotic efficiency, strategy-proof, preserving IR property (with a single object) while *almost satisfying budget-balanced*, but still not verifying it. It therefore still suffers from this drawback, though at a lesser extent, and at the expense of a longer computation of costs.




Figure 2: Empirical cumulative distribution of the budget imbalance in VCG auction

4.4.2 The budget-balanced approaches

In this subsection, we concentrate on the budget-balanced property. We indeed believe that it is an important one, if not a requirement, towards self-organized inter-domain pricing. We are going to describe several approaches. Each time,




Figure 3: Empirical cumulative distribution of the social welfare in VCG auction

another property (at least) has to be relaxed. Note also that several can be, but can be *almost* verified, which can have some interest if the computational complexity is reduced (for instance).

M. Dramitinos et al. [?] also applied VCG reserve pricing for bandwidth allocation *over links*, relaxed efficiency property and using multi-units Dutch auction for agents to bid the links. Their mechanism can be formulated in our Inter-domain model as follows. Each intermediate node $v \in V$ is characterized by a pair $(x_v^T, \hat{c}_v^{T_0})$ where x_v^T is the maximal capacity node v willing to sell for forwarding external traffic and $\hat{c}_v^{T_0}$ is the initial price per unit. Then, following Dutch auctions principle, the price on intermediate nodes is reduced as time elapses. Source providers place their bids when prices are reasonable to them and are instantly allocated bandwidth over a bundle of links (to form y_r). The scheme continues until price reaches zero which helps to compute payment since first price payment is not considered. Providers with higher bids are therefore allocated first. The game is therefore different from those in previous sections: we rather have a game where players are traffic senders competing for (limited) resources, while sellers, intermediate nodes just sell to the highest bidders (no cost is involved here). Price reduction speed can be different, according to demand and residual capacity, in order to solve the game faster. Senders submit their bids when the prices (which are then different for each link) are such that their utility for the whole path is reached. Efficiency is shown to be verified, but only under full information about market demand, which is unlikely. The payment for intermediate nodes is determined as the opportunity cost it imposed on others and based on their reserve quantities while bidding for the bandwidth, i.e., the bids of users excluded from the allocation. Also routes are awarded exactly that amount of money. Overall, the scheme is not efficient (provided no full knowledge of demand), incentive-compatible thanks

to the second price principle, budget-balanced because senders pay exactly for what they have (which is less than their bids), and scalable with respect to the number and capacities of links and the number of users.

B. Faltings also made some changes in VCG payments by randomly selecting an excluded coalition of an agent or a group of agents whose bids are rejected during the optimization [?]. The excluded coalition can be chosen by any mechanism that does not depend on the utility declarations of the successful agents, and will pay for the (negative) collected surplus. Their work is fine when surplus is positive, but, it is unlikely to ask excluded players (getting no resource) to pay for others and in order to reach budget balance. Consequently, if applied to Inter-domain, the mechanism is neither efficient nor individually rational.

The AGV (Arrow, d'Aspremon and Gerarld-Varet) mechanism [?] also provides some modifications to Vickrey payments. The AGV unit payment is formulated as

$$p_v^r = \sum_w \bar{c}_w \mathbf{1}_{(w \in (y_{-v}^*)_r \text{ and } (y_{-v}^*)_r > 0)} - \frac{1}{N-1} \sum_{w \neq v} \left[\sum_{u \neq w} \bar{c}_u \mathbf{1}_{(u \in (y_{-w}^*)_r \text{ and } (y_{-w}^*)_r > 0)} \right]. \quad (5)$$

For each node v , it calculates the minimal cost when v is excluded,

$$\sum_w \bar{c}_w \mathbf{1}_{(w \in (y_{-v}^*)_r \text{ and } (y_{-v}^*)_r > 0)},$$

and each agent v is charged a $\frac{1}{(N-1)}$ share of the payments made to each of the other agents. It guarantees that the mechanism is budget-balanced but two sacrifices are made in exchange for this property. AGV is truthful only in Bayes-Nash equilibrium rather than dominant strategies and is *only ex ante* individually rational. Scalability is still an issue though. Variations of the redistribution, instead of being uniform among the $N-1$ other agents in (??), can be more subtle [?, ?, ?]. With this type of mechanisms, some agents may not be the winner (in terms of getting traffic) and still have to pay to make the system balanced. It seems unlikely in the case of Inter-domain.

5 Alternative solutions

We propose here alternative mechanisms for inter-domain pricing.

5.1 Combinatorial double-sided auctions

Most auctions applied to network allocation problems are one-sided auctions: either multiple buyers compete for commodities sold by one seller, or multiple sellers compete for the right to sell to one buyer. Combinatorial double-sided auctions are auctions in which both sellers and buyers submit their demand or supply bids. In some applications, combinatorial double-sided auctions are more efficient than several one-sided auctions combined [?]. Inter-domain is actually a natural application with multiple sellers (providers could transit traffic by request) and multiple buyers (providers which want to forward traffic to destination for their customers). Our aim is therefore to apply combinatorial double-sided auctions to Inter-domain pricing. R. Jain et al have applied

combinatorial double-sided auctions to network resource allocation [?, ?]. The mechanism is claimed to be *ex ante* individually rational, efficient and strongly budget-balanced, but truth-telling is dominant strategy for all players except the highest matched sellers over each traffic flow [?]. Therefore, incentive compatibility property is in some sense *almost* verified. The mechanism is defined briefly and more formally now. While in [?, ?] buyers and sellers are separate players, in our case they are both at the same time, which is more complicated as require an adaptation.

Each provider v places buy-bids on a set of routes \mathcal{R}_v (note that routes are here initiated from v , but this can be extended without difficulty). For each route $r \in \mathcal{R}_v$, the bid is made of the maximum per unit price $\hat{c}_v^I(r)$ AS v is willing to pay and the maximum amount \hat{y}_r he is willing to get. He additionally places sell-bids $(\hat{c}_v^T(r), \hat{x}_v(r))$ for routes r AS v is on, where $\hat{c}_v^T(r)$ is the minimum unit price he wants to sell resource and $\hat{x}_v(r)$ the maximum amount he agrees to sell.

The auctioneer then tries to allocate resource on routes, $y = (y_r)_r$ such that

$$\max_y \sum_v \sum_{r \in \mathcal{R}_v} \hat{c}_v^I(r) y_r - \sum_v \hat{c}_v^T(r) x_v^T \quad (6)$$

such that

$$\begin{aligned} 0 &\leq y_r &&\leq \hat{y}_r \quad \forall r \\ x_v(r) &\leq \hat{x}_v(r) &&\quad \forall v \\ x^T &= Ay &&\leq \hat{x}^T \end{aligned}$$

where

$$\begin{aligned} A &= (A_{vr})_{v,r} \text{ with } A_{vr} = 1 \text{ if } v \in r, 0 \text{ otherwise} \\ y &= (y_r)_r \\ x^T &= (x_v^T)_v. \end{aligned}$$

Following [?], the (settlement) price imposed to a buyer is the highest ask price among matched sellers. Buyers' (resp. sellers') bids are larger (resp. smaller) than the settlement price. The payment made by buyers is exactly what sellers get, thus making the mechanism budget-balanced. The mechanism is shown to be efficient, individually rational, strongly budget-balanced but not incentive compatible, though *almost* incentive compatible. The payment rule applied to our example is illustrated Table ???. The scheme is particularly attractive due to all those properties. On the other hand, solving the allocation problem is an NP-hard problem, making it complicated to be applied to large topologies. To make the mechanism decentralized (and scalable), each provider would better interact with its neighbor ASes (according to its routing table). Then it becomes a double-sided auction in case of one buyer and its neighbor sellers. In other words, the above mathematical optimization problem is decomposed into several local optimization problems (relation between an AS and its neighbors), and those local problems have to be match such that the whole allocation for a path is the same on all nodes. In Jain et al's work, the mechanism needs a central authority for the calculation of the surplus maximization problem (??). So, we would like to apply a distributed constraint optimization problem using only a partial information between every pair of nodes to find y_r .

Another option is to apply VCG double-sided auction to preserve incentive compatibility property [?]. Quoting [?], if \tilde{y}^{-v} (respectively \tilde{y}^{-v}) is the allocation vector when $\hat{y}_r = 0 \forall r \in \mathcal{R}_v$ (respectively $\hat{x}_v(r) = 0$ for all route r v is on) the money transfer (the payment) to be made by buyer v is

$$\tilde{C}_v = \sum_{u \neq v} \sum_{r \in \mathcal{R}_u} \hat{c}_u^I(r) (\tilde{y}_r^{-v} - y_r) - \sum_u \sum_{r \in \mathcal{R}_u} \hat{c}_u^T(r) ((\tilde{x}_r^T)^{-v} - x_v^T) \quad (7)$$

and the money transfer to be made by seller v (negative would means transfer to the seller)

$$\bar{C}_v = \sum_u \sum_{r \in \mathcal{R}_u} \hat{c}_u^I(r) (\tilde{y}_r^{-v} - y_r) - \sum_{u \neq v} \sum_{r \in \mathcal{R}_u} \hat{c}_u^T(r) ((\tilde{x}_r^T)^{-v} - x_v^T). \quad (8)$$

Payment can alternatively be made by *route*, by just removing the considered route, instead of the whole set of route of AS v . This mechanism seems to give us another approach for inter-domain pricing but it still suffers from some problems (budget balance being the main one) requiring to be solved in order to be applied to inter-domain.

David Parkes et al. [?] have modified Vickrey payment in order to achieve an *fairly efficient, fairly incentive-compatible* in VCG combinatorial double-sided auction scheme while remaining BB and IR properties. Budget balance is achieved in bilateral trading by redistributing the surplus to all agents involved in the trade. Payment to each node v could be transformed from opportunity cost in (??-??) to Vickrey discount [?] by

$$p_v^r = \hat{c}_v^T(r) 1_{(v \in y_r^* \text{ and } y_r^* > 0)} - \Delta_{VCG}, \quad (9)$$

$$p_v^r = \hat{c}_v^I(r) - \Delta_{VCG}, \forall r \in \mathcal{R}_v \quad (10)$$

where Δ_{VCG} is the *Vickrey double-sided auction discount* and Δ_{VCG}^* is modified Vickrey discount defined in [?] to get mechanism's budget balance. Vickrey discount is the marginal utility in which the agent v contributes to the market. Vickrey discount is always non-negative, representing smaller payments by buyers and higher payments to sellers and is calculated by

$$\begin{aligned} \Delta_{VCG} &= V^* - V_{-v}^* \\ \Delta_{VCG}^* &= \max(0, \Delta_{VCG} - \frac{(\sum_{v \in r^*} \Delta_{VCG} - V^*)}{|r^*|}) \end{aligned} \quad (11)$$

where

$$V^* = \sum_{v \in y^*} \sum_r (\hat{c}_v^I(r) \cdot y_r - \hat{c}_v^T(r) \cdot x_v^T)$$

is the total gain from trade under optimal allocation of goods.

$$V_{-v}^* = \sum_{u \in y_{(-v)}^*} \sum_r (\hat{c}_u^I(r) \cdot \tilde{y}_r^{-v} - \hat{c}_u^T(r) \cdot (\tilde{x}_r^T)^{-v})$$

is the sum of all gain from trade under optimal allocation of goods without agent v . So budget balance holds when $\sum_{\forall r} \Delta_{VCG}(y_r) \leq V^*$. The authors have provided multiple modified Vickrey discounts Δ_{VCG}^* and have compared their

results in experiment. The mechanism is interesting for us because it brings out the fairly good performance in terms of social welfare ($\approx 90\%$ of the optimal one) with low difference between real and declared costs in providing their services according to simulations (i.e., the model is fairly incentive compatible). Table ?? illustrates the different payments in combinatorial double-sided auctions applied to our example defined Figure ?. We assume that the buyer/source provider specifies its willingness to pay $c_{y_r}^I = 7$ cost per unit traffic for the whole path to destination⁵. In Table ??, we notice that the difference

Node	Valuation	$(\Delta_{VCG}, \Delta_{VCG}^*)$	VCG[?]	Modified VCG[?]	Highest price[?]
Src	7	$(2, \frac{1}{3})$	5	$7 - \frac{1}{3}$	5
A	-5	(0, 0)	0	0	0
B	-4	(0, 0)	0	0	0
C	-2	$(2, \frac{1}{3})$	-4	$-(2 + \frac{1}{3})$	-2
D	-4	(0, 0)	0	0	0
E	-3	$(3, \frac{4}{3})$	-6	$-(3 + \frac{4}{3})$	-3
F	-6	(0, 0)	0	0	0

Table 3: Costs and payments to nodes for the example of Figure ?.

between declared valuation and rewards in VCG double-sided auction is -5 , negative payment means the agent receiving the money from the trade. In that case, the mechanism has to inject money to make it work. On contrary, if we apply modified VCG double-sided auction by D. Parkes, the surplus has been redistributed only to agents involved in the trade to make it budget-balanced. The highest ask-price rules by R. Jain in combinatorial auction also lead to the budget-balanced payments. We also run simulations for double-sided auctions using Matlab on the topology described in Figure ?, with again node transit costs uniformly and independently distributed over $\{2, \dots, 10\}$. Notice first that double-sided VCG auction generates the same values of budget imbalance than in the single-sided case, whatever the bidding values of the buyers. Furthermore, although modified VCG payment rules by D. Parkes *et al.* induce different payments than with the highest matched price proposed R. Jain *et al.*, both pricing schemes generate exactly the same social welfare. An empirical cumulative distribution of the social welfare for VCG and the VCG mechanism modified by D. Parkes *et al.* is depicted in Figure ? for the different payment rules in double-sided auction.

The difference between the pricing rules in Jain's works defined in [?, ?] is an illustration of the trade-off in mechanism design following the Myerson impossibility theory. In [?], the resulting allocation is efficient, *ex ante* individual rational, budget-balanced but not incentive-compatible to the highest asked sellers. So, it could possibly lead to a price war between sellers if they both try decrease their prices to win the game rather than truth-telling to their services. Furthermore, the pricing scheme proposed by David Parkes et al. [?] achieves budget balance in the market but the pricing scheme is only *fairly efficient*, *fairly incentive-compatible* in VCG combinatorial double-sided auction while remaining BB and IR properties. On the other hand, while using VCG double-sided auction in [?] the resulting allocation is efficient, individual ratio-

⁵The price is approximated from the average path cost per unit traffic from all sellers




Figure 4: Empirical cumulative distribution of social welfare with double-sided auctions for 50 independent experiments

nal, incentive-compatible to all agents but the scheme does not hold the budget balance property anymore.

5.2 Incentive and learning mechanisms in distributed computing

In this section, we present some incentive mechanisms⁶ that provide incentive for collaboration in distributed systems such as inter-domain, ad hoc and wireless networks, etc. When the nodes do cooperate, they make communication in distributed systems viable. It could be the operation of ad hoc network, in which the infrastructure needs the cooperation of wireless nodes to forward traffic to destination. Inter-domain is similar since neighboring nodes should control how traffic enters or leave the network based on inter-domain traffic volumes, and routes. Pricing mechanism have been applied in rate control in wired network and resource control in wireless network to provide a means to stimulate cooperation within the network [?, ?, ?, ?, ?]. To simplify those approaches, we present a very general incentive algorithm based on pricing mechanisms.

In our general algorithm, we can see that the strategies and pricing update functions used is the main ideas and the difference of these approaches. Those strategies are designed to create the proper incentives for cooperation. The difference comes from the pricing models/ game strategies applied to the Inter-

⁶Incentive in cooperation protocols has a different meaning from *incentive-compatible* property in mechanism design. It means providing incentive for selfish nodes to follow exactly protocols by (usually) giving rewards to them. Incentive-compatible in mechanism design is the property to describe how the agents truthfully revealed their services in cooperation.

Algorithm 1 A general incentive algorithm based on pricing

- 1: Pricing initialization at each node;
 - 2: **repeat**
 - 3: Nodes: Update(*strategies*);
 - 4: Nodes: Update(*price*);
 - 5: **until** Nash equilibrium
-

net. Let us begin with recent incentive mechanisms in which we focus on their strategies and pricing function and try to adapt those to inter-domain.

Relay selection and resource control in wireless networks is quite similar to the cooperativity problem of buyer/seller in inter-domain pricing. B. Wang et al. has applied Stackelberg game [?] to stimulate cooperation between relay nodes. The Stackelberg leadership model [?] is a strategic game in economics in which, the leader (the source-buyer provider) moves first and then the follower (the intermediate-seller providers) move after. The motivation of the authors is to help the source/buyer to smartly choose relay/seller nodes and competing relay nodes to ask a reasonable price to maximize their utilities. A Stackelberg game is solved by *backward induction* in order to find sub-game perfect Nash equilibria. The buyer/ leader picks a quantity that maximizes its payoff, anticipating the predicted response of the sellers/ followers. The followers actually observe this and in equilibrium pick the expected quantity as a response. To place it into our context, the optimization problem for the source provider/buyer *source* can be formulated as:

$$\max U_{source} = c_{source}^I \cdot \sum_v x_v^T - \sum_v c_v^T \cdot x_v^T$$

while as a seller the goal is

$$\max U_{seller} = c_v^T \cdot x_v^T \quad (12)$$

where c_v^T and c_{source}^I are the declared costs. For the pricing update function used, the seller nodes first propose a low asking pricing at the beginning and then, increase the price according to optimal bandwidth purchase from the source. In this kind of Stackelberg competition model, buyers and sellers are considered different players. The resulting incentives (by giving rewards and maximizing player's utility function) yields an efficient, individually rational, budget-balanced scheme because the sellers are paid exactly what they provide for their services. On the other hand, the mechanism is not incentive-compatible.

A similar step-by-step approach has been applied to inter-domain by D. Barth et al. [?, ?]. First, they propose a transit price negotiation model for each provider, then a learning algorithm is used by each provider to update node's strategies. The transit price negotiation of nodes depends on the traffic flow entering and leaving ASes, $x_v^I = \sum_{r \in \mathcal{R}^S(v)} y_r$ and $x_v^T = \sum_{r \in \mathcal{R} \setminus (\mathcal{R}^S(v) \cup \mathcal{R}^D(v))} y_r$.

If the operator has no transit traffic, it reduces its price to try to attract some. If the operator overuses its capacity regardless from serving transit traffic, in this case it does not update its price. The pricing strategies are based on the

two parameters (μ, λ) where μ is a ratio between 0 and 1 and λ is an integer.

Moreover, a distributed learning algorithm [?] is applied and analyzed as a repeated game where providers update their strategies at each stage in order to learn their optimal prices. Each node's strategy profile is defined as a probability vector $x_v^t = (x_{v,0}, \dots, x_{v,N}) \in \mathcal{X}_v^t$, where $x_{v,j}^t$ is the probability node v choose the price $c_{v,j}^t \in \mathcal{C}_v^t$ (discrete set) at stage t . Similarly to [?, ?, ?], the algorithm is driven by the following steps:

1. at step 0, an initial strategy profile $x_{v,0}^t$ is chosen. One can for instance a uniform distribution of no information is available.
2. At each step k , the transit price $c_{v,k}^t$ by node v is chosen according to probability $x_{v,k}^t$.
3. paths with smallest costs are computed based on prices announced by neighbors and own declared costs.
4. The outcome of the game is computed.
5. Each node computes its utility function $u_{v,k}$ resulting from the outcome of the game at step k .
6. Each node updates its strategy profile strategy according to

$$x_{v,j}^{k+1} = \begin{cases} x_{v,j}^k - bu_{v,k}x_{v,j}^k & \text{if } c_{v,j}^t \neq c_{v,k}^t \\ x_{v,j}^k + bu_{v,k} \sum_{\forall l, c_{v,l}^t \neq c_{v,k}^t} x_{v,l}^k & \text{otherwise.} \end{cases} \quad (13)$$

This step adjusts the probability distribution for the transit price, considering the utility brought from currently adopted strategy. If the utility is high then the probability of the current strategy is increased, otherwise it is decreased.

7. If the algorithm has not converged to an equilibrium point, go to step 2., otherwise stop.

In the algorithm, $b \in [0, 1]$ is the step-size of the updating rule and $u_{v,k}$ is the normalized utility $u_{v,k} = \frac{U_{v,k} - A_{v,t}}{B_{v,t} - A_{v,t}}$, with $A_{v,t} = \min_{k \leq t} U_{v,k}$ and $B_{v,t} = \max_{v,t} U_{v,k}$. From Theorem 3.2 in [?] and Theorem 2 in [?], the algorithm always converges to a point which is a Nash equilibrium of the game. This game model does verify very few of the important properties. It is not necessarily incentive-compatible, nor efficient. But on the other hand, it is budget-balanced, individually rational, computationally tractable, distributed, which are very relevant characteristics for inter-domain.

J. Crowcroft et al. [?, ?] also explore the selfish problem in ad hoc networks using pricing a mechanism in which relay nodes are rewarded for their power consumption and traffic transfers. Each source node s may choose an amount to pay per unit time, ω_s denoted as its willingness-to-pay parameter, and receives in return a traffic flow y_r , on route r , proportional to ω_s . The price rate, as a function of time, node v charges for forwarding a unit traffic flow along route r

is denoted by $c_v^T(t)$. Then, the total flow rate generated by source s is given by expression

$$\sum_{r \in \mathcal{R}^s} y_r(t) = \frac{\omega_s(t)}{\min_{r \in \mathcal{R}^s} \sum_{v \in r} c_v^T(t)}. \quad (14)$$

This model selects the lowest cost path as usually in communication networks. Pricing strategies are dynamically updated according to each transit traffic flow y_r and maximum supply \mathcal{X}_v from intermediate/ relay nodes.

$$\frac{d}{dt} c_v^T(t) = \frac{\kappa c_v^T(t)}{\mathcal{X}_v} (x_v^T(t) - \mathcal{X}_v). \quad (15)$$

The dependence of the right-hand side of (15) on price per unit of traffic and bandwidth capacity is an attempt to scale the dynamics of the prices in a network with different prices and capacities. Incentives for collaboration are achieved by using *credit balance*. Credit balance is a virtual money and is provided to each node when they first joint to the network. Each node gains credits by forwarding traffic for other nodes and using it for payments to other nodes to ensure end-to-end delivery of its won traffic. By doing so, incentives to forward are provided. Each node s controls its credit balance, $b_s(t)$ by dynamically adjusting its willingness-to-pay parameter according to the level of its credit balance by using a formula: $w_s(t) = \alpha_s b_s(t)$ for some parameter $\alpha_s > 0$ [?, ?]. In conclusion, the authors claim that incentives to cooperate are the natural outcome of the game that emerges from the pricing mechanism. The mechanism is not incentive-compatible, but efficient (in term of optimal resources consumed at transit nodes and routes selected, decentralized, budget-balanced and individual rational. Global stability of the price and credit balance of the system can be established if the networks structure is static and have been verified through simulation.

A related work on optimization of Inter-domain traffic with multiple ISPs has been conducted by Shakkotai and Srikant [?]. Rather than providing a new pricing model, the authors examine the existence of equilibrium price strategies in an Internet model in which Internet is separated by local and transit ISPs. These ISPs are co-located in a small geographical region and compete for the same customers and transit ISPs are responsible to transfer traffic between local ISPs. Interactions between ISPs are investigated by using repeated games. Pricing strategies between local ISPs are examined by using Bertrand competition which is a model of price competition between duopoly firms which results in each charging the price that would be charged under perfect competition [?]. Interactions between local and transit ISPs are investigated by using Stackelberg competition in which each ISP plays knowing that the next ISP would optimize his play based what he does currently [?]. ISPs play in a predefined sequence to study how they act as price transfer agents between economic levels.

A close but different setting is the optimization of multiple traffic flows between a pair of providers is investigated by G. Shrimali2007 [?]. In this game model, two providers with respective utilities U_i and U_j use an iterative procedure to optimize traffic distribution among different direct paths between them, instead of multi-hop routes involving independent ISPs (in our model, if there are several routes among two nodes, those routes are merged). ISPs utility functions measure some performance indices such as average delay or maximum load on a link. Using the axiomatic theory of cooperative games,

a fair and Pareto efficient outcome usually referred as Nash solution can be obtained by maximizing the Nash product given by $U_i U_j$. The authors then show that global optimization can be separated and solved into separable optimization problem using shadow prices [?] and *Lagrangian multipliers* [?]. Under this approach, ISPs myopically optimize local objectives while iteratively playing best response to each other. The scheme is efficient, fair⁷ but relaxes the IC property because Nash bargaining is not incentive compatible.

6 Discussion and conclusions

6.1 Summary

The goal of this paper was to describe inter-domain relationships between independent and selfish providers, and review the pricing mechanisms that have been proposed to enforce cooperation. The (original) orientation of our paper is to look at the properties verified by mechanisms. The relevant properties when dealing with inter-domain are

1. Individual rationality
2. incentive compatibility
3. efficiency
4. decentralization
5. budget balance
6. robustness to collusion.

In the literature, most works were looking at the promising VCG auctions because it verifies three properties: incentive-compatible, individually rational and efficient. On the other hand, we have shown that in that case, it is impossible to verify at the same time budget balance: someone needs to inject money in the system permanently. The question we wanted therefore to answer was: what are other mechanisms, and their associated properties? We have reviewed the most relevant mechanisms. Some of them were applied to other kinds of problems, and we have adapted to our inter-domain context, described at the beginning in a general form. We have also remarked that if some properties are not fully verified, there exist some “almost” versions for which the property is close to be verified. We have looked at some schemes in this way.

6.2 Discussion about most relevant schemes and research trends

Now that the basic results have been described, we can discuss the most relevant schemes and properties that should be verified, and the next steps to be investigated towards better self-managed inter-domain relationships.

⁷fair means that cooperation should reach equitable performance gains to all agents when compared to their default strategies

6.2.1 Discussion on properties

First, we would like to stress that budget-balance is a requirement. Indeed, it seems important to us that the system can live in autarchy. Even if some improvement of VCG have been proposed, with almost budget balance, this will probably not be acceptable from regulators. This circumvents from using directly VCG and some of its variations (but some could be applied).

Individual rationality is better to be verified too, ensuring that participating to the game brings non-negative utility. On the other hand, of the expected value of the utility is necessarily positive (depending on the random unknowns on the system), this *could* be sufficient.

Incentive compatibility makes sure that players do not cheat when declaring resources, costs or prices. This is expected, but could somehow be relaxed; an “almost version” would be acceptable.

Efficiency ensures that the outcome is the best for the whole system. This is again something expected for the whole system, and important property. Here again, an “almost” version, close to the optimal one, would be acceptable, if the gain in terms of other properties is large.

Decentralization, or scalability, is an important issue. Indeed, the relevant situation looks at thousands of ISPs if considering a network such as the Internet. Computing quickly the outcome of the game is an importance factor. This is also one of the main properties. Even if quantifying the maximum amount of computations is difficult, an NP-hard procedure cannot be considered.

Collusion of nodes has been widely investigated in distributed environment such as grid computing, P2P, wireless and ad hoc networks[?, ?]. Such coalitions can improve the social welfare of the collude nodes, but they can also harm the welfare of nodes outside the coalition [?]. When dealing with telecommunication networks pricing, robustness to collusion is better being verified, although it is often difficult to check in practice. Even VCG could lead to an efficient outcome, a group of collude nodes could increase their utilities while decreasing others. Thus, designing a mechanism which is resilient to collusion could be an interesting property [?, ?], if the resulting allocation is fairer. Otherwise, why not being satisfied with collusion if it can lead (or help to lead) to an efficient outcome?

6.2.2 What about previously described schemes?

Among the previously described schemes, none is totally satisfactory, since an interesting feature comes at the expense of loosing on another property. Though, the most relevant mechanisms, in their current form, are probably (in each case budget balance is verified):

- The double-sided auction of Jain et al. in [?] is efficient, individually rational, budget-balanced but sacrifices incentive-compatible property when highest-ask price rule is used. In order to have IC property, VCG double-sided auction can applied [?], but at the expense of budget balance. Highest-ask pricing rule is preferred in inter-domain due to the (we believe) more important budget balance property. Those double-sided auctions can though be only applied to topologies of small to moderate size, due to the required amount of computations. They are therefore advised in that case.

- In a different direction, the work of Barth et al. is easy to implement, computational feasible. As a trade-off, many properties are not verified, but tractability is, which is important for large network topologies.

6.2.3 Hints on future research activities

In order to drive to a better scheme, here are the research directions we aim at pursuing:

- Investigate distributed versions of the most important above schemes (double-sided auctions...). This might mean losing in terms of efficiency, but if the loss is limited, this would be valuable. Most auction mechanisms listed in our are solved in a centralized way. A research direction would be to look at algorithms to design a decentralized inter-domain pricing, with maybe the help of schemes such as DPOP [?] and ADOPT [?]. These algorithms can deal with the allocation problem as long as the influence of each agent on the solution is limited to a number of variables. For instance, following the same principle, Parkes et al. have proposed some principles for the distributed, faithful implementation of VCG mechanism [?, ?]. It is also robust to collusion because the algorithm reveals limited information to the mechanism. Thus the collusion of manipulative players could be prevented. Dealing with the case where nodes only have contacts with neighbors is worth being investigated.
- The double-sided auctions should be better studied for our situation where providers are sellers and buyers at the same time, especially when ASes have relations with their neighbors only and need to buy bandwidth for transferring the traffic previous node on a route.
- In order to have budget balance, other important properties have to be relaxed. If efficiency is relaxed, we need to analyze in the detail the efficiency loss (in term of social welfare).
- The economic views of ISPs in Internet have been investigated by Shakkotai et al. using repeated games in several competition models. But only flat-based pricing is used to analyze in the model, so making use of other pricing mechanisms could lead to new results.

Acknowledgement

This work was done in the framework of the INRIA and Alcatel-Lucent Bell Labs Joint Research Lab on Self Organized Networks.

References

- [1] ALTMAN, E., BOULOGNE, T., EL-AZOUZI, R., JIMÉNEZ, T., AND WYNTER, L. A survey on networking games in telecommunications. *Computers and Operations Research* 33, 2 (2006), 286–311.
- [2] ANDEREGG, L., AND EIDENBENZ, S. Ad hoc-VCG: a truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents.

- In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking* (New York, NY, USA, 2003), ACM, pp. 245–259.
- [3] AUSUBEL, L. M., AND MILGROM, P. *The lovely and lonely Vickrey Auction*. The MIT Press, 2006.
- [4] BARTH, D., BLIN, L., ECHABBI, L., AND VIAL, S. Distributed cost management in a selfish multi-operators BGP network. In *In Proceeding of Next Generation Internet Networks, 2005* (April 2005), pp. 24–30.
- [5] BARTH, D., ECHABBI, L., AND HAMLAOUI, C. Optimal Transit Price negotiation: The distributed learning perspective. *Journal of Universal Computer Science* 14, 5 (2008), 745–765.
- [6] CAESAR, M., AND REXFORD, J. BGP routing policies in ISP networks. *Network, IEEE* 19, 6 (Nov.-Dec. 2005), 5–11.
- [7] CAVALLO, R. Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems* (New York, NY, USA, 2006), ACM, pp. 882–889.
- [8] CLARKE, E. H. Multipart pricing of public goods. *Public Choice* 11, 1 (September 1971).
- [9] COSTAS COURCOUBETIS, R. W. *Pricing Communication Networks: Economics, Technology and Modelling*. Wiley, 2003.
- [10] CROWCROFT, J., GIBBENS, R., KELLY, F., AND ÖSTRING, S. Modelling incentives for collaboration in mobile ad hoc networks. *Performance Evaluation* 57, 4 (2004), 427–439.
- [11] CROWCROFT, J., GIBBENS, R., KELLY, F., AND ÖSTRING, S. Providing incentives in providerless networks. *Ad Hoc Networks* 2, 3 (2004), 283 – 289.
- [12] D’ASPREMONT, C., AND GRARD-VARET, L.-A. Incentives and incomplete information. *Journal of Public Economics* 11, 1 (1979), 25 – 45.
- [13] DRAMITINOS, M., STAMOULIS, G. D., AND COURCOUBETIS, C. An auction mechanism for allocating the bandwidth of networks to their users. *Computer Networks* 51, 18 (2007), 4979 – 4996.
- [14] EGER, K., AND KILLAT, U. Fair resource allocation in peer-to-peer networks (extended version). *Computer Communications* 30, 16 (2007), 3046–3054.
- [15] EIDENBENZ, S., RESTA, G., AND SANTI, P. COMMIT: a sender-centric truthful and energy-efficient routing protocol for ad hoc networks with selfish nodes. In *Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International* (2005), p. 10 pp.

-
- [16] FALTINGS, B. A Budget-balanced, Incentive-compatible Scheme for Social Choice. In *Agent-Mediated Electronic Commerce VI* (2005), P. Faratin and J. A. Rodriguez-Aguilar, Eds., vol. 3435, Springer, pp. 30–43.
- [17] FEAMSTER, N., BORKENHAGEN, J., AND REXFORD, J. Guidelines for interdomain traffic engineering. *ACM SIGCOMM Computer Communication Review 33* (2003), 2003.
- [18] FEIGENBAUM, J., PAPADIMITRIOU, C., SAMI, R., AND SHENKER, S. A BGP-based mechanism for lowest-cost routing. *Distributed Computing 18*, 1 (2005), 61–72.
- [19] FEIGENBAUM, J., RAMACHANDRAN, V., AND SCHAPIRA, M. Incentive-compatible Interdomain routing. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce* (New York, NY, USA, 2006), ACM, pp. 130–139.
- [20] FLAVIO M. MENEZES, P. K. M. *Internet Routing Architectures*. Sam Halabi, September 2000.
- [21] GIBBENS, R. J., AND KELLY, F. P. Resource pricing and the evolution of congestion control. *Automatica 35* (1998), 1969–1985.
- [22] GREEN, J., AND LAFFONT, J.-J. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica 45*, 2 (March 1977), 427–38.
- [23] GROVES, T. Incentives in teams. *Econometrica 41*, 4 (July 1973), 617–631.
- [24] GUO, M., AND CONITZER, V. Worst-case optimal redistribution of VCG payments. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce* (New York, NY, USA, 2007), ACM, pp. 30–39.
- [25] GUO, M., AND CONITZER, V. Optimal-in-expectation redistribution mechanisms. In *AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems* (Richland, SC, 2008), pp. 1047–1054.
- [26] HAYRAPETYAN, A., TARDOS, E., AND WEXLER, T. The effect of collusion in congestion games. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing* (New York, NY, USA, 2006), ACM, pp. 89–98.
- [27] HERSHBERGER, J., AND SURI, S. Vickrey pricing in network routing: Fast payment computation. In *In Proc. of the 42nd IEEE Symposium on Foundations of Computer Science* (2001), pp. 252–259.
- [28] HURWICZ, L. The design of mechanisms for resource allocation. *American Economic Review 63*, 2 (May 1973), 1–30.
- [29] HURWICZ, L. On the existence of allocation systems whose manipulative Nash equilibria are pareto-optimal. In *3rd World Congress of the Econometric Society* (1975).

-
- [30] JAIN, R. Designing a strategic bipartite matching market. *46th IEEE Conference on Decision and Control, ICDC 2007* (Dec. 2007), 139–144.
- [31] JAIN, R., AND VARAIYA, P. Efficient market mechanisms for network resource allocation. *4th IEEE Conference and European Control Conference on Decision and Control, CDC-ECC '05* (Dec. 2005), 1056–1061.
- [32] JAIN, R., AND WALRAND, J. An efficient Nash-implementation mechanism for network resource allocation. *submitted* (2009).
- [33] KALAGNANAM, J., AND PARKES, D. C. Auctions, bidding and exchange design. In *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, D. Simchi-Levi, S. D. Wu, and M. Shen, Eds., Int. Series in Operations Research and Management Science. Kluwer, 2004, ch. 5.
- [34] LIAN, Q., ZHANG, Z., YANG, M., ZHAO, B. Y., DAI, Y., AND LI, X. An empirical study of collusion behavior in the maze P2P file-sharing system. *Distributed Computing Systems, International Conference on O* (2007), 56.
- [35] LIAO, R. R.-F., WOUHAYBI, R. H., AND CAMPBELL, A. T. Incentive engineering in Wireless LAN based access networks. In *Proceeding of 10th International Conference on Network Protocols (ICNP)* (2002).
- [36] MAILLE, P., AND TUFFIN, B. Why VCG auctions can hardly be applied to the pricing of inter-domain and ad hoc networks. *3rd EuroNGI Conference on Next Generation Internet Networks* (May 2007), 36–39.
- [37] MARBACH, P., AND BERRY, R. Downlink resource allocation and pricing for wireless networks. In *Proceeding of Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM 2002* (2002), vol. 3, pp. 1470–1479.
- [38] MATSUBARA, S. Adaptive pricing that can withstand buyer collusion of false-type-declaration. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce* (New York, NY, USA, 2001), ACM, pp. 249–252.
- [39] MODI, P. J., MIN SHEN, W., TAMBE, M., AND YOKOO, M. ADOPT: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence 161* (2005), 149–180.
- [40] MOULIN, H. An efficient and almost budget balanced cost sharing method. *Games and Economic Behavior In Press, Corrected Proof* (2008).
- [41] MOULIN, H. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory* (2009).
- [42] MOULIN, H., AND SHENKER, S. Strategyproof sharing of submodular access costs: Budget balance versus efficiency. Working Papers 96-31, Duke University, Department of Economics, 1996.
- [43] MYERSON, R. B. Optimal auction design. Discussion papers, Northwestern University, Center for Mathematical Studies in Economics and Management Science, Dec 1978.

-
- [44] MYERSON, R. B., AND SATTERTHWAITTE, M. A. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 2 (April 1983), 265–281.
- [45] NISAM, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. V. *Algorithmic Game Theory*. Cambridge University Press, September 2007.
- [46] NISAN, N., RONEN, A., AND RONEN, A. Algorithmic mechanism design. *Games and Economic Behavior* (2001).
- [47] OSBORNE, M. J., AND RUBINSTEIN, A. *A Course in Game Theory*. The MIT Press, July 1994.
- [48] PARKES, D. C. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *In Proceedings of the 17th International Joint Conference on Artificial Intelligence* (2001), pp. 1161–1168.
- [49] PARKES, D. C. Distributed implementations of Vickrey-Clarke-Groves mechanisms. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems* (2004), pp. 261–268.
- [50] PETCU, A., AND FALTINGS, B. DPOP: A Scalable Method for Multiagent Constraint Optimization. In *Proceeding International Joint Conferences on Artificial Intelligence 2005* (2005), pp. 266–271.
- [51] PETCU, A., FALTINGS, B., AND PARKES, D. C. MDPOP: Faithful distributed implementation of efficient social choice problems. *Journal of Artificial Intelligence Research* 32 (2008).
- [52] REKHTER, Y., AND LI, T. A Border Gateway Protocol. *Internet Engineering Task Force (IETF)*, 1771.
- [53] SANDHOLM, T. Issues in computational Vickrey auctions. *International Journal on Electronic Commerce* 4, 3 (2000), 107–129.
- [54] SASTRY, P., PHANSALKAR, V., AND THATHACHAR, M. Decentralized learning of Nash equilibria in multi-person stochastic games with incomplete information. *IEEE Transactions on Systems, Man and Cybernetics* 24, 5 (May 1994), 769–777.
- [55] SHAKKOTTAI, S., AND SRIKANT, R. Economics of network pricing with multiple ISPs. *IEEE/ACM Transactions on Networking* 14 (December 2006).
- [56] SHAPLEY, L. S. A Value for n-person games. *Annals of Mathematical Studies* 28 (1953), 307–317.
- [57] SHOHAM, Y., AND LEYTON-BROWN, K. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, 2008.
- [58] SHRIMALI, G., AKELLA, A., AND MUTAPCIC, A. Cooperative Inter-domain traffic engineering using Nash bargaining and decomposition. *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE* (May 2007), 330–338.

-
- [59] STAAB, E., AND ENGEL, T. Collusion detection for grid computing. *Cluster Computing and the Grid, IEEE International Symposium on 0* (2009), 412–419.
- [60] TUFFIN, B., AND MAILL, P. How many parallel TCP sessions to open: A pricing perspective. 2 – 12.
- [61] VICKREY, W. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1 (1961), 8–37.
- [62] VRAGOV, R. Implicit consumer collusion in auctions on the Internet. *Hawaii International Conference on System Sciences* 7 (2005), 174c.
- [63] WANG, W., YANG LI, X., AND WANG, Y. Truthful multicast routing in selfish wireless networks. In *In Proceedings of MobiCom* (2004), ACM Press, pp. 245–259.
- [64] WIKIPEDIA. Bertrand competition.
- [65] WIKIPEDIA. Depth-first search.
- [66] WIKIPEDIA. Lagrange multipliers.
- [67] WIKIPEDIA. Shadow prices.
- [68] WIKIPEDIA. Stackelberg competition.
- [69] XIA, M., STALLAERT, J., AND WHINSTON, A. B. Solving the combinatorial double auction problem. *European Journal of Operational Research* 164, 1 (2005), 239 – 251.
- [70] XING, Y., AND CHANDRAMOULI, R. Stochastic learning solution for distributed discrete power control game in wireless data networks. *Networking, IEEE/ACM Transactions on* 16, 4 (Aug. 2008), 932–944.



Centre de recherche INRIA Rennes – Bretagne Atlantique
IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399