



UNIVERSITY CARLOS III OF MADRID

Department of Telematics Engineering

Master of Science Thesis

Game Theory Application to Interdomain Routing

Author: **Andra Lutu**
Dipl.Eng. in Telecommunications

Supervisor: **Marcelo Bagnulo Braun, Ph.D.**

Leganés, September 2010

Abstract

This thesis proposes a game theoretic analysis of interdomain routing. In the two following chapters we try to capture some of the intricacies of the current Internet routing protocol, the *Border Gateway Protocol* (BGP).

The first chapter of the thesis paper presents a survey of recent advances in the application of game theory to model the behaviour of interdomain routing. In these models, the participants of the interdomain routing game are represented as strategic agents seeking to improve their benefits through the manipulation of the interdomain routing protocol BGP. The main results achieved over the last few years in this field include models to analyze the stability of the interdomain routing and the design of mechanisms that guarantee BGP to be incentive-compatible with or without monetary transfers.

However, over the last years, the research community has been deeply concerned about the scalability issues that the Internet routing is facing. As the Internet popularity grows, so do the network resources needed in order to sustain its worldwide availability. In the second chapter of this thesis we consider a commons model in which the Global Routing Table (GRT) is a public resource. We use this model to study the economic incentives the ASes have for deaggregating their assigned address blocks. We evaluate the efficiency of the global routing system, the properties of the game equilibria and we examine its relation to the social welfare point of the considered game setup. We find that the strategy adopted by the ASes in the interdomain is not an overall optimum strategy and it leads to an inefficient exploitation of the common resource. Therefore, we prove that the GRT, just like any common natural resource, “remorselessly generates tragedy”, following Hardin’s game theoretic analysis on the *tragedy of the commons*. Finally, we introduce in the model a pricing mechanism that aims to avoid the tragedy of the Internet routing commons.

Table of Contents

1	Advances in Game-Theoretic Models for Interdomain Routing	1
1.1	Introduction	1
1.2	A Game Theoretic Approach to BGP Safety	2
1.3	Incentives for BGP-compliant behaviour	5
1.3.1	Incentive-compatibility without traffic attraction	5
1.3.2	Incentive compatibility with traffic attraction	7
1.4	Payment Mechanisms for Incentivizing ASes	8
1.5	Conclusions	9
2	The Tragedy of the Internet Routing Commons	10
2.1	Introduction	10
2.2	The Game Theoretic Model	11
2.3	Game Equilibria and Overall Optimality	13
2.3.1	Game Equilibria	13
2.3.2	Social Optimality	15
2.3.3	Evaluation of the Model	16
2.4	A Payment Mechanism	17
2.5	Conclusions	19
	References	20
	Appendix	21

List of Figures

1.1	Examples of BGP system with more than one or no stable states	3
1.2	BGP is not incentive-compatible even if Gao-Rexford constraints hold. . . .	6
1.3	Route Verification or Policy Consistency alone is not enough in a network with traffic attractions.	7
2.1	The generic shape of the value function	13
2.2	The general shape of the value function	22

Chapter 1

Advances in Game-Theoretic Models for Interdomain Routing

1.1 Introduction

The Internet is the interconnection of a large number of independently-managed domains called Autonomous Systems (ASes). ASes are driven by economic forces and they obtain benefits from routing users' traffic. The process of establishing routes between ASes is called *interdomain routing*. This is a distributed task that involves all the ASes and it is a very complex one, due to the restrictions imposed by the business relationships between ASes. ASes exchange routing information and express their preferences using the Border Gateway Protocol (BGP).

Game Theory provides a collection of analytical and modeling tools that helps understanding the interaction between decision-makers. Consequently, it can be used to gain a deeper understanding of the ASes behaviour. In this paper we present a survey of game-theoretic approaches to analyze interdomain routing.

The first aspect of interdomain routing that has been modeled using game theory is the stability of the system. BGP allows for the configuration of private preferences and complex routing policies. Ideally, a BGP system would use only *stable routes* to forward traffic, thus reaching a *stable state* of the network, a state in which ASes would not change their selected paths [9]. However, it is possible that in some cases the routing policies of different ASes interact in a dynamic way, creating unwanted route oscillations. A BGP system is said to be *safe* when it has a unique stable state and that state is guaranteed to be reached by the network. Assuring BGP safety is a challenging task and it has been a long time desired feature of any BGP system [10, 6]. As we describe in the section 1.2, game theory can be used to model the BGP and provide some insight with respect to the safety of the system.

Another area where game theory has made significant contributions is related to the incentive-compatibility in the interdomain routing. By modeling ASes as *rational agents*, whose sole purpose is to increase their benefit even with the cost of harming the other entities in the network, game theory proves that it may be in the self interest of the ASes to misbehave. Misbehaving in the interdomain takes numerous forms, such as lying, reporting inconsistent information or denying routes. In section 1.3, we present some BGP protocol and operational changes that remove the incentives for rational ASes to misbehave. An

alternative approach for guaranteeing incentive-compatibility in the interdomain routing is presented in section 1.4, where payment mechanisms are introduced.

1.2 A Game Theoretic Approach to BGP Safety

The BGP route selection process in a network is a distributed and asynchronous operation, triggered by advertisements and withdrawals of routes. The state of the network is determined by the routes chosen by all the ASes involved in the process. We say that the BGP system *converges* when it reaches a stable state in which ASes would not change their routing choices. However, a BGP system may not converge to a stable state, even if one actually does exist for the network [9]. There are then two important features in a BGP system: *solvability* (the BGP system has at least one stable state) and *safety* (the BGP system has a stable state and converges to it in all possible scenarios).

Game theory can be used to model BGP and provide a different perspective on the stability and safety aspects of the interdomain routing. The game theoretic model for interdomain routing proposed in [16] defines two games: the ONE-ROUND GAME and the CONVERGENCE GAME.

The ONE-ROUND GAME is a static game with complete information¹. The network is modeled as a connected graph $G = (N, L)$, where N is the set of nodes and L is the set of links. The ASes (the nodes of the graph) represent the players in the game and they are considered as *strategic agents*, who try to improve their outcome. The actions that are available to each player is the selection of a route among the different ones available towards the destination. The payoff functions of the players are the preferences of the different ASes over the different routes they can use, meaning that the actual payoff for a player is be equal to the preference of the actual route that it has finally selected. Since the ONE-ROUND GAME is a static game with complete information, it basically assumes that all ASes are aware of the preferences (i.e. the routing policies) of every other ASes and that they do not consider the history of the announcements.

Each of the players has to chose a *strategy* that specifies a feasible action for the player in every eventuality in which the player might be called upon to act. The combination of the strategies for all the payers in the game form a “*strategy profile*” that determines the outcome of the game (including the payoff for each player). In the ONE-ROUND GAME the strategy of a player is the procedure to chose an outgoing edge from all the available permitted ones for forwarding traffic. In this sense, executing the BGP selection process is a strategy.

A pure Nash equilibrium in the ONE-ROUND GAME corresponds to a “stable state” [9]. It corresponds to a set of route selection performed by all the ASes that maximizes their route preferences and that modifying their selection would result in obtaining a less preferred route.

The existence of a pure Nash equilibrium implies *solvability* of the BGP system and it is a necessary condition for its safety. However, not all networks have a pure Nash equilibrium. A particular case of a network setup without a stable state has been presented in [10] and it is shown in figure 1.1.a. It is easy to see that in that particular setting, for any route selection,

¹A game with complete information means that each player’s payoff function is common knowledge among all the players.

there is at least one AS that will be presented with a more preferred route than the one is currently using, therefore it will increase its benefit by making a different selection.

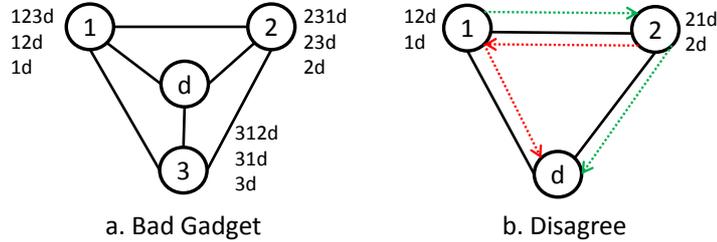


Figure 1.1: Examples of BGP system with more than one or no stable states

Unfortunately, identifying if a BGP system has a solution is a harder problem than expected.

The results obtained by Griffin et al. [10] imply that determining whether a pure Nash equilibrium in the ONE-ROUND GAME exists is NP-hard.

This basically means that assuming we have a central entity that has all the policy information for every AS, determining whether the resulting BGP system has a stable state is NP-hard.

The ONE-ROUND GAME we have presented and analyzed for modeling BGP essentially gives information about the solvability of a given BGP system but it cannot provide a realistic description of the protocol dynamics. For this reason, the ONE-ROUND GAME is only used for benchmarking. In order to capture more of the intricacies of a BGP system, we analyze a different more complex game, called the CONVERGENCE GAME that help us to understand the safety aspects of a BGP system.

The CONVERGENCE GAME is a dynamic/multiple-round game with imperfect² and incomplete information. The players and the payoff function are defined in the same way as in the case of the ONE-ROUND GAME. The CONVERGENCE GAME is a incomplete information game, which means that the payoff function of the ASes is not known by the other ASes. This implies that the route preferences of each AS (i.e. its routing policy) is private information.

In each round, the players that participate are chosen by the *Scheduler*. The *Scheduler* may also delay update messages and/or remove links and nodes from the AS graph. In each round, selected players perform the following set of actions: they apply their import policies to the route information received from neighbours, they select one route, and they announces routes to some or all of their neighbours. The action set defined for the CONVERGENCE GAME is much wider than the execution of BGP. In particular, the possible actions include announcing a route that is different from the one that has been selected and used by the local AS.

Performing BGP is a strategy in which the players repeat the following actions: read update messages from neighbors, choose the most preferred route for traffic forwarding and announce it to all neighbors. The BGP behaviour is modeled by the *best-reply dynamics*

²A multiple-round game with imperfect information means that all players do not know all the previous actions taken by all other players in the game

strategy-profile. *Best-reply dynamics* is defined in the following way: starting with an arbitrary strategy profile, in every round of the game some players switch their strategies to be the best reply to the current strategies of the other players. This best-reply dynamics process can converge to an equilibrium point, i.e. to a point where the decision of all players is invariant.

The ONE-ROUND GAME and the CONVERGENCE GAME are related through the best-reply dynamics in the following way.

If there is an unique Nash equilibrium point in the ONE-ROUND GAME, then the outcome of performing best-reply dynamics in the CONVERGENCE GAME is that equilibrium point[19].

However, there are networks that have more than one stable states. For example, the network in figure 1.1.b has two stable states: $(1d, 21d)$ and $(12d, 2d)$. Depending on the order in which the *Scheduler* assign turns to the different players, the network may end up in one of the two possible stable states. It is even possible to find a *Scheduler* that assigns turns in a way that the network permanently oscillates between the two solutions.

Nisan et. al proved in [14] that if there are two or more Nash equilibria in a game, then best reply dynamics can potentially oscillate.

Thus far we have learned that there is unique Nash equilibrium in the ONE-ROUND GAME, the best reply dynamics of the CONVERGENCE GAME converge to it and if there is more than one pure Nash equilibrium point in the ONE-ROUND GAME, the CONVERGENCE GAME may oscillate. A necessary and sufficient condition that should guarantee BGP safety has not yet been found. However, in [9] the authors propose a sufficient condition for safety in the interdomain routing, i.e. *No Dispute Wheel*, which has become the broadest condition known to guarantee that BGP is safe. A *Dispute Wheel* is a structure consisting of a cyclic dependency between nodes' routing preferences that cannot be simultaneously satisfied. They prove that the lack of such a structure in a network is enough to guarantee that the network has a unique stable state and that BGP converges to this state.

This can be expressed in game theoretic terms as: if No Dispute Wheel holds for a specific network, then best-reply dynamics of the CONVERGENCE GAME converge to the unique pure Nash equilibrium of the corresponding ONE-ROUND GAME. This happens for every asynchronous schedule considered in the CONVERGENCE GAME.

For the case of dynamic games with incomplete information that are executed in distributed environments (such as interdomain routing), Shneidman et al. have proposed a different equilibrium concept from the pure Nash equilibrium, namely the ex-post Nash equilibrium[23]. The ex-post Nash equilibrium is a strategy profile from which no node wished to deviate in the CONVERGENCE GAME, since doing so would result in reduced payoff, regardless of the private information of all the other nodes. Ex-post Nash needs weaker knowledge assumptions than the ones made in the case of the pure Nash equilibrium, since it does not requires information on the network structure, the players' private preferences, or the chosen schedule in the game and it only assumes that the rationality of the nodes is common knowledge between the nodes.

Ideally, executing BGP as specified would be the ex-post Nash equilibrium of the CONVERGENCE GAME. That would mean, that it is in the best interest of every AS to comply with the BGP specifications. However, in the following sections, we show that this is not true in the general case and further constraints are needed to achieve that.

1.3 Incentives for BGP-compliant behaviour

In this section, we try to answer the question whether it is in the best interest of the ASes to follow the prescribed BGP behaviour or whether they could improve their payoff by intentionally deviating from it.

Game theory is the right tool to answer this set of questions since it allows us to model the ASes as rational agents which goal is to improve their benefits, with selfish (and often conflicting) interests. For this, we introduce the concept of *incentive-compatibility* in the BGP framework. Intuitively, incentive-compatibility means that no unilateral deviation from BGP by any AS can strictly improve the outcome of that particular AS. More formally, in the context of the CONVERGENCE GAME, we defined that the ex-post Nash equilibrium point is a strategy profile from which no player wants to deviate.

Hence, executing BGP is said to be incentive-compatible in ex-post Nash if this strategy profile is an ex-post Nash equilibrium[18] of the CONVERGENCE GAME.

For the incentive-compatibility analysis we use the same model we used in the previous section, but we consider multiple payoff functions. In order to do that we consider that every node i in the system has a private valuation function v_i , which specifies the value for each available route for the current node towards the destination. The valuation function should be interpreted as representing the nodes ability to express their preferences over routes in a unit that is comparable to other components of the payoff function. In the simplest model that we consider in the first subsection, the utility function is equal to the valuation function (like in the previous section). Later on, in the second subsection, we consider a more complex payoff function that takes into account both the preference of the selected route and the impact that the announcement of that route to other ASes has on the attraction of traffic from those ASes.

1.3.1 Incentive-compatibility without traffic attraction

In this sub-section we consider the case where the payoff function of the players is equal to the valuation function of the routes.

For that restricted case, Levin et al. prove that BGP is not incentive-compatible in ex-post Nash, even if No Dispute Wheel holds[16].

This means that even if the network is guaranteed to be safe, there may exist one node which has the incentive to *unilaterally* deviate from the BGP prescribed behaviour, because by doing so he would be improving his benefit. In figure 1.2 we have an example of a network in which node m has the incentive to deviate in order to increase its payoff. It does so by announcing the non-existing route with AS path " md ". The result is that m 's most preferred route with AS path " $m1d$ " is now available to m , increasing its payoff.

Furthermore, this negative result is valid even if the network is consistent with the Gao-Rexford constraints[6]. In a nutshell, Gao-Rexford applies to networks where the only types of business relations between ASes are peer-to-peer and customer-provider relations, and where there are no customer-provider cycles. In this type of networks, safety is guaranteed if ASes set their preferences in such a way that they prefer customer routes over peering routes and peering routes over provider routes[6]. However, these constraints are not enough to guarantee that no node in the network will *unilaterally* deviate from the BGP prescribed behaviour.

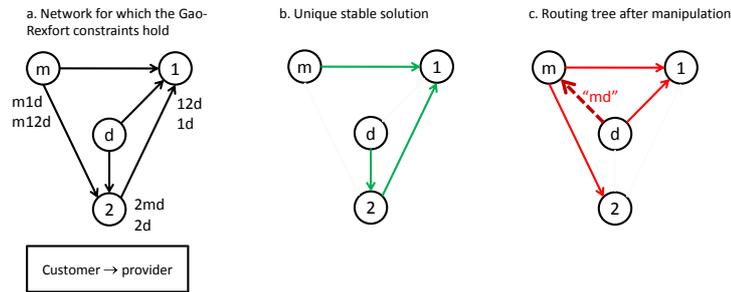


Figure 1.2: BGP is not incentive-compatible even if Gao-Rexford constraints hold.

Given the negative results on incentive-compatibility on BGP, it is natural to ask whether it is possible to change BGP so that the result is incentive-compatible. This is the goal of a related area of study, called *Mechanism Design*[18] which can be seen as the reverse engineering approach to game theory. The purpose of *Mechanism Design* is to create games in which the desired behaviour emerges as an equilibrium of selfish participants, independently of the participants' unknown true preferences. Incentivizing ASes to adhere to BGP can be done in two ways: by restricting ASes' policies, which are considered in the rest of this section or by providing monetary motivations, which are considered in section 4.

We focus for now that use *different types of constraints* (e.g. routing policy constraints, security constraints, control-plane verification mechanism) to provide the necessary motivation for ASes not to deviate from BGP.

The mechanism proposed in [5] uses a condition called Policy Consistency that, together with No Dispute Wheel, guarantees that BGP is incentive-compatible in ex-post Nash. Policy Consistency means that given two routes to the destination, any two neighbouring nodes share the same preference about which one is better.

So far we have focused in incentive-compatibility which can be understood as the case in which if all nodes but one adhere to BGP, then the remaining node's best choice is to also follow the prescribed behaviour. It may happen that more than one node have the incentives to coordinate in their misbehaviour. For this reason, we wish that the network would be *collusion-proof*. We call a network collusion-proof when no deviation from BGP by any group of ASes of any size can strictly improve the routing outcome of even a single AS in the coalition without strictly harming the others. It is proven that a BGP system complies both aforementioned conditions is also collusion-proof.

This is however a very restrictive condition imposed on the interdomain routing, that may not hold in real-life scenarios.

For this reason, Levin et. al [16] prove that a less restrictive condition, called Route Verification, together with No Dispute Wheel, guarantees the incentive-compatibility of a BGP system.

Route Verification means that a node can verify whether a route announced by a neighboring node is indeed available to that particular node. *Route Verification* can be implemented in the real-world BGP by adding new security policies to the protocol, like in the case of Secure BGP (S-BGP)[15]. The previous conditions also guarantee collusion-proof.

1.3.2 Incentive compatibility with traffic attraction

We now consider a more complex but also more realistic model, where the payoff function of the players also takes into account the benefit resulting from attracting traffic [8]. The rationale for this is that ASes increase their benefit not only by choosing their most preferred path towards the destination but also from providing transit services for other ASes. A clear case where this applies is for a transit ISP that charges its customers per traffic volume.

In order to take incorporate transit considerations into the model, Golberg et. al [8] extends the payoff function to have two terms, namely, the valuation function which represents the preferences for outgoing routes and the attraction function that expresses the benefits a node could receive from attracting traffic.

Unfortunately, when traffic attractions are considered, neither Route Verification[16] nor Policy Consistency[5] are sufficient to guarantee that BGP is incentive-compatible.

This is true, even if safety condition hold i.e. if there is no-Dispute wheel and/or the Gao-Rexford conditions hold.

The network configuration presented in figure 1.3.a shows that, when traffic attractions are considered, a node may have incentives to make dishonest announcements in the network, even if Route Verification holds. The manipulating node m has incentives to announce the path md to node c , even if it is actually using path $m1d$. By doing so, node m gains by attracting more traffic from its customer node c . The announcement can be made even if Route Verification is implemented, because node 1 announced $1d$ to m .

In figure 1.3.b we see here that in a policy consistent network (for which Route Verification does not hold), a manipulator m can have an incentive to announce a nonexistent path in order to attract traffic from its customer c . In this case, the manipulator m announces a false path md that is not available to it (because d does not export this path to m). By doing so, he deviates from the prescribed BGP-compliant strategy and gains a traffic attraction that he could not have otherwise.

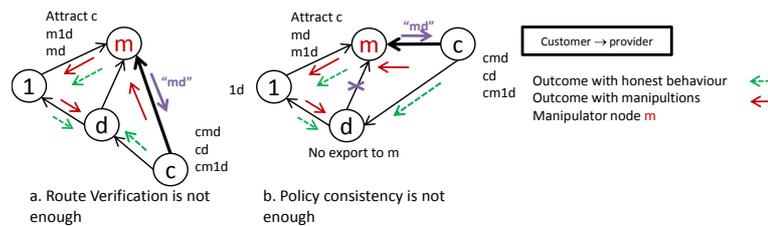


Figure 1.3: Route Verification or Policy Consistency alone is not enough in a network with traffic attractions.

Since the previous conditions are not enough to guarantee incentive compatibility for BGP when traffic attraction is considered, [8] defines the following new constraints: First, they define that a node a uses *next-hop policy* if the preference for a given route is solely determined by its next hop. Second, they introduce *Loop Verification*, which is a mechanism that verifies that if a route is discarded by BGP because it contains a loop in the AS path, the loop actually exists. The goal is to prevent misbehaving ASes from using the BGP loop prevention mechanism as a tool to direct a remote AS to discard a route based on a non-existent loop. *Loop Verification* is a form of *Route Verification* restricted to the routes that

cause loops.

In addition, the *the all-or-nothing export* is introduced. The *the all-or-nothing export* rule assumes that for each neighbor, a node can either export all admitted paths or none at all.

Using these definitions, Golberg et. al [8] prove that Next-hop Policies and Loop Verification are enough to guarantee that BGP is incentive-compatible when any form of traffic attraction is considered as long as all the ASes use the the all-or-nothing export policy.

The previous result applies to a generic form of traffic attraction function. Golberg et. al [8] also provide different sets of constraints for the case that a more detailed traffic attraction function is considered (e.g. when the payoff of the ASes attracting traffic only increases when it attracts traffic from their customers).

1.4 Payment Mechanisms for Incentivizing ASes

In this section we focus on mechanisms that make use of monetary transfers in order to make sure that the rational players behave in the game. Because of that, the model used significantly differs from the ones used in the previous sections.

In the case of the mechanism design problem for interdomain routing [4], each agent is supposed to have a private information that accounts for the transit cost c_k that it incurs when it forwards a transit packet. The goal in [4] is to design a mechanism that *results in sending each packet along the lowest-cost path (LCP), according to the true cost vector $c = (c_1, c_2 \dots c_n)$* . The mechanism they propose achieves overall routing efficiency by minimizing the total *real* cost of routing all the packets.

It is assumed that a node k incurs a cost c_k for a packet sent from i to j if and only if k lies on the selected route from i to j . For computing the total path cost in a distributed manner, the ASes have to input all their *real costs*. Consequently, the mechanism wants that the agents report their private information about their costs truthfully. In order to assure the truthfulness of the mechanism, is permitted to use monetary transfers as economic incentives.

The proposed mechanism define a payment p_k that each node k receives to compensate for carrying transit traffic. The payoff for agent k is the payment p_k minus the total cost incurred by a node k . The authors design a *payment mechanism* which takes as a input the AS graph and the vector c of declared costs, and outputs the set of LCPs and prices. The payments take the form of a per-packet price.

The mechanism proposed in [4] is part of the so-called Vickrey-Groves-Clarke (VGC) mechanism class. The VGC mechanisms are ones for which the goal is to maximize the summation of the valuations for all the agents in the network. The most important characteristic of this type of mechanism is the fact that ***a VGC mechanism is always incentive-compatible.***

This mechanisms can be implemented as an extension to BGP that carries the resulting prices, causing only a minor increase in routing table size and convergence times. The price computation algorithm computes the prices at each node from the prices and costs at neighboring nodes.

However, the resulting mechanism requires payments for each node along to path towards the destination, which differs from the current Internet pricing model. In [11] the

authors design a mechanism that achieves the same goal as the previous one, but it only requires per-packet compensation between adjacent nodes, more similar to the current Internet payment model.

1.5 Conclusions

In this survey we have presented the most relevant results obtained from modeling interdomain using game theory. Such approach provides a new perspective as ASes are modeled as *strategic agents* seeking to improve their benefits through the manipulation of BGP.

A key result obtained is that BGP in its current form is not incentive-compatible, which means that ASes can increase their benefits by not complying with the specified BGP behaviour. In order to guarantee incentive-compatibility in the interdomain routing, there are two available approaches: impose additional constraints in the behaviour of the ASes, or incentivize ASes to adhere to BGP via VGC payments.

The combination of constraints imposed on the operations of the BGP protocol depends on the payoff function considered in the game theoretic model analyzed. However, it seems that a practical corollary of the incentive-compatibility analysis is that the implementation of security extensions for BGP such as S-BGP[15] would help to remove the incentives for ASes to misbehave.

Chapter 2

The Tragedy of the Internet Routing Commons

2.1 Introduction

The Internet community has been facing important challenges concerning the scalability of the global routing system [1], [17], [13]. The interdomain routing relies on the Border Gateway Protocol (BGP) to exchange reachability information between the Autonomous Systems (ASes). All the routing information is stored in the *Global Routing Table* (GRT). Thus, routers are permanently involved in a constant exchange of network prefixes and other routing information to keep every router informed on how to reach hundreds of thousands of other networks on the Internet.

It has been argued that the computation power and memory requirements for a router that could sustain the growing Internet can be met in light of the advances being made in multi-core processor design and large memories [3]. However, even if the technology can cope with the scalability challenges, the necessary periodic upgrades resulting from the dramatic increase of the routing table may significantly increase Internet operator's expenditures on new equipment, challenging the economic viability of the Internet as we know it.

The rapid size amplification of the GRT can be partially explained as a result of the increasing popularity of the Internet and, consequently, of the large number of new users that attach to it[2]. Even if this is certainly true, the growth due to the augmentation in number of reachable networks is not the only cause for the BGP routing table expansion[17]. There are different other factors that triggered this rapid growth, like multi-homing and traffic engineering [13].

Evidently, the increased number of prefix allocations partially explains the growth of the Global Routing Table. However, a single address space allocation may translate in multiple routing table entries in the GRT as a result of *address space fragmentation*. Therefore, when a new block is allocated by an Internet Registry, this prefix may be broken into several smaller prefixes that are then announced in the Internet, inducing a faster growth rate of the GRT than the rate with which the new address blocks are being allocated.

This process is commonly known as *deaggregation*. Such behaviour is not without controversy, as it acts counter to the goals of the current Classless Inter Domain Routing (CIDR) architecture, which encourages aggressive address aggregation [21]. The technical reasons

for which the networks are using these more-specific prefixes have been studied in more detail in [1, 17].

It has been shown that, in theory, the current BGP routing tables size can be reduced by a factor of more than 2 if the more specific prefixes could be integrated in the already existing covering routing entries[2, 1]. Therefore, by announcing these more specific prefixes, the ASes are contributing to the explosive inflation of the GRT. The resulting additional cost generated by this type of behaviour is described as an *externality*. Generally, an externality is the additional cost that the behaviour of a single agent may bring to the other agents, without any voluntary agreement between the two that would allow for a negotiation of the distribution of these costs.

The size of the routing table is an externality because when a network is deaggregating it obtains a benefit which is far greater than the cost this operation incurs and it does not consider the additional cost it brings to the other networks. In other words, ASes deaggregate on the expense of all the members of the Default-Free Zone.

This problem has been described as a *tragedy of the commons*[12] in [13, 21]. In this paper we propose a game theoretic model of the Internet that aims to study the dynamics of the interaction between ASes in the Internet and to evaluate to which extent the Internet public resources generate tragedy.

In Section 2.2 we consider a *commons model* in which the ASes sharing the same BGP routing space have to make the choice of deaggregating or not the assigned address blocks. One important result of our analysis is that the economic incentive for the ASes in the interdomain is to engage a detrimental behaviour of heavily deaggregating their prefixes, thus leading to a *tragedy of the commons*.

In Section 2.3 we study the properties of the *game equilibria* and we examine its relation to the *social welfare* point of the considered game setup.

Finally, we try to evaluate one of the possible techniques of avoiding the tragedy of the commons and of bringing the equilibrium point closer to the social optimum in Section 2.4. We propose a pricing mechanism to internalize the additional costs. We conclude the paper in Section 2.5.

2.2 The Game Theoretic Model

Since the seminal article by Hardin [12], “*The Tragedy of the Commons*” has been used to model the problems of overuse and degradation of resources, with a wide range of applications. The majority of the research work using this model was focused on the degradation of natural resources including water resources or oil, the destruction of fisheries, deforestation and more. According to Hardin[12], a tragedy of the commons occurs when a group of entities abuses a common resource, thus harming the other individuals with whom it shares this resource and themselves. In the paper the author argues that any common natural resource “remorselessly generates tragedy”, since the gain an individual user receives from increasing its consumption from the common outweighs its own cost and all the users of the resource will evenly pay for this.

In Hardin’s model, each player is making an *individual* decision that is actually creating a *global* effect. The players do not adequately pay for the impact of their decision. This cost

is *externalized* and is supported by all the others. The result of such events is an inefficient exploitation of the resources, which leads to the abusive consumption of the common good.

The game theoretic model presented in this paper is based on the previous analysis of the tragedy of the commons made in [7]. It is a symmetric one shot game with perfect and complete information, in which the players have to simultaneously choose the degree of individual exploitation of the common natural resource.

In real life settings, the games played are typically quite complex and have large strategy spaces that are not specified beforehand. Here we present a simplified model of the interdomain routing commons problem. We consider a number of N ASes in the interdomain, interconnected to form the Internet as we know it. For simplifying purposes, we assume that each AS is represented by a single router.

We analyze a symmetric one-shot complete-information game. The players are the N ASes that have to simultaneously decide the number of prefixes to announce in the Internet. A strategy for a player i is the number of prefixes it originates in the interdomain, p_i . We assume that the strategy space is $[0, \infty)$ to cover all the possible choices that could be of interest to the AS. Considering the game setup presented above, we propose the following expression for the payoff each player i receives from announcing p_i prefixes, given that the other players are announcing $(p_1, p_2, \dots, p_{i-1}, p_{i+1}, \dots, p_N)$ prefixes is:

$$u_i(p_i, p_{-i}) = p_i v(P) - c \min\{P_{max}, P\}, \quad (2.1)$$

where p_{-i} is the vector of strategies chosen by all the other players but player i , $P = p_1 + p_2 + \dots + p_{i-1} + p_i + p_{i+1} + \dots + p_N$ is the total number of announced prefixes in the Internet, P_{max} is the maximum number of prefixes that can be stored in a router in the interdomain, c denotes the cost incurred by an AS storing one prefix and, finally, $v(P)$ denotes the value each AS receives each prefix announced. We explain all these concepts in more detail next.

The utility function is the difference between the total gain that an AS receives from announcing p_i prefixes ($p_i v(P)$) in the interdomain and the cost for storing all the announced routing information up to the maximum permitted limit ($c \min\{P_{max}, P\}$).

In this model, we consider that the *Global Routing Table* is a *finite common resource*. There is a direct dependency between the size of the GRT and the *router memory size*. We model the hardware constraints imposed by the routing equipment developers by assuming that P_{max} is the maximum number of prefixes that can be fitted in the memory of the largest capacity state-of-the-art router on the market.

Assuming that each routing entry needs the same amount of memory in order to be stored in the routing table, we consider that every prefix has assigned an individual memory slot in the routing equipment. The cost each AS has to pay for storing a prefix is the price of this memory slot, c .

We assume that every AS stores all the routes and upgrades its routing equipment until the maximum capacity of the routers is reached. The limitation in router memory capacity is reflected in the payoff function for each AS i by the term $c \min\{P, P_{max}\}$, i.e. the total routing cost. We assume that when the total number of announced prefixes is higher than the maximum number of possible entries in the GRT, the ASes start randomly discarding prefixes.

By $v(P)$ we denote the value each AS receives from announcing one prefix in the inter-domain. The total benefit each AS receives is equal to $p_i v(P)$. Deaggregation proponents claim that this provides a very effective and fine-grained method of performing the traffic engineering much needed in the current Internet. By increasing the granularity of the advertisements through the use of variable prefix lengths, the Internet Service Providers (ISPs) achieve a better control of the distribution of their traffic over transit links. Certain networks may choose to announce more-specific prefixes than the ones allocated with the purpose of increasing their security. In particular, the ASes sourcing deaggregated prefixes protect the network from prefix hijacking or other types of malicious attacks.

Formally, for $P < P_{max}$, $v(P) = const.$ and for $P > P_{max}$, $v'(P) < 0$ and $v''(P) < 0$. An example of such a function is depicted in fig. 2.1.

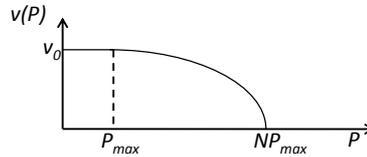


Figure 2.1: The generic shape of the value function

As we can see in fig. 2.1, when the total number of routing entries P is lower than P_{max} , the value per prefix is constant, every router receiving the same benefit from announcing one more prefix in the Internet (v_0). Therefore, while below this threshold, announcing one more prefix in the Internet does not affect the received per prefix value, imposing only the additional cost of storing that new routing entry.

When the number of the entries in the GRT grows to more than the maximum number admitted, and the ASes are randomly discarding prefixes from their routing tables. Thus, certain destinations are not reachable by some of the ASes, thus modifying the value that the agents expected to receive. This physical limitation is reflected in the value function $v(P)$: when $P > P_{max}$, the decrease in value is notable. We assume that this is the only factor that is influencing the benefit received per announced prefix.

When the number of prefixes announced by each of the N ASes in the Internet is equal P_{max} (thus, the total number of prefixes in the Internet being $P = NP_{max}$), the value received per prefix is zero, as the ASes will use their entire capacity only to store their own announced prefixes and, consequently, no network would be able to reach any other one in the Internet.

2.3 Game Equilibria and Overall Optimality

In this section we analyze the equilibria properties of the interdomain routing game. We focus on the relation between the *Nash Equilibrium* of the game and the *Social Welfare* point.

2.3.1 Game Equilibria

We consider $(p_1^*, p_2^*, \dots, p_i^*, \dots, p_N^*)$ to be the Nash Equilibrium strategy profile of the routing game. Thus, every p_i^* is maximizing the individual payoff $u_i(p_i, p_{-i})$ for each player i

in the game, given that the other players choose $p_{-i}^* = (p_1^*, p_2^*, \dots, p_{i-1}^*, p_{i+1}^*, \dots, p_N^*)$:

$$\max_{p_i} u_i(p_i, p_{-i}^*) = \max_{p_i} \{p_i v(P) - c \min\{P_{max}, P\}\} \quad (2.2)$$

In order to find the p_i value that solves the above-mentioned equation, we first consider the strategy profile of the game so that the total number of prefixes in the Internet is below the maximum router capacity. When $P < P_{max}$, the cost for the routing equipment is depending of the total number of prefixes announced in the Internet. The maximum available capacity of the routers is not reached, so announcing one more prefix will not have any negative impact on the per prefix value and it will only incur the extra cost of upgrading the routing equipment. Moreover, the value function is constant: $v(P) = v_0$. Therefore, the payoff function has the following form:

$$u_i(p_i, p_{-i}) = p_i v_0 - cP, \quad (2.3)$$

The first order condition for the optimization problem is:

$$v_0 - c = 0. \quad (2.4)$$

Thus, we have to consider now the relation between the cost of memorizing a prefix and the value per announced prefix. If the cost of a memory slot c is higher than the value received per announced prefix v_0 , then the ASes have no economic incentive for announcing prefixes in the Internet. Hence, we practically have no routing operation in the interdomain. Therefore, this case does not present interest for our analysis.

However, when the cost of a memory slot c is below the value received per announced prefix v_0 , the ASes have no incentive to stop announcing more specific prefixes. Therefore, they will keep deaggregating the assigned address blocks, as they will always receive a positive benefit ($v_0 - c > 0$). The lack of incentives for the ASes to reduce their deaggregating behaviour eventually leads to using the entire available capacity of the routers. Therefore, assuming that the ASes are greedy economic entities whose only purpose is to increase their benefits, we can easily conclude that the equilibrium point of the considered game is not a strategy profile that verifies $P < P_{max}$. Hence, we search next for the equilibrium profile strategy when $P \geq P_{max}$.

We are relying our analysis on the assumption that every AS is willing to pay for upgrading its routing equipment up to the maximum limit. All the ASes are paying for the highest capacity router, thus the equipment expenditure is the same for all of the players in the game. This is reflected in the payoff function by the presence of the cP_{max} factor.

The inefficient exploitation of the router memory is captured in the value each network receives from announcing a prefix in the Internet. When the total number of prefixes P exceeds the maximum admitted number of memory cards in the router P_{max} , the value per announced prefix is not constant anymore, as the ASes start randomly discarding prefixes from the routing table. As a result, the benefit for those particular networks decreases. Therefore, the value function considered when $P \geq P_{max}$ is a fast decreasing function, as shown in fig. 2.1. Consequently, the payoff function has the following particular expression:

$$u_i(p_i, p_{-i}) = p_i v(P) - cP_{max}. \quad (2.5)$$

In this case, the first order condition for the optimization problem corresponding to the Nash equilibrium, when all the other players choose the equilibrium strategies $p_{-i}^* = (p_1^*, p_2^*, \dots, p_{i-1}^*, p_{i+1}^*, \dots, p_N^*)$ is:

$$\frac{du_i(p_i, p_{-i}^*)}{dp_i} = 0 \Rightarrow v(p_i + P_{-i}^*) + p_i v'(p_i + P_{-i}^*) = 0, \quad (2.6)$$

where $P_{-i}^* = \sum_{j \neq i} p_j^*$. Substituting p_i^* into (2.6), summing over all the N ASes and dividing by N outputs

$$v(P^*) + \frac{P^*}{N} v'(P^*) = 0. \quad (2.7)$$

Due to the symmetry of the game, we can conclude that the Nash equilibrium number of prefixes each AS i is originating in the Internet is the same for all the players, i.e. $p_i^* = \frac{P^*}{N}, \forall i$. Therefore, the equilibrium point of the game will be reached when the total number of announced prefixes is equal to the value P^* that satisfies equation (2.7). In other words, an AS will stop announcing prefixes in the interdomain (and, thus, reach an equilibrium point) when the value received from originating one prefix would be equal to the harm done only to its own already stored prefixes, i.e. $p_i^* v'(P^*)$.

2.3.2 Social Optimality

In this section we analyze the properties of the social welfare strategy profile and its relation with the equilibrium strategy profile. In contrast with the Nash Equilibrium, the Social Welfare strategy profile, denoted by $(p_1^{**}, p_2^{**}, \dots, p_i^{**}, \dots, p_N^{**})$, maximizes the sum of payoffs. When the number of announced prefixes in the Internet is lower than the maximum limit ($P < P_{max}$), substituting the expression for $u_i(p_i, p_{-i})$ with the expression from (2.3) outputs:

$$\max_{P \geq 0} \{Pv(P) - NPc\}. \quad (2.8)$$

As the value function is constant for $P < P_{max}$, similar to case of determining the Nash equilibrium, the strategy chosen by the players depends on the ratio between the per prefix value and the cost implied by storing a new routing entry in all the routers in the Internet. If the benefit received for announcing one more prefix is lower than the cost incurred ($v_0 < NC$), then the ASes have no incentive to announce prefixes in the Internet. Otherwise, the players will announce all that they can, thus reaching the maximum capacity of the routing equipments. Therefore, we can easily conclude that the overall optimum is not reached when the total number of announced prefixes is strictly below the maximum limit.

When the number of announced prefixes exceeds the maximum router capacity ($P \geq P_{max}$), the first order condition for the optimization problem corresponding to determining the Social Welfare strategy profile has the following expression:

$$\frac{d}{dp} \{Pv(P) - NP_{max}c\} = 0 \Rightarrow v(P) + P v'(P) = 0. \quad (2.9)$$

Substituting the Social Welfare strategies in (2.9) we obtain

$$v(P^{**}) + P^{**} v'(P^{**}) = 0, \quad (2.10)$$

where $P^{**} = \sum_i p_i^{**}$. Therefore, in (2.10) we can see that an AS will stop announcing more prefixes when the value per prefix is equal to the harm done to *all* the prefixes announced in the Internet by all the ASes, i.e. $-P^{**}v'(P^{**})$.

Due to the fact that ASes are modeled as rational agents, it is not in their benefit to consider the harm that their strategies are doing to the other players in the game. For this reason, the strategies chosen by the players in the equilibrium point may lead to an inefficient equilibrium, that might not coincide with the social optimum strategy profile.

2.3.3 Evaluation of the Model

In the economic model of the Internet proposed in this paper, we assume that the ASes only pursue their own material self-interest and do not care about “social” goals. The main concept of rational behaviour, the Nash equilibria, is known not to always optimize the social outcome. It should be noted that the payoff each player receives in the Nash equilibrium is *lower* than the social optimum payoff, underlying the inefficient consumption of the common resource. Given the game setup presented in section 2.2, we evaluate which is the ratio between the the overall optimum and the Nash equilibrium of the game.

As reflected in the first order condition from (2.6), the incentive for an AS in the Nash equilibrium point is to announce one more prefix considering only the harm caused to its own already announced routing information. The common resource, i.e. the router memory, is over-utilized because each player is rational and is only considering its own benefits and costs. Hence, in (2.7) we have the value of $\frac{P^*}{N}v'(P^*)$ in opposition with $P^{**}v'(P^{**})$ in (2.10).

Moreover, comparing (2.7) with (2.10) we can easily prove¹ that the number of prefixes announced in the Nash equilibrium is strictly higher than the total number of prefixes announced in the social optimum:

$$P^* > P^{**}. \quad (2.11)$$

The ratio between the values for the total number of prefixes at the equilibrium point and at the social optimum depends on the specific form of the value function. Next, for the particular case of a value function $v(P) = a - P^2$, where $a = \text{const.}$, we analyze the ratio between the two above-mentioned values when $P \geq P_{max}$. It can be easily checked that this particular expression of value function complies with the set of rules imposed on the value function defined in the game setup. Thus, since the particular expression of the value function must verify $v(NP_{max}) = 0$, we can express parameter a as $a = N^2P_{max}^2$.

Following the Nash equilibrium analysis in 2.3.1 and the social welfare analysis in 2.3.2, we obtain the undermentioned values for the total number of prefixes in each case:

$$(P^*)^2 = \frac{aN}{N+2}; \quad (P^{**})^2 = \frac{a}{3}. \quad (2.12)$$

¹Extending the proof in [7] for our model, let us assume, by the contrary, that $P^* \leq P^{**}$. Since $v'(P) < 0$, then $v(P^*) \geq v(P^{**})$, and since $v''(P) < 0$, then $0 \geq v'(P^*) \geq v'(P^{**})$. Finally, since $\frac{P^*}{N} < P^{**}$, the left term of equation (2.7) strictly exceeds the left term of equation (2.10), which is impossible since both are equal to zero.

Therefore, the ratio between the total number of prefixes in the social optimum (P^{**}) and in the equilibrium (P^*) is

$$\left(\frac{P^{**}}{P^*}\right)^2 = \frac{1}{3} \left(1 + \frac{2}{N}\right) \Rightarrow \frac{P^{**}}{P^*} \approx \sqrt{\frac{1}{3}}. \quad (2.13)$$

Both the Nash equilibrium and the overall optimum strategy profiles result in a total number of prefixes which exceeds the maximum limit. Consequently, all the ASes are paying the same routing cost, cP_{max} .

As a measure of the global routing system efficiency, we evaluate in this section the price of uncoordinated utility-maximizing decision between the players, also known as *the price of anarchy*[20]. By analyzing the price of anarchy we are trying to determine how much better would the Internet run if there existed a central authority who could coordinate the operations of every entity in the Internet.

The *price of anarchy* (PoA), a standard measure of the suboptimality introduced by self-interested behavior, is defined as the ratio between the optimal “centralized solution” value of the total benefit ($p_i^{**}v(P^{**})$) and the worse equilibria value of the total benefit for a single player ($p_i^*v(P^*)$)[22].

Therefore, given the game setup in section 2.2, the Nash equilibrium strategy profile studied in 2.3.1, the social optimum studied in 2.3.2 and using the results from (2.7) and (2.10), the price of anarchy has the following expression:

$$PoA = \frac{p_i^{**}v(P^{**})}{p_i^*v(P^*)} \Rightarrow PoA = N \left(\frac{P^{**}}{P^*}\right)^2 \frac{v'(P^{**})}{v'(P^*)}. \quad (2.14)$$

where, for each AS i , p_i^{**} denotes the social optimum strategy and p_i^* denotes the Nash equilibrium strategy. For the value function with the above-mentioned particular expression, $v(P) = a - P^2$, where $a = const.$, the PoA is

$$PoA = N \left(\frac{N+2}{3N}\right)^{\frac{3}{2}}. \quad (2.15)$$

Therefore, when $N \rightarrow \infty$ then $PoA \rightarrow \infty$. This means that when the number of ASes grows, the Nash equilibrium moves further away from the social optimum. Both results in (2.11) and (2.14) prove that the self-interested behaviour of the agents in the Internet leads to a suboptimal outcome, with an inefficient exploitation of the common resource.

2.4 A Payment Mechanism

In this section, we propose a solution for avoiding the tragedy of the Internet routing commons. One classic approach is to internalize the costs incurred by the commons problem using monetary exchange in order to obtain a certain outcome.

We consider the implementation of a payment mechanism in the current Internet so that the social welfare outcome is achieved.

Given the optimal strategy profile ($p_1^{**}, p_2^{**}, \dots, p_i^{**}, \dots, p_N^{**}$) studied in 2.3.2, we set the payment to exactly reflect the harm one AS causes to all the other ASes when announcing

one more prefix in the Internet. The payment is a per prefix monetary amount introduced into the payoff function so that the outcome of the new game coincides with the previous social optimum. Using the payoff function that has the expression in (2.5) and the overall optimum strategy profile, the per prefix payment introduced by the mechanism is:

$$x_i = - \sum_{j \neq i} p_j \frac{dv(P)}{dp_i} \Rightarrow x_i = -(P^{**} - p_i^{**})v'(P^{**}). \quad (2.16)$$

The new strategy profile of the ASes has to maximize the individual payoff function, taking into account the additional cost paid by each AS for each announced prefix:

$$\max_{p_i} \{u_i(p_i, p_{-i}) - p_i x_i\}. \quad (2.17)$$

Now we prove that the strategy profile corresponding to the social welfare is a solution for the previous equation. Consequently, by including this payment mechanism into the payoff function, the ASes have the incentives to play the strategy corresponding to the social optimum of the original game setup. The first-order condition for this optimization problem is:

$$\frac{du_i(p_i, p_{-i})}{dp_i} - x_i = 0. \quad (2.18)$$

We show next that the overall optimum strategy profile is a solution for (2.18). Using (2.16) and (2.5) in equation (2.18) yields:

$$v(P) + p_i v'(P) + (P^{**} - p_i^{**})v'(P^{**}) = 0 \quad (2.19)$$

Therefore, substituting p_i^{**} in the equation above outputs:

$$v(P^{**}) = -P^{**}v'(P^{**}). \quad (2.20)$$

We can easily see that the previous result coincides with the social optimum value in (2.10). Hence, we can conclude that the social optimum is solving equation (2.18) and thus maximizing the payoff function for all the agents in the game.

Even if by introducing this payment mechanism we bring the game equilibria in the social optimum point, the real-life implementation of such a system is not without controversy, as it would imply that the ASes would engage in monetary exchange that can influence their strategy choice.

The monetary incentives can be included in the game setup as an enforced tax or by creating a market with flexible prices and many traders. If the payment is included in the game as a fixed per prefix tax, the difficulty of the implementation would be calculating its exact value.

However, if implemented as a pricing mechanism in a market with multiple traders, the above-mentioned inconvenience of introducing payments is avoided. The economic equilibrium of the market would be easily reached through the pricing mechanism. In such a scenario, the networks would have to pay to their providers for announcing more specific prefixes. Hence, the ISPs could actually benefit from the deaggregating behaviour of their customers. Such a modification of the game could imply changing the payoff each agent receives, thus altering the incentives each entity has for announcing prefixes in the Internet and possibly changing the final outcome of the game.

2.5 Conclusions

In this paper we have proposed a simplified economic model of the interdomain routing commons problem. Considering the GRT to be a public resource, we have modeled the ASes as rational agents who have to make a strategy choice of how many prefixes to announce in the Internet so that they maximize their own benefit. We found that the game reaches an inefficient outcome because the common resource is over utilized when the total number of prefixes exceeds the maximum limit, facing a case of *the tragedy of the commons*. Finally, for avoiding the tragedy of the commons, we introduce a taxing mechanism where each AS has to pay a cost proportional with the number of prefixes it announces.

References

- [1] T. Bu, L. Gao, and D. Towsley. On characterizing BGP routing table growth. *Computer Networks*, 45(1):45–54, 2004.
- [2] L. Cittadini, W. Mhlbauer, and S. Uhlig. Evolution of Internet Address Space Deaggregation: Myths and Reality. *IEEE Journal on Selected Areas in Communications, special issue on Internet Routing Scalability*, 2010.
- [3] K. Fall and P. B. Godfrey. Routing Tables: Is Smaller Really Much Better? In *HotNets-VIII*, October 2009.
- [4] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A bgp-based mechanism for lowest-cost routing. *Distrib. Comput.*, 18:61–72, 2005.
- [5] J. Feigenbaum, V. Ramachandran, and M. Schapira. Incentive-compatible interdomain routing. In *Proceedings of the 7th Conference on Electronic Commerce*, pages 130–139. ACM Press, 2006.
- [6] L. Gao and J. Rexford. Stable Internet Routing without Global Coordination. *IEEE/ACM Transactions on Networking*, 9:681 – 692, December 2001.
- [7] R. Gibbons. *A Primer in Game Theory*. Pearson Education, 2001.
- [8] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright. Rationality and Traffic Attraction: Incentives for Honest Path Announcements in BGP. In *SIGCOMM*, 2008.
- [9] T. G. Griffin, F. B. Sheperd, and G. Wilfong. The Stable Path Problem and Interdomain Routing. *IEEE/ACM Transactions on Networking*, 10:232–243, April 2002.
- [10] T. G. Griffin and G. Wilfong. An analysis of BGP convergence properties. *Proceedings of the ACM SIGCOMM’99*, pages 277–288, 1999.
- [11] A. Hall, E. Nikolova, and C. Papadimitriou. Incentive-compatible interdomain routing with linear utilities. In *WINE’07: Proceedings of the 3rd international conference on Internet and network economics*, pages 232–244. Springer-Verlag, 2007.
- [12] G. Hardin. The Tragedy of the Commons. *Science*, 162:1243–1248, December 1968.
- [13] G. Huston. Analyzing the Internet BGP Routing Table. *The Internet Protocol Journal*, 4(1), 2001.

-
- [14] A. D. Jaggard, M. Shapira, and R. N. Wright. Towards a Unified Approach to (In)Decision: Routing, Circuits, Consensus, and Beyond. *Working paper*, 2009.
- [15] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo. Secure border gateway protocol (s-bgp). *IEEE Journal on Selected Areas in Communications*, 18:103–116, 2000.
- [16] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *STOC 08*, pages 57–66, 2008.
- [17] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang. IPv4 address allocation and the BGP routing table evolution. *SIGCOMM Comput. Commun. Rev.*, 35(1):71–80, 2005.
- [18] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Games and Economic Behavior*, pages 129–140, 1999.
- [19] N. Nisan, M. Shapira, and A. Zohar. Best - Reply Mechanism. *Draft*, 2007.
- [20] C. Papadimitriou. Algorithms, games, and the internet. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 749–753, 2001.
- [21] Y. Rekhter, P. Resnick, and S. Bellovin. Financial Incentives for Route Aggregation and Efficient Address Utilization in the Internet. In *Coordinating the Internet*. MIT Press., 1996.
- [22] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 513–522, 2009.
- [23] J. Shneidman and D. C. Parkes. Specification faithfulness in networks with rational nodes. In *Twenty-Third Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, 2004.

Appendix

The Tragedy of The Commons

In his paper, Hardin argues that any common resource “remorselessly generates tragedy”, since the gain an individual user receives from increasing its consumption from the common will eventually outweighs the cost and all users of the resource will evenly pay for this[12].

The example used by the author in [12] is that of land shared by farmers for grazing their goats. In this scenario, if one of the farmers in the group decides to add one more cow for grazing, this will benefit him. However, the rest of the farmers support a part of the cost incurred by this decisions, because the common is harmed by diminishing the available land for grazing all the animals.

We present the game theoretic model included in [7]. Consider n farmers in a village who share the village green to graze their goats. Each farmer i owns g_i goats and the total number of goats in the village is $G = g_1 + g_2 + \dots + g_{i-1} + g_i + g_{i+1} + \dots + g_n$.

The cost of buying and caring for a goat is c , independent of the number of goats a farmer owns. The value each farmer receives from grazing a goat on the green considering the total number of G goats is $v(G)$. Since the author considers a finite resource and since each goat needs a certain amount of grass in order to survive, a maximum number of goats that can be grazed on the green can be defined, G_{max} . Consequently, if $G < G_{max}$ the value received for each goat verifies $v(G) > 0$ but if $G \geq G_{max}$, the value is $v(G) = 0$.

The value function, $v(G)$ is chosen so that it underlines the common resource consumption dynamic. Since the first few goats have plenty of grass, adding one more goat does very little harm on the goats already grazing. However, when many goats are grazing on the green so that they are barely surviving ($G \approx G_{max}$), bringing one more may endanger the other ones. Formally, for $G < G_{max}$, $v'(G) < 0$ and $v''(G) < 0$, like you can see in fig. 2.2.

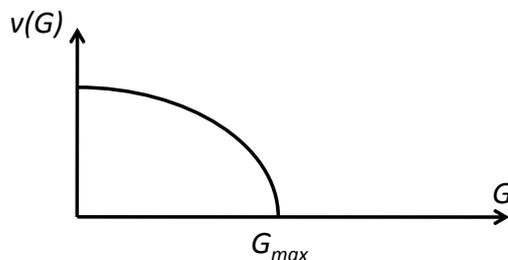


Figure 2.2: The general shape of the value function

The farmers are the ones who simultaneously have to decide how many goats to graze on the village green. A *strategy* for a farmer i is the number of goats that he owns. Assuming that the strategy space is $[0, G_{max})$, the payoff received by each farmer who is grazing g_i goats on the green when the other farmers own $(g_1, g_2, \dots, g_{i-1}, g_{i+1}, \dots, g_n)$ is:

$$g_i v(G) - c g_i, \quad (\text{A-1})$$

where $G = g_1 + g_2 + \dots + g_{i-1} + g_i + g_{i+1} + \dots + g_n$.

Considering $(g_1^*, g_2^*, \dots, g_n^*)$ to be the Nash equilibrium strategy profile, then each g_i^* must maximize (A-1) if the other farmers choose $(g_1^*, g_2^*, \dots, g_{i-1}^*, g_{i+1}^*, \dots, g_n^*)$. We obtain the following first order condition for this optimization problem:

$$v(g_i + G_{-i}^*) + g_i v'(g_i + G_{-i}^*) - c = 0, \quad (\text{A-2})$$

where G_{-i} denotes $g_1 + g_2 + \dots + g_{i-1} + g_{i+1} + \dots + g_n$.

Substituting g_i^* into A-2, summing over all n farmers' first-order conditions and dividing by n yields:

$$v(G^*) + \frac{G^*}{n} v'(G^*) - c = 0, \quad (\text{A-3})$$

where $G^* = g_1^* + g_2^* + \dots + g_{i-1}^* + g_i^* + g_{i+1}^* + \dots + g_n^*$.

The social welfare strategy profile denoted by $G^{**} = (g_1^{**}, g_2^{**}, \dots, g_i^{**}, \dots, g_n^{**})$, solves the following equation:

$$\max_{G \geq 0} G v(G) - G c, \quad (\text{A-4})$$

for which the first-order condition is

$$v(G^{**}) + G^{**} v'(G^{**}) - c = 0. \quad (\text{A-5})$$

Comparing (A-3) and (A-5) yields that $G_* > G_{**}$. The first-order condition depicted in (A-2) shows the incentives faced by a farmer who is grazing g_i goats but is considering to bring one more goat on the village green. The value of the additional goat is $v(g_i + G_{-i}^*)$ and its cost is c . The harm done on the goats already grazing on the village green is $v'(g_i + G_{-i}^*)$ to each goat, or $g_i v'(g_i + G_{-i}^*)$ to all the goats. The common is over utilized because each farmer is considering only his own benefit and not the effect of his actions on the other individuals involved. This is why the model accentuates the presence of $\frac{G^*}{n} v'(G^*)$ in (A-3) and of $G^{**} v'(G^{**})$ in (A-5).