Transit Prices Negotiation: Combined Repeated Game and Distributed Algorithmic Approach*

Dominique Barth[1] , Johanne Cohen[2], Loubna Echabbi[3,**], and Chahinez Hamlaoui[1]

[1] PRiSM, 45, avenue des etats-Unis 78035 Versailles Cedex, France
[2] lORIA, campus scientifique, BP 239, F-54506 Vandœuvre les Nancy, France
[3] INPT, Madinat Al Irfane, avenue allal el Fassi, Rabat, Maroc
echabbi@inpt.ac.ma

Abstract. We present both a game theoretic and a distributed algorithmic approach for the transit price negotiation problem in the interdomain routing framework. The analysis of the centralized transit price negotiation problem shows that the only one non cooperative equilibrium is when the lowest cost provider takes all the market. The perspective of the game being repeated makes cooperation possible while maintaining higher prices. We consider then the system under a realistic distributed framework and simulate its behaviour under a simple price adjustment strategy and analyse whether it matches the theoretical results.

Keywords: interdomain routing, repeated games, distributed algorithmic.

1 Introduction

Today inter-domain market plays on two different time scales: A long term time scale (months or even years) where economic contracts are negotiated and a short term (seconds) where routing decisions are made based on the concluded business relationships. Some recent works [1,4,5] propose to couple those two processes more tightly by enabling a more dynamic interaction between transit price propositions and routing decisions. In order to capture the dynamic aspect of such interaction, authors of these papers propose to employ a repeated game approach. The repeated game framework enables to capture how the threat of a future behaviour can impact the current actions of players.

In [1], the repeated routing game is introduced and a price matching strategy is proposed and analysed. The difference between the analysis in this work and our proposal is that we consider that the traffic dedicated to a given destination can be routed only through a single provider. This assumption is made in order to maintain a coherence with the Border Gateway Protocol where only one path

* This work was supported by The RNRT ACTRICE and the ANR SOGEA french projects.
** Corresponding author.

T. Chahed and B. Tuffin (Eds.): NET-COOP 2007, LNCS 4465, pp. 266–275, 2007.

is chosen to each destination. In our proposal, we still consider an interaction between transit price negotiation and routing decision process. The bilateral economic nature of the Internet is still maintained by a cascade like pricing where each agent negotiates low prices only with his immediate neighbours. This is different from the source based pricing approach taken in [1]. We propose an adequate model to capture the different dimensions of the problem then we focus on the analysis of the related game on some specific scenario mainly by considering the simple but not simplistic case of one source and one destination. In the game analysis, we assume a full knowledge of the different parameters of the problem, which is not very realistic but gives an idea about the nature of the game. Further, we will analyse the problem from a distributed point of view taking into account realistic considerations.

2 The Transit Price Negotiation Model

The network is given by a graph $G(N, E, cost)$ where N represents the ASs, *cost* is the unitary cost related to managing the transit over the AS intra-network and E are the physical inter-domain links. We will focus on an only one traffic flow between a node S which is the source of the traffic, it can be an access network and the node $Dest$ its destination. The rest of the nodes are providers. A random variable P represents the period on which inputs (graph, traffic Matrix) are stable. We assume that P follows an exponential law with mean D, that can be obtained by some statistical knowledge or stochastic analysis. We consider that the source has an upper bound on price under which she accepts to send the traffic. Otherwise, she does not send the traffic. We will denote it p_{max}. We consider discrete transit prices. Price discretisation depends on the encoding format in control packets, for instance here we take a unit discretization. That is provider transit price can take values as $1, 2, 3, \ldots$.

During the period P, the inputs of the problem are stable and a stationary environment game can model interactions between ASs during the transit price negotiation. Each AS announces its transit price to his neighbours with the corresponding route into the destination. When an AS decides to buy a route from its neighbour, he can itself announce this route to his own neighbours while proposing an adequate transit price. Thus, the negotiation follows a cascade like model from the destination backward to the source, where each AS in the path plays both the customer and the provider role. The objective of each provider is clearly to maximize its own benefit by proposing attractive transit prices but also by choosing itself the lowest providers. In case of identical announces, an AS can choose a provider following a pre-order on his providers. Our goal is to analyse equilibrium situations where ASs do not have the incentive to deviate from their proposed prices and to check whether such situations are beneficial to the sender (the source).

The game proceeds in series of stages of identical duration d a constant of common knowledge. Then d/D models the probability of the game coming to an end. We consider δ such that $d/D = 1 - \delta$ the probability that the game is still taking place. Hence, the game arrives to stage k with probability δ^k. At the start of each stage : Each player advertises its per packet price . We suppose that each AS is aware of the history of the game; that is after each stage all ASs are aware of proposed prices and consequent outcomes on the previous stages.

A powerful notion while analysing a repeated game is the *subgame perfect equilibrium*[6], where the played strategies represent a Nash equilibrium in each subgame. That is given any history of the game given by past plays, the adopted strategies still represent a Nash equilibrium trough the rest of the game. A set of strategies can be proved to induce a subgame perfect equilibrium if they satisfy the **one deviation principle**. This principle ensures that no player can increase its utility by deviating from its original strategy at a single stage. The intuition behind this principle is that improving the utility of a player supposes that at least at one stage the pay-off obtained by deviating is greater than the one in the original strategy. Thus, in order to prove that the set of strategies form a subgame perfect equilibrium, it is sufficient to prove that they satisfy the one deviation principle.

Now, let us analyse our game under these considerations on a simple scenario. First, we consider the simple case where there is a single communication in a network of 4 vertices, one source and one destination and two intermediate providers. We consider that providers have identical costs. There exists an analogy with the Bertrand game [3]. Bertrand game models interactions between duopoly firms that propose homogeneous products and compete only on price. The consumers buy all products from the cheaper firm or half at each when the price is equal. In the Bertrand game, firms are supposed to have the same marginal cost, when the customer demand is supposed to be linear in the price. A monopoly price p is given and represents the price that the firm will charge if she had the monopole on the market. In our simple scenario providers have identical costs when demand is constant. The monopole price is given by p_{max} since it is the maximum price that can be charged.

There are two possible outcomes in Bertrand competition : Both firms decide to not cooperate and price the only non-cooperative Nash equilibrium which is to charge the marginal cost c. Indeed for each price p_1 proposed by firm1, the best response of firm2 consists for every $p_1 > c$ in lowering slightly the price to win the market. The only one equilibrium is (c,c). Note that when marginal costs are different, the firm with lower marginal cost can win all the market. Otherwise, both firms can cooperate and charge the monopoly price p and thus share the market. Since BGP requires a single routing, splitting the traffic can be done on time by alternating $(p_{max}, p_{max}+1)$ announces. Hence, for instance player 1 wins over even stages and player 2 over odd ones. In order to make this threat credible a punishment should be added to avoid deviations. Thus, a possible strategy can be to alternate $(p_{max}, p_{max}+1)$ announces and if player 1 for instance deviates by playing $p' = p_{max} - 1$[1] in stage $2k$ then player 2 plays $p' - 1$ in stage $2k+1$. This strategy satisfies the one stage deviation principle for sufficiently patient players.

[1] p' should be lower than p_{max} in order to win the game but the higher possible to make the maximum benefit. Given the unit discretisation, that price is $p_{max} - 1$.

Indeed, suppose without loss of generality that player 1 decides to deviate by playing $p_{max}-1$ in an odd stage $2k$ then player 2 will play $p_{max}-1$ in the stage $2k+1$. Considering an only one stage deviation, the game will continue with the original strategy. Then pay offs will differ only in stage $2k$ and $2k+1$. In stage $2k$ player 1 wins with price $p-1$ and in stage $2k+1$ he will loose which give him a total benefit of $p_{max}-1$.[2] When in the original strategy he would get a total pay-off of p_{max}. The expected improvement is given by discounting the pay off by the probability of the game taking place at the corresponding stage. Hence the deviation is profitable iff: $(p_{max}-1)*\delta^{2k} > p_{max}*\delta^{2k+1}$ that is when $\delta \geq \frac{p_{max}-1}{p_{max}}$ the one deviation principle is satisfied for sufficiently patient players and the proposed strategy is then a subgame perfect equilibrium. Note, that this strategy is profitable to both providers but is not profitable to the source since it induces a flip-flop like routing. This behaviour can easily be generalized to the case of n providers with identical costs.

Let us consider now the case of one source and one destination with n possible intermediate providers having different costs. For the purpose of simplicity, let us assume that $cost_1 < cost_2 < \ldots < cost_n$ where $cost_i$ is the cost of provider i.[3] The utility of the provider is the difference between its price and its cost if he wins and 0 otherwise. Again, similarly with the Bertrand game the only one non cooperative equilibrium is when the lowest cost provider (here provider 1) takes all the market by proposing $cost_2-1$. Again the perspective of the game being repeated makes cooperation possible in order to maintain higher prices and strategies can be constructed with the same intuition to prevent deviations. However, given that costs are different, many cooperations are possible. For example, provider 1 can announce a price in $[cost_2..cost_3-1]$[4] in order to invite provider 2 to join him and share the market. Actually, they can cooperate by both setting the price at $cost_3-1$ the maximum price such that they can get all the market. In such a situation, we will talk about coalition and denote the set of providers joining it $coalition_2=\{1,2\}$. And so on, provider 1 can set a price in $[cost_{i-1}..cost_i-1]$ in order to invite provider $i-1$ to join him and share the market then forming $coalition_i=\{1..i\}$.

Obviously $coalition_i \subset coalition_j$ iff $i<j$ that is if a provider j can join a given coalition then every provider $i<j$ can do. Also, each provider i can only join a $coalition_j$ where $j>=i$. Note that if providers $i=1,\ldots,j$ decide to cooperate thus forming $coalition_j$ and given that they are all utility maximisers they should announce $cost_{j+1}-1$. Hence, we will talk about strategy of joining coalition j when the strategy consists on setting price equal to $cost_{j+1}-1$.

Now, the question that arises is which coalition will be chosen and would all providers necessary to form it have actually incentive to join it. Let denote s_i^j

2 Of course multiplied by the amount of traffic, but here we consider without loss of generality a unit of traffic.

3 When $cost_1 \leq cost_2 \leq \ldots \leq cost_n$, we obtain the same results, we need just to consider class of providers having the same cost.

4 Integer values in the corresponding interval.

the strategy of player i that consists in choosing to join $coalition_j$ where $j \geq i$. This can be done by provider i setting a price $p_i^j = cost_{j+1} - 1$.

The utility of a player i when choosing $coalition_j$ is given by:

$$u_i(s_i^j, s_{-i}) = \begin{cases} 0 & \text{if } \exists i', j' < j \text{ s.t } s_{i'} = s_{i'}^{j'} \\ \frac{p_i^j - cost_i}{|\{i'/s_{i'} = s_{i'}^j\}|} & \text{otherwise} \end{cases}$$

That is, the utility of provider i is 0 if another provider proposes a lower price, otherwise he shares the market with the other providers that have proposed the same price as him. Hence the utility of player i choosing a $coalition_j$ given that all other providers $i' \leq j$ have also join that coalition is $(p_i^j - cost_i)/j$. Each player is expected to choose the coalition that maximizes such utility. We will denote the corresponding strategy (price announced) s_i^*.

For instance for provider 1 $s_1^* = max\{cost_2 - cost_1 - 1, \ldots, \frac{cost_n - cost_1 - 1}{n-1}, \frac{p_{max} - cost_1}{n}\}$ and we denote $coalition_{j^*}$ the corresponding coalition. Now, the question is whether providers $\{2, \ldots, j^*\}$ will choose to join the same coalition. That is $coalition_{j^*}$ is the coalition that maximizes their utility too. The answer is given by the following theorem:

Theorem 1. *If $coalition_{j^*}$ is the coalition that maximizes the first provider utility then it maximizes providers $2, \ldots, j^*$ utilities:*

$$\forall i' \in \{2, \ldots, j^*\} \quad s_{i'}^* = s_{i'}^{j^*}$$

The proof skipped due to space limit can be found in the research report [2].

That is the lowest cost provider chooses his preferred coalition and the involved providers follow him. When different best coalitions are possible (with the same utility) a problem of coordination can arise. A dominant strategy for the lowest cost provider is to choose the lowest coalition and for the other members to follow him. When a player k deviates players $1, .., k-1$ punish him by playing according to $coalition_{k-1}$. For sufficiently patient players, this is a subgame perfect equilibrium. The intuition, is that the punishment will exclude the deviator from the coalition for the rest of the game.

Let us consider now the case where instead of direct connection, providers are connected via disjoint routes to the source. Let us denote i the direct provider connected to the destination, i' the corresponding provider connected to the source and l_i the length of the route between i and i'. Without loss of generality we consider $l_1 < l_2 < \ldots < l_n$. The benefit of a provider is the difference between the price at which he has bought the route (the price of his provider) and the price at which he proposes the route to his customer. The net benefit is obtained by subtracting the transit cost. For simplicity assume first that there are no transit costs. That is a customer is interested to buy a route if at least he can make a benefit of 1.

The game can be separated into two different games: the sequential game that each provider i plays with his predecessors on the route to s and the simultaneous game that players $i = 1 \ldots n$ are playing in order to fix their price.

This separation is possible and relevant only because paths are disjoint. Indeed on each route, players between i and i' including i' are completely dependent on i the owner of the route toward the destination. This game is known as the ultimatum game [6] where some value is to be divided between some players. A given player (called the first mover) proposes a division of the value and the others can only accept the division or refuse it inducing utility of 0 to everyone. The optimal strategy for the player proposing the division is to take the maximum portion and let to the others the minimum such they are still interested (in a continuous setting ϵ and in a discrete frame setting 1) which is an Nash equilibrium. The player who is proposing the division has an advantage because he is the first mover in the sequential game. In our case, if there is for example only one route, the provider 1 is the first mover because he announces a route to the destination first and he should propose the route at $p_{max} - l_1$ letting each of the other intermediate providers get a benefit of 1 (the route is proposed then to the source at p_{max}). The following providers will then accept since they prefer to get a benefit of 1 rather than to lose the market.

When there are several routes the provider i has to fix his price depending on the simultaneous game he is playing with the other direct providers. Providers directly connected to the destination ($\{1, \ldots, n\}$) have to take into account that each provider on the correspondent path should at least make a benefit of 1. Hence each of the n possible routes can be proposed to the source at least at $l_i + 1$ by each corresponding i'. The problem can be viewed then as n providers proposing to connect directly the source and the destination as in the former case with each provider i having a cost $l_i + 1$. The lowest cost provider who has the market power is the one on the shortest path. As we have argued above, he chooses his optimal coalition and the other involved providers follow him. Intermediate providers have to propose the price at which they have bought the route +1, otherwise their route will not be chosen by the source. Note that when internal transit costs are not null then we can obtain the same results by considering as metric the sum of transit costs. We give an example that help to understand how the situation can be different in the general case from the special cases.

Example 1. Let us consider the network given in Fig. 1. Suppose that the source has a $p_{max} = 8$. Provider 1 prefers to join $coalition_2$ since $(6/2 > 8/3)$ and provider 2 follows him. They will then propose $p_1 = 6$ and $p_2 = 5$. Provider 4

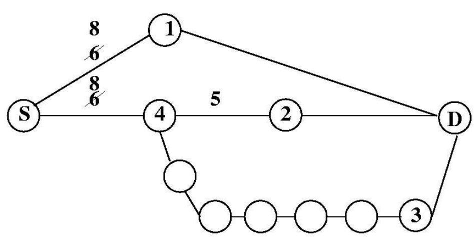

Fig. 1. Cases of non disjoint routes between the source and the destination

should announce 1 to insure that the second route will be chosen. His advantage is that he blocks the third route. he can propose a coalition to provider 1 where they can improve their benefit. The first and the second route still shares the market but at a higher price (p_{max}) as depicted in Fig 1. Provider 2 conserves his benefit and have non incentive to punish provider 1 or 4. Provider 4 acts as a stopper of the third route and thus can propose a second coalition that improve his benefit without decreasing benefit of provider 2.

This is the intuition we have used to propose an algorithm that computes prices in the general case [2]. It consists in computing successive coalitions to improve intermediary providers benefit while respecting precedent coalitions.

3 The Dynamic Distributed Game

In this section we try to analyse how the system behaves in a distributed framework. Indeed, in reality nodes have only a local view of the game including the topology and thus the nature and the length of the possible routes. We simulate the distributed game and investigate if some specific local strategies can lead to a similar results than the one expected by the theoretical analysis. For this purpose, we need to introduce first the distributed algorithmic model. The system is still modeled by a graph linking the source S and the destination $Dest$ with m nodes labeled $i \in \{1 \dots n \dots m\}$ where the direct nodes are $\{1 \dots n\}$. We denote $Successor(i)$ the set of possible providers of node i and $Predecessor(i)$ the set of possible customers of the node i. A Node i is characterized by the following variables: **Current price** per unit of traffic denoted p_i which is announced by the node i to his neighbours in $Predecessor(i)$, **Current Provider** denoted $provider(i)$ which is one of node's neighbours that can reach the destination. It is the one who proposes the best price to i. For $j = provider(i)$ we have $p_j \leq p_k \; \forall k \in Successor(i)$, **State** denoted $state(i)$ that indicates whether the node is crossed by the transit traffic (O) or not (N). That means that the node belongs to the chosen route. For the special case of the source the state indicates whether the source has received at least an acceptable route (its price is lower than p_{max}) or not. We define the set $Customer(j) = \{k/provider(k) = j\}$. We have $p_i > p_{provider(i)} \; \forall i$ that is a node proposes a route at a price higher than the price at which he has bought it. Every node chooses the provider that proposes the lowest route price. If a node chooses a provider j and thereafter it receives from a provider k a proposal of a route with lower price it switches toward k.

In a distributed setting, each node is informed of all the variables of its neighbours using traffic control. However all routes may not be visible at every node: the set of routes learned at one node depends on route selection at its provider. Node's state depends on the route chosen by the source. At the beginning each node's state is equal to N because routes are not established yet. Node's state is updated when he receives a state update message from its neighbours as following: state(i)= O if $\exists j \in Customer(i)$ such that $state(j) = O$ or if $i = S$ and $p_{provider(S)} \leq p_{max}$. and N otherwise. Hence when the source chooses an acceptable route, its state changes to O and then it sends an update message to its provider who in turns changes his state and so on until the destination. When the source switches on a new received route with a better price, the state of nodes on the new route is updated iteratively into O when the state of the nodes on the old route is updated iteratively into N.

We propose to test a simple strategy that all the nodes can use to update their price depending on their state: if $state(i) = O$ then $p_i \longleftarrow p_i + 1$ otherwise $state(i) = N$ and then if $(p_i - p_{provider(i)}) > 1$ then $p_i \longleftarrow p_i - 1$.

The intuition behind this strategy is that providers with no transit traffic decrease their prices in order to attract the traffic. Each provider accepts to transit the traffic if he has at least a benefit of 1 and does not decrease its price under this limit. When a provider gets the transit traffic, he tries to increase his price in order to reach the maximum possible benefit.

4 Simulation Analysis

Our objective here is to study the stabilizing behaviour of the distributed system under the above price adjustment strategy and whether it matches the theoretical results. Simulation is done using OMNET [7] simulator. We consider different topologies (Fig. 2). Links have the same propagation delay equal to 0.31 ms. Neither queuing nor scheduling delays are considered in the simulation. Node state messages are generated automatically when the node state is updated and are sent as traffic control messages. In our simulation, the stage game duration is $d = 50ms$. We implement the simple price adjustment strategy explained above and consider different scenarios:

Scenario 1. We consider topology 2 and simulate the price adjustment strategy when transit prices starts from a high price (chosen $> p_{max} = 20$ as depicted in Fig. 3 at the left side. Then prices are adjusted until t= 150 ms (stage 4) where routes proposed to the source become acceptable. Both routes share the market but at a higher price. Direct providers have an advantage over intermediate ones, the first one taking the maximum benefit.

When a direct provider chooses to start at a price lower than p_{max} as in the scenario depicted in Fig. 3 at the right side then his route is selected during

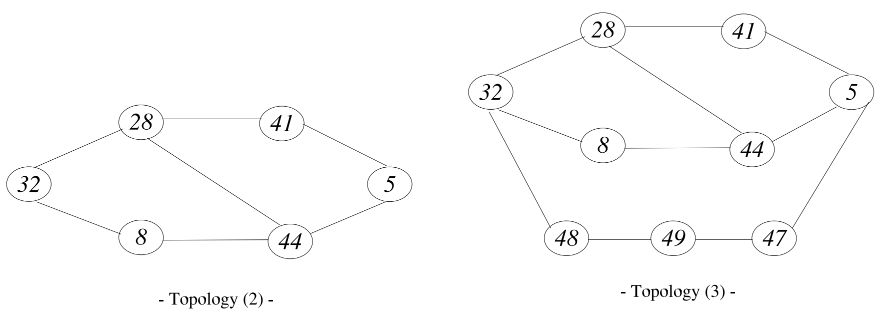

Fig. 2. The different simulated topologies

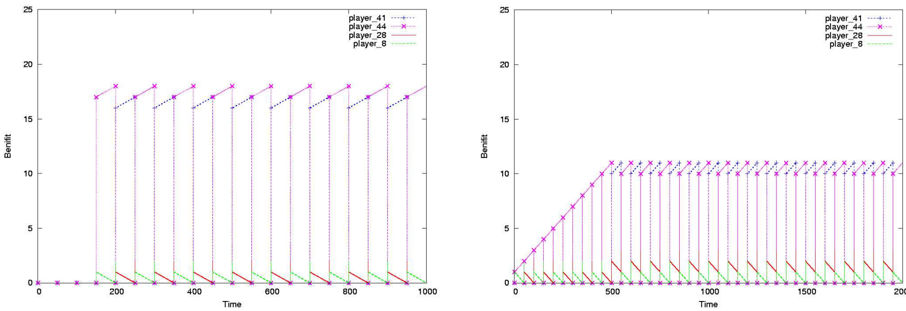

Fig. 3. Providers benefit in scenario 1

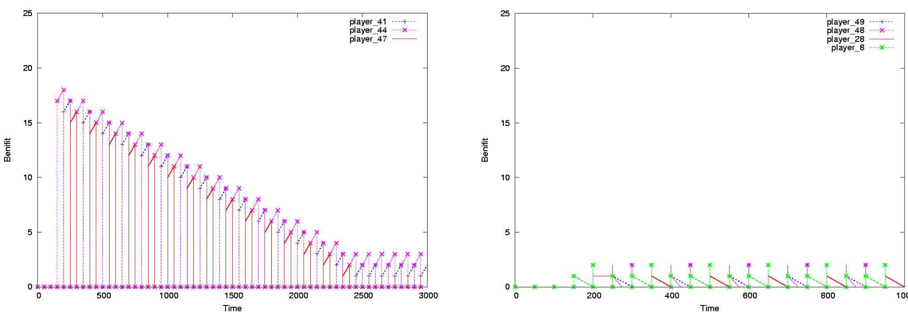

Fig. 4. Direct providers and intermediary providers benefits in scenario 2

few steps. Prices are then adjusted until a situation where both routes share the market. Note that provider 41 and 44 would have better benefit with the precedent scenario where they started both at p_{max}. In summary, this simple strategy gives similar results to those expected by the theoretical analysis using only local information. Indeed, when providers start from high prices they can share the market while maintaining higher prices than when they do not cooperate.

Scenario 2. We consider now topology 3, where three different direct providers compete for the market. We simulate the price adjustment strategy where all transit prices starting from $p_{max} = 20$. Direct providers benefits and intermediate providers benefits are depicted in Fig 4. As with topology 1, prices are adjusted until step = 150 ms (stage 4) where routes proposed to the source become acceptable. However, direct providers do not succeed in maintaining high prices. This can be explained by the way prices are updated. Indeed, the three routes share the market but each provider lowers it price at least two times (when the other routes are chosen) but increases only one time its price. This leads the prices to decrease drastically. We need a more elaborated strategy in order to obtain a behaviour which is similar to the theoretical expected behaviour.

5 Conclusion

We present a combined game theoretic and distributed algorithmic approach to the transit price negotiation problem. We highlight situations where cooperation is possible in order to maintain higher prices. However such situations lead to a flip flop routing. An interesting issue is to investigate how the source can avoid such behaviour for example by adding some penalties when its provider changes its price. A more elaborated local strategies are currently tested mainly based on stochastic learning of optimal strategies. Finally we are investigating how to generalize the proposed approach to a network where there are different sources and destinations for the traffic while considering coherent routing.

References

1. M. Afergan. Using repeated games to design incentive-based routing systems. In 25th Conference on Computer Communications, IEEE INFOCOM, 2006.
2. D. Barth, J. Cohen, L. Echabbi, and C. Hamlaoui. Transit price negotiation : a repeated game approach. In PRiSM Report 2006/76, www.prism.uvsq.fr, 2006.
3. J. Bertrand. Theorie mathematique de la richesse sociale,. Journal des Savants, 67:499 – 508, 1883.
4. J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In Proceedings of the 2002 ACM Symposium on Principles of Distributed Computing., 2002.
5. R. La and V. Anantharam. Optimal routing control: Repeated game approach. IEEE Transactions on Automatic Control, 47(3):437 – 450, 2002.
6. Martin J. Osborne. An Introduction to Game Theory. Oxford University Press, USA, 2003.
7. A. Varga. The omnet++ discrete event simulation system. In In European Simulation Multiconference, 2001.