Policy-Based Pricing for Heterogeneous Wireless Access Networks

Javier Baliosian, Joan Serrat, Matias Richart, Juan Saavedra, Mariela Borba, Jose Melus

Policy-based Pricing for Heterogeneous Wireless Access Networks

Javier Baliosian[1], Joan Serrat[2], Matias Richart[1], Juan Saavedra[1], Mariela Borba[1], and Jose Luis Melus[2]

[1] School of Engineering, University of the Republic, Montevideo, Uruguay, javierba@fing.edu.uy,
[2] Polytechnic University of Catalonia (UPC), Barcelona, Spain, serrat@tsc.upc.edu

Abstract. Our cities are already covered by a myriad of diverse wireless access networks. The most ubiquitous access networks are the well organized homogeneous and centralized operator-based cellular networks that sustain their business model on a captive client basis. However, a new billing paradigm is rising, where a client can choose to connect to the provider that best comply with his/her current requirements and context. Inside this paradigm, this paper presents a distributed, rule-based pricing strategy aimed to improve the quality of service and to increase the global income of a service provider. The performance and reliability of the rule-based decisions is supported by a Finite State Transducers-based inference machine specially designed to manage networking systems. We show, with simulations, that, using our strategy, the operators can make the new billing paradigm profitable while the clients benefit from the economic advantages of competition and of the quality given by a pricing–based network balance mechanism.

1 Introduction

Our cities are already covered by a myriad of diverse wireless access networks. The most ubiquitous are the well organized homogeneous and centralized operator-based cellular networks that base their business on a captive client model. However, a more heterogeneous and free network of WiFi, WiMAX and 3G access-providers has appeared and it is starting to challenge the dominant billing paradigm. In this new world, clients can choose the best or the cheapest provider for the time and place they are trying to connect from, setting the grounds of a future free market of Internet connectivity-providers. In this context, pricing will have a fundamental role. Using the right pricing strategy, an operator will try to obtain the highest possible revenue while the users will try to get a service that fits their requirements at the minimum possible price. As stated in [1] "From an economic point of view, pricing plays an important role in trading any resource or service. The most important objective of trading is to provide benefits to both the sellers and the buyers. Therefore, the price must be chosen so that the revenue of the sellers is maximized while the highest satisfaction is achieved by the buyers. There are two main factors influencing the price setting, namely, user demand and competition among service providers. Price and demand are functions of each

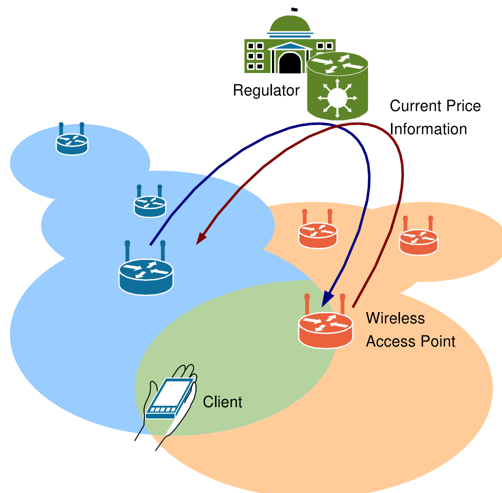

Fig. 1. Scenario

other." Following these ideas, if demand is high the service provider can charge a high price in order to earn more. However, if demand is low, the price must be reduced to attract more mobile users. Competition between service providers also impacts on the price of the service. Typically, if the services are substitutable (even though different), users buy a service that provides the highest satisfaction at the lowest price.

As explained by Courcoubetis and Weber in [2] the communication services that we are considering in this work can be considered simply as means for the transport of data with a given quality, characterized as a certain error rate, delay and jitter. Obviously, the network access providers will want to recover their investment charging a price for their services, but that is not the only reason for which pricing is important. The price of simple goods it is often determined by only one parameter such as the number of copies, their weight or the length of a lease. However, communication services are specified by several parameters such as peak rate, average bandwidth or loss rate. Moreover, multimedia service contracts are specified with additional parameters such as tolerance to bursts and adaptability to network changes. Since connectivity services can be specified in terms of so many variables, the number of different possible contracts is enormous and complicates the design of a reasonable and coherent pricing strategy. On the other hand, contracts are more than simple price agreements. For example, a contract may be an incentive for the user to produce traffic conforming to the agreed parameters. This, at time, will impact positively on the service quality and the price paid by the clients in general. All this motivates an effort to develop a pricing technique simple enough to be implemented by the operators, but, at the same time, sophisticated enough to compete successfully with other strategies and work as a scalable feedback mechanism to control how the network is used. A provider can reduce the price of a service during off-peak hours to incentive the use of idle resources or charge an extra price to a user that exceeds the agreed traffic. Additionally, pricing may be seen as an alternative to TCP for congestion control. Thus, in a similar manner to TCP and its signaling, a higher price may induce users to reduce the packet transmission rate or to stop it completely. However, to be useful as a congestion-control mechanism, a pricing technique must be very dynamic. It should be able to change the price of a particular service in real time and in a particular region of the network, for example in a particular access point that is suffering of congestion.

This paper presents a distributed, rule-based pricing system that implements exactly the same intuitive ideas in the shape of policy-rules to be enforced on the price charged by each provider. Those rules are aimed to improve the quality of service and to increase the global income of a service provider in that world in which users are free to choose every time they connect.

The paper is organized as follows: in §2 we introduce the rationale of our solution and and overview of its architecture. In section 2.2 and 2.3 we present the set of rules governing our system and the kind of conflicts that arise using them. In §3 we present details of the system design and the manner in which policy conflicts are solved. Finally, in §4 we depict a simulation-based evaluation and we conclude and present some future work in §5.

2 A Policy-based Pricing Solution

The rationale behind the proposal of this paper is twofold: first, it tries to offer an autonomic, scalable pricing system that provides the operators with simple and business-understandable means to set prices. In the examples presented in this paper, we follow the principles of demand and competition mentioned above. Second, our proposal is aimed to demonstrate that the process of solving the conflicts between the pricing policies constitutes an optimization procedure that is capable of, for example, balance an access network using the price of the access service as an incentive-based tool to drive the users choices. The following sections describe the overall scenario and the design of the proposed system.

2.1 Overall System Architecture

We envisage scenarios where three types of actors coexist on a geographical area: users, providers and a regulator (see Figure 1). Users are persons in possession of some wireless-capable device that are willing to establish a connection to the Internet with some quality of service requirements and at the minimum possible cost. As we envision it, users are not attached to a provider by a contract. Instead, they pay for the provider's service on demand, by, for example, a credit card or pre-paid means, as it is common in current cellular 3G services. In our model, in front of equal or similar services, users will always choose the services of less expensive provider. The providers have an access network conformed for a set of access devices that we will call generically Access Points (AP). A provider's objective is to sell access services to users while maximizing their revenue. The third entity, the regulator, is a neutral entity, probably played by a governmental agency in a real setup, with the objective of enforcing the sharing of pricing information between the providers to allow pure market competition.

In our scenario, there are two different providers offering their connectivity services through a network of APs over the overlapping geographical regions. This means that at any point of this region, a user has network coverage from one or more APs of each provider. In the particular world we are envisioning, each AP has its own price for a connectivity service of a given bandwidth and duration. Those prices may vary from one AP to another even if they are operated by the same provider and from time to time. In this manner, an AP in a popular part of the city may have a higher price than another in a neighborhood with few users or the same AP can diminish the service price during the low-usage parts of the day. The mobile user's terminal can connect to any of those APs and run an agent that is able to communicate with an AP and get information about the current price for a certain service the AP provides (later we will discuss about the technical solutions for this issue). Once it has the price from different APs, the agent connects to the AP with the lowest price, with or without human assistance (another possible criteria will be discussed bellow in §5). The established connection will last for the contracted time or until the connection is lost, for example, because the user moved out of the AP's coverage region. For this paper, the manner in which the regulator collects and shares price information between APs is not relevant but it may be made with a standard publish-subscribe mechanism with some integrity and no-repudiation guaranties.

2.2 Governing Policies

Policies are intended to allow each AP to decide which is the most advantageous price for its own connectivity service, having into account its context, user's demand and potential competitors. These decisions are made by a Policy Decision Point [3] installed in the same AP, independently from the others regardless whether they belong to the same or competing providers, following a set of policy-rules modeling the economic criteria of demand and competition mentioned in §1. It is worthy to mention that these sets of policies can be freely established by each provider according to best practices founded on past experience, forecasts based on economic models, etc and that by no means are imposed by the regulator. Even more, a provider could participate without making use of pricing policies, just fixing a flat rate for its services. Nevertheless, for the sake of illustrating the approach, and also to allow a quantitative evaluation we present hereafter specific sets of policies; one concerning demand and the other competition.

The set of rules driving the decisions regarding demand is presented in Table 1. The rationale behind this set of rules is simple; the price of the service is increased, kept constant or decreased depending on the number of users served and its gradient of change. In this way, the price will be adapted to stimulate or inhibit service demand and its adaptation rate will track the evolution of such demand. In practice we accomplish it by classifying the number of users in three categories (few, mid and lots), the gradient of change in two (slow and fast) and also allowing two rates of price change (slow and fast).

On the other hand, the set of rules facing the competition is presented in Table 2. Here the global objective is to accommodate the price to the evolution of the competitors to avoid users' migration. In particular we have classified the price of the competition in two categories (lower and higher) and considered two adaptation rates (slow and fast). In this context, a competitor is an individual AP and two APs are competitors between them when their coverage regions overlap and they belong to different providers.

Rule 1:	**if** few_users and users_steady **then** decrease_price_slow
Rule 2:	**if** few_users and users_decreasing_slow **then** decrease_price_slow
Rule 3:	**if** few_users and users_decreasing_fast **then** decrease_price_fast
Rule 4:	**if** few_users and users_increasing_slow **then** keep_price
Rule 5:	**if** few_users and users_increasing_fast **then** keep_price
Rule 6:	**if** mid_users and users_steady **then** increase_price_slow
Rule 7:	**if** mid_users and users_decreasing_slow **then** decrease_price_slow
Rule 8:	**if** mid_users and users_decreasing_fast **then** decrease_price_fast
Rule 9:	**if** mid_users and users_increasing_slow **then** keep_price
Rule 10:	**if** mid_users and users_increasing_fast **then** increase_price_slow
Rule 11:	**if** lots_users and users_steady **then** keep_price
Rule 12:	**if** lots_users and users_decreasing_slow **then** decrease_price_slow
Rule 13:	**if** lots_users and users_decreasing_fast **then** decrease_price_fast
Rule 14:	**if** lots_users and users_increasing_slow **then** increase_price_slow
Rule 15:	**if** lots_users and users_increasing_fast **then** increase_price_fast

Table 1. Set of rules driving an access point behavior on demand bases

Rule 16:	**if** competitor_price_lower and competitor_price_decreasing_slow **then** decrease_price_slow
Rule 17:	**if** competitor_price_lower and competitor_price_decreasing_fast **then** decrease_price_fast
Rule 18:	**if** competitor_price_lower and competitor_price_steady **then** decrease_price_fast
Rule 19:	**if** competitor_price_lower and competitor_price_increasing_slow **then** decrease_price_slow
Rule 20:	**if** competitor_price_lower and competitor_price_increasing_fast **then** decrease_price_fast
Rule 21:	**if** competitor_price_higher and competitor_price_decreasing_fast **then** decrease_price_slow
Rule 22:	**if** competitor_price_higher and competitor_price_steady **then** increase_price_slow
Rule 23:	**if** competitor_price_higher and competitor_price_increasing_slow **then** increase_price_fast
Rule 24:	**if** competitor_price_higher and competitor_price_increasing_fast **then** increase_price_fast

Table 2. Set of rules driving an access point behavior on competitors' bases

2.3 Conflicting Policies

Conflicts, or logic contradictions, are an intrinsic phenomenon of the policy-based management. The process to solve those conflicts is the way that policy-based network management has so as to optimize the configuration of the services and devices of a communications network.

In the particular case of the this work, conflicts arise mainly because the diversity of objectives of our system. For example, lets consider the following simple –and quite intuitive– version of the rules presented above:

Rule 1: `if few_users then decrease_price_slow`

Rule 2: `if competitor_price_higher then increase_price_fast`

Rule 1 models a pricing policy based on demand, and Rule 2 models a pricing policy based on competition. Although in a simple manner, they are the most patent case in which our previous set of rules may present a conflict: which action has the system to trigger if there are a small number of users connected to an AP but the price of the competitor is higher than the local price?

The process that our system uses to decide in front of this kind of situations is described bellow in §3.1.

3 A Policy-based System Design

Our solution combines centralized, hierarchical and fully-distributed management structures to address different challenges with the most appropriate approach. Our basic de-

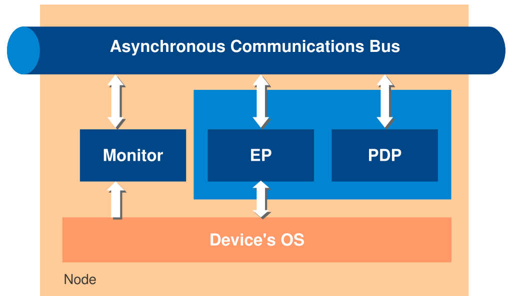

Fig. 2. Architecture of a Node

sign criteria is to maximize the distribution of tasks over the nodes as much as possible. Only the tasks that inherently require a centralized organization, such as global optimizations, or those tasks which perform better on a weakly-distributed structure, are carried out using hierarchical structures. Management tasks such as local optimizations follow a fully distributed approach.

The system can be seen as mainly composed by independent units or nodes that work as peers deployed at each managed network device (a wireless access point). From a high level architectural point of view those nodes follow a classical policy-based architecture constituted by a Policy Decision Point (PDP), a Policy Enforcement Point (PEP) and an agent that monitors the state of the device and the behavior of the network. Policies follow the Event Condition Action (ECA) model, where the condition may be the occurrence of some event (e.g., an alarm or a service request) or certain network state, and the action is the desired response when the condition is true or just it occurs. The monitoring agent of any of the nodes of the network generates notifications, which are the source of events for the PDP of the same or any other node. All the communications between nodes, including control messages, notifications and request for actions are carried out by a content-routed asynchronous communication bus that organizes the nodes in a hierarchical overlay. The communications bus' design is inspired by existent notification services such as Siena [4]. Finally, a centralized management station is in charge of the edition, optimization and distribution of rules into the nodes.

The reconfiguration actions decided by the PDP may also involve other network elements, in which case the node may be the manager of the activity and delegate responsibility for some actions to other nodes. As it will be described in the next section, policies have to be managed as finite state transducers (FST). The translation from policy rules expressed in natural language into the FSTs implementing those rules is currently made by programmer assisted algorithms presented in [5]. Despite the fact of being a software assisted process, it is still a complex and work-intensive task.

3.1 The Logic behind the Pricing Management System. Finite State Transducers for Conflict Resolution

To better understand how the FST-based conflict-resolution process works, we will describe it with the help of the rules presented above in §2.3.

As stated above, Rule 1 models a pricing policy based on demand, and Rule 2 models a pricing policy based on competition. The FST in Figure 3(a) models both rules at once.

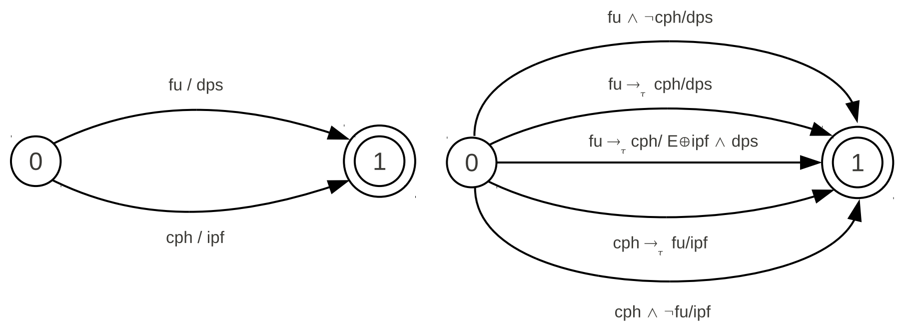

(a) TFFST Modeling two Conflicting Pricing Rules

(b) TFFST Modeling two Pricing Rules with their Conflicts Solved

Fig. 3. Finite State Transducers for Conflict Resolution

In the figure, the TF `few_users` is depicted as `fu`, the TF `competitor_price_higher` as `cph`, action `decrease_price_slow` is depicted as `dps` and `increase_price_fast` as `ipf`.

For simplicity, we are considering the case in which a single notification arrives, then the PDP makes a decision based on it, and nothing else happens. (This makes the FST simpler because only has to accept strings made of a single symbol).

Now, consider the case in which a notification arrives stating that the price of a competitor has changed. At that time the PDP evaluates the FST in Figure 3(a) as follows: starting in the initial state "0" it evaluates the TFs on the input label of each of the state's outgoing edges. This is, the PDP evaluates `few_users` and `competitor_price_higher`. In the case that there are few users connected to the local AP and that the price of the competitor is lower or equal to the price charged by the local node, `few_users` is a positive number and `competitor_price_higher` a negative number. Therefore, the PDP chooses the upper edge of the FST, which produces `decrease_price_slow` as an output, causing local EP to decrease the price slowly. However, in the case that the price of the competitor is higher than the local price, we have two edges going out of state "0" with positive TFs and not a straightforward choice to be made. To solve this problem, before being deployed to the AP's PDP, the FST is determinised and the FST in Figure 3(b) is computed using the algorithm mentioned before. Now, when the PDP evaluates the determinised transducer, at each state it has only one possible edge to follow. Let us assume that, when the PDP evaluates it, the TF `few_users` takes a value of 0.3 and `competitor_price_higher` a value of 0.5. We are in the case in which `few_users` is "tauter than" `competitor_price_higher`, expressed in Figure 3(b) as $fu \rightarrow_\lambda cph$, therefore, the transducer produces a single output, in this case the order to decrease the price charged by the local AP slowly (action decrease_price_slow).

4 Evaluation

The simulation platform presented hereafter was designed with the aim to evaluate the viability of the proposed solution and, above all, the viability of the new pricing paradigm. In this section we present the result of two sets of experiments in which we compare the behavior of our proposal against the most popular pricing strategy for WiFi and 3G access networks.

4.1 Evaluation Setup

In order to evaluate our solution we simulated an environment with two competing providers offering access services on the same geographical region and with identical networking infrastructure resources. However, one of the providers follows our pricing strategy and the other follows the prevailing fixed-rate pricing structure presented in §2.2. Although it will not happen in a real deployment, to ensure a fair comparison between competitors in our experiments, both providers have identical sets of APs installed on the same coordinates and with identical coverage areas. A large number of clients move along the area covered by both providers choosing to connect to the AP with the lower price for a randomly specified service. To perform the experiments in this section we developed a simple discrete events simulator based on the JavaSimulation library.

For each experiment, each provider has nine APs deployed in a 1200x1200m region. Those APs are distributed in a mesh at approximately 400m of each other (see Figure 4). As said before, APs are by pairs, one owned by the provider following our strategy and the other owned by the competitor provider following a fixed-rate strategy. The price charged by each AP is a random variable that follows a uniform distribution with values between 50 and 60 units. Each AP has an available bandwidth of 3.5 Mbps, and a homogeneous circular coverage area with a radius of 285m. For the experiments, there are 450 mobile users. Each user tries to hire a connectivity service with a random bandwidth of between 50 and 150 Kbps with any of the APs in its connectivity range. These users move inside the square area mentioned above following one of two mobility models. In the first set of experiments we assume an urban environment and use the Manhattan mobility model [6] to capture the movements of the users. Following this model, each user moves along a mesh of parallel, perpendicular streets placed every 100m. For all the experiments using the Manhattan model, the speed of a terminal varies randomly with a pedestrian mean of 1m/s and a standard deviation of 0.2. At each intersection, a terminal decides whether to turn left, right or keep the direction with a probability of 0.2.

The second set of experiments were made to study the balancing capabilities of our strategy assuming that the clients move following the Reference Point Group Mobility (RPGM) model [7]. The RPGM models a set of clients moving in groups but with some individual freedom to move inside and between groups. Each group has a logic center (group leader) that determines the behavior of the rest of the group. At the beginning of the simulation each client is uniformly distributed in the vicinity of the leader, then, each client moves with random speed and direction. All the experiments run during a simulated time of 120 minutes.

The metric used to compare the behavior of both providers is the total amount of money made by each provider (adding the money made by all its APs) within the simulation time.

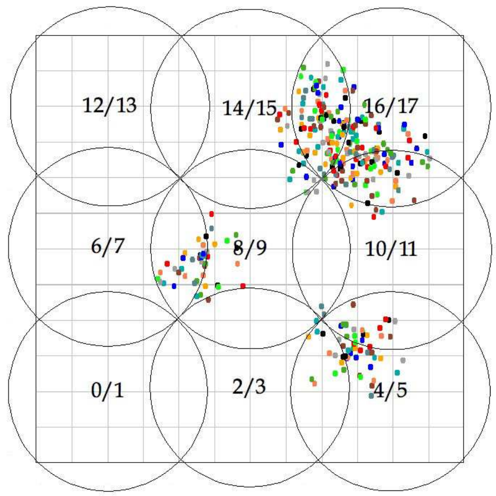

Fig. 4. Position of APs and Users (RPGM Mobility)

4.2 Evaluation results

Figure 5 shows the price evolution for the 9 APs of the provider implementing our pricing strategy. In this experiment the users move following the Manhattan model. In the figure, it is possible to see how the prices change getting similar to the competitor price but slightly below. This makes the clients to prefer our provider. The higher number of clients connected to APs owned by our provider explains the bigger revenue despite charging lower prices for the same service.

The Reference Point Group Mobility left regions of the geography without users and regions with high concentrations of them. This situation makes more visible the price dependence on the number of users trying to connect to an AP. In this experiment, APs 0 and 12 have no users in their coverage regions (see Figure 4), consequently, the price charged by those APs decreases until it reaches the minimum possible cost defined by the provider. The APs 2, 6, and 8 have a small number of users in range, therefore, they moderately reduce the price. Finally, those APs with most clients in their coverage area maintain their price close to the price charged by the competitor provider. This divergence between the prices charges by different APs of the same provider causes that some clients migrate from their closest APs to an AP with a lower price. In this manner, the pricing strategy works as a network balancing tool. We can see that the global earnings made by our provider are also higher than the earnings of the competitor (see Figure 8).

This experiment also shows how rule-based pricing can be used as a congestion control tool. As more users start connecting to a given AP, the price rises following for example Rule 15, naturally pushing some users out to other APs. Finally, an accessory, hard access-control prevents the connection of too many users to an AP in case the pricing mechanism is not enough.

The stability of the connections is related with the users' movement and velocity. When a user that is connected to a given AP is informed of the existence of a lower price AP in its current position, it terminates its current connection and tries to connect

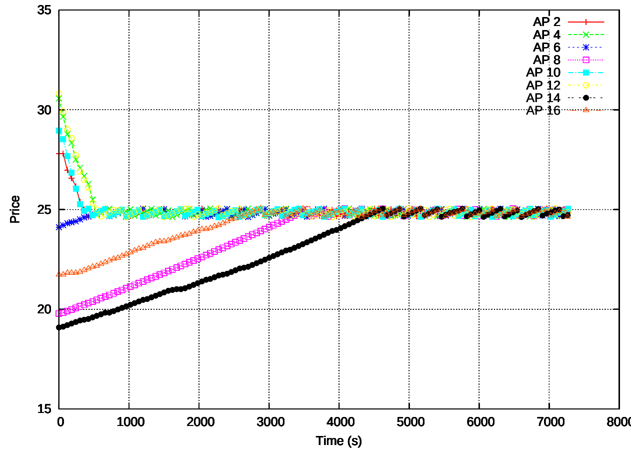

Fig. 5. Prices by AP for the Provider using our Strategy (Manhattan Mobility)

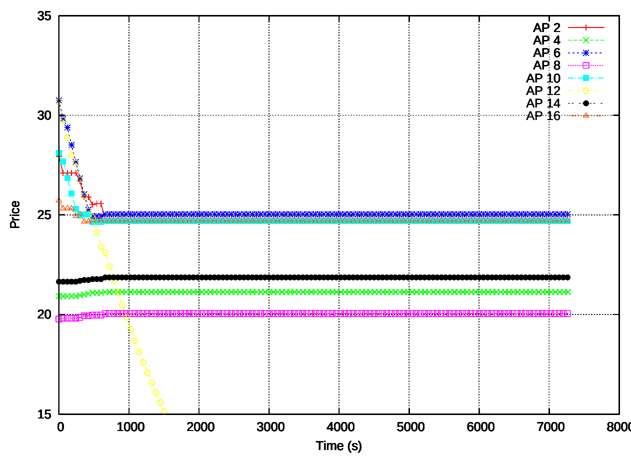

Fig. 6. Prices by AP for the Provider using our Strategy (RPGM Mobility)

to the AP with the lower price. However, a high percentage of users are connected along the simulation to any of the AP with the only exceptions of the intervals in which the users are switching from one AP to another.

5 Conclusions And Future Work

In this work we present a distributed, rule-based pricing system that improves the quality of service –by means of a smart access control and dynamic network balance– while increasing the global income of a service provider. This pricing mechanism assumes a business model, not yet existing today, where users can freely select any service provider with radio coverage at their location. However, we believe that this work supports the thesis of the plausibility of such a model. The depicted simulations encourage the idea that operators can make the new billing paradigm profitable while the clients

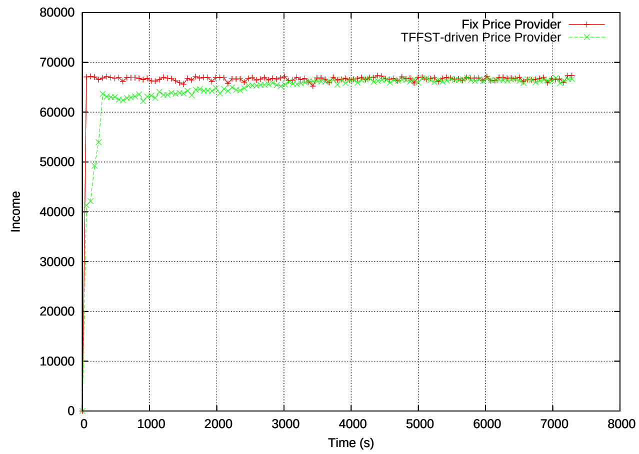

Fig. 7. Revenue by Provider. (Manhattan Mobility)

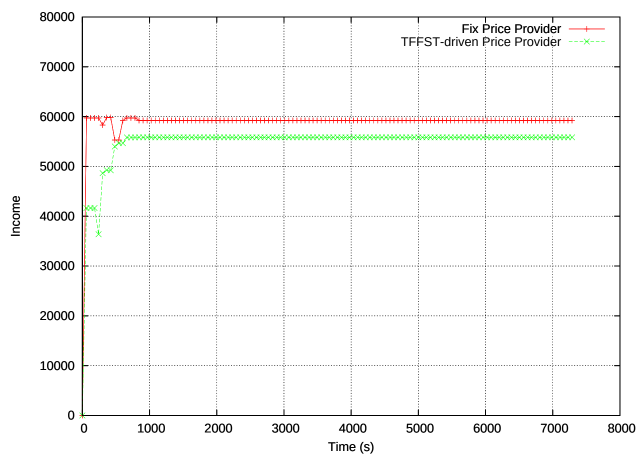

Fig. 8. Revenue by Provider. (RPG Mobility)

benefit from the economic advantages of competition and of the quality given by a pricing–based network balance mechanism. First, a solution must be designed to let the mobile nodes and the access point to communicate the price between the AP and the mobile node, even if there is no established connection between them. There are several possible solutions for this issue, for example, in the case of WiFi APs the price information can be transmitted as part of the AP beacon, we have made experiments that support this solution, however, using this method it is at least difficult to offer different prices to different clients but other solutions may be envisioned. Another issue to address is the stability of the solution. As can be seen in the graphs, even if the dynamic price of the rule-based APs converges to a region near an optimum, it keeps fluctuating above and below it. This issue can be alleviated with a proper study on the functions representing fuzzy concepts such as "decrease price slowly´´. This fine kind of tuning is a whole research issue that will be part of our future work. Finally,

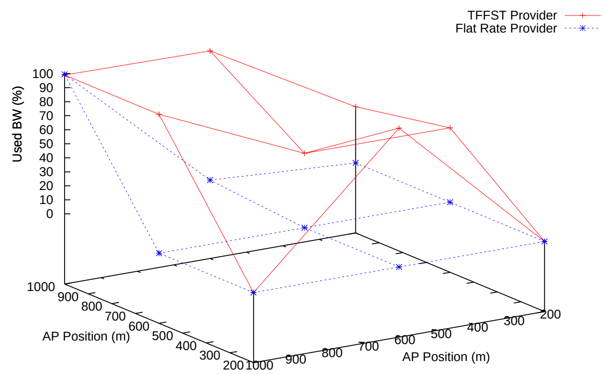

Fig. 9. Distribution of BW usage per Provider. (RPGM Mobility)

another future work issue is the inclusion of additional AP selection criteria such as signal strength, preferred provider or handover reduction.

Acknowledgments

This work has been done in the context of the MICINN project TEC2009-14598-C02-02.

References

1. Niyato, D. and Hossain, E. "Competitive Pricing in Heterogeneous Wireless Access Networks: Issues and Approaches", IEEE Network, Volumen: 22, Issue: 6, pages: 4-11; November-December 2008.
2. Costas Courcoubetis, Richard Weber "Pricing Communication Networks: Economics, Technology and Modelling" Copyright © 2003 John Wiley & Sons, Ltd. ISBN: 0-470-85130-9
3. B. Moore, E. Ellesson, J. Strassner, and A. Westerinen, "Policy Core Information Model – Version 1 Specification," RFC 3060 (Proposed Standard), Feb. 2001, updated by RFC 3460. [Online]. Available: http://www.ietf.org/rfc/rfc3060.txt
4. Caporuscio, A. Carzaniga, and A. L. Wolf, "Design and evaluation of a support service for mobile, wireless publish/subscribe applications," Department of Computer Science, University of Colorado, Tech. Rep. CU-CS-944-03, Jan. 2003. [Online]. Available: http://www.inf.unisi.ch/carzaniga/papers/
5. J. Baliosian and J. Serrat, " Finite State Transducers for Policy Evaluation and Conflict Resolution," in Proceedings of the Fifth IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2004), June 2004, pp. 250–259.
6. ETSI, Universal Mobile Telecommunication System (UMTS), "Selection procedures for the choice of radio transmission technologies of the umts," UMTS 30.03 Version 3.2.0, 1998-04. [Online]. Available: http://www.3gpp.org/ftp/Specs/html-info/3003U.htm
7. X. Hong, M. Gerla, G. Pei, C.-C. Chiang, "A Group Mobility Model for Ad Hoc Wireless Networks", ACM International Workshop on Modeling, Analysis, and Simulation of Wireless and Mobile System (MSWiM), 1999.