

Peer-to-Peer Wireless LAN Consortia: Economic Modeling and Architecture

Panayotis Antoniadis, Costas Courcoubetis, Elias Efstathiou, George Polyzos

Department of Computer Science, Athens University of Economics and Business

Ben Strulo

BTextact Technologies

{antoniad, courcou, efstath, polyzos}@aueb.gr, ben.strulo@bt.com

Abstract—In this paper we explore the incentive and architectural issues that arise in Consortia of Peer-to-Peer Wireless Local Area Networks. A P2P WLAN Consortium (PWC) is a community of WLAN Administrative Domains (ADs) that offer network access to each other’s registered users. The great benefit of ubiquitous access that these roaming members enjoy compensates for their AD’s cost of providing access to visitors. Existing roaming schemes utilize central authorities or bilateral contracts to control the parties’ behavior. In contrast, a PWC forms a P2P community in which participating ADs are autonomous entities. They make independent decisions concerning the amount of resources (e.g. access bandwidth) they contribute. As a result, similarly to existing P2P systems, a PWC will suffer from abusive behavior (free riding) if no incentive mechanisms exist to ensure that ADs offer the amount of resources that is economically justified. We explore the use of flexible rules on reciprocity to guide domain policies. We develop a suitable economic model and derive rules that would bring the system to a near-optimum equilibrium by forcing peers to contribute in order to consume. We then discuss certain key implementation issues related to the selection of critical parameters and rules enforcement in such a distributed environment.

I. INTRODUCTION

Internet services such as e-mail and the Web are, for many, more valuable than the telephone. However, Internet access is still nowhere near as ubiquitous as access to the telephone network. On the other hand many portable devices, such as smart-phones, palmtops, and tablet computers are becoming perfectly capable of handling end-to-end Internet protocols and applications. The users of these portable devices would greatly benefit from Internet access that is wireless, always-on, ubiquitous, high-speed and cost-effective. However, deploying infrastructure with wide coverage to support this is a non-trivial task.

Wireless Local Area Networks (WLANs) are an important component of this infrastructure in-the-making. Specifically, the IEEE 802.11 WLAN standard [14] has grown steadily in popularity since its inception and is now

well positioned to complement much more complex and costly technologies such as 3G (at least in metropolitan areas). This is already happening. WLAN signals already pervade many cities and WLAN cells frequently cover greater areas than was intended with their installation.

Fortunately, this fact, combined with how easy it is to gain access to a WLAN has drawn much attention from security experts and network administrators [11]. Successful WLAN vendors now proudly advertise the fact that their equipment can deny access to unauthorized users. However, artificially limiting network coverage or broadly denying network access is not helping to create a ubiquitous communications environment.

What we propose is an access network infrastructure that can incorporate existing and future WLANs and whose main function is to allow flexible access without compromising security. We name this infrastructure a *Peer-to-Peer WLAN Consortium* (PWC). Simply put, a PWC is a community of peer WLAN Administrative Domains (ADs) that offer network access to each other’s registered users. Users roaming to other ADs within the Consortium can enjoy various services such as Internet access, intranet services and other higher level services, thus benefiting from the community formed and, hopefully, compensating for their AD’s cost of providing similar access and services to visiting members of other domains.

Existing alternative schemes like Wireless ISP (WISP) associations, e.g. *Pass-One* [4], or large WISPs, such as *Cometa Networks* [1], have similar goals with the PWC. WISP associations attempt to standardize technologies, protocols and behavior among existing WISPs in order to make WLAN roaming as seamless as possible. Cometa and other large WISPs attempt to set up new WLAN APs in hotspots and create their own standards, usually by investing a substantial amount of capital in the process.

A distinctive characteristic of the PWC is that it allows the ADs to make independent decisions concerning the amount of resources (e.g. access bandwidth) they contribute. In that sense, PWC is a ‘pure’ P2P system, similar

in principle to existing P2P file sharing applications such as Gnutella, Kazaa, etc. No central entity controls the interaction between the peers (the ADs), which dynamically enter and leave the system having full control of their participation level in the community.

This characteristic of the PWC enables a more scalable, flexible, low-cost and economically efficient solution for global broadband wireless coverage than existing schemes. In a PWC, however, without the appropriate incentives, actions are taken by individual ADs without taking into account the costs and benefits to other ADs in the system. The result of this is, in general, inefficient usage of the system. In the most simple and extreme case, free riding is complete and each AD offers no resources (in order to minimize its cost—a decrease in the quality of the service provided to its own local customers) while consuming as much as possible of other ADs' resources. Altruism (i.e., non-self-interested behavior) can go some way to correcting this inefficiency; this may be (part of) the explanation of why existing P2P systems (such as Gnutella) operate with some degree of success even if relevant studies [5] [17] indicate that the majority of participating peers in such systems are free riders. It is unlikely, however, that altruism will be sufficient to correct all of the inefficiency present in a P2P system.

So what are the appropriate mechanisms that are needed in order to give peers the correct incentives to contribute to the P2P system? In standard markets, prices provide the appropriate incentives. However in P2P systems where no global information is available (peers acquire information only by communicating with other peers) and there are externalities in the actions of peers (the contribution and consumption of resources of one peer affects the utility and cost of all the others), freely-determined (unregulated) prices would not lead to efficient behavior. Moreover, the complexity of implementing price mechanisms involving real money in a highly distributed P2P system, in which there is no central controlling entity, motivates the search for simpler to implement incentive mechanisms. Price mechanisms require price dissemination which by itself is often costly to provide and hence requires additional incentives. Involving real money (or even tokens which may be converted to real money) in transactions between peers adds substantial complexity in terms of security mechanisms and also requires the existence of a central authority who will do the clearing between peers. Also, human factors may discourage the use of P2P software, either directly, or because the mental burden of making cost decisions makes such a system less pleasant to use. So, in P2P, there may be no explicit prices, but implicit ways to account for production and consumption of resources by

individual peers. Specific system rules, implemented as part of the P2P software running on each agent, could be used to restrict the behavior of the peers and influence their decisions in order to achieve more efficient usage of the system.

Existing P2P file sharing applications have recently started to incorporate system-specific rules into their applications. In most cases these rules are very simple and compensating in nature. In Kazaa [2] for example, the contribution of each peer is computed and, according to its level peers have the corresponding priority in case of congestion. Note that in the absence of any central control, these mechanisms are fairly easy to subvert for a technically minded user; he can simply modify his client to claim a greater contribution than is true.

On the other hand, P2P systems following the free market approach seem to be not as successful, as in the case of Mojo Nation. Mojo Nation, was a file sharing P2P system that was developed in 2000 but is no longer in commercial development. It was the first (and so far only) 'real' P2P system to follow the free market approach (using an internal currency called Mojo). Note that the Mojo currency was managed by a central bank. The commercial model of the developers was not successful although a related open-source development continues [3] without the micro-payment scheme. There are also a few proposals to use prices for creation of the suitable incentives to peers to offer their resources, either set by a central entity that has full information of peers' interactions [12] or forming a free market [8].

There is also significant work on providing incentives in the context of ad-hoc networks, which have many characteristics in common with P2P systems as the one presented in this paper. The work in [18] considers a rule based mechanism, where the percentage of packets a node forwards depends on the percentage of its own packets that are forwarded by neighboring nodes. The works in [7], [19] consider a credit-based or virtual currency approach, which require mechanisms for guaranteeing the integrity and protection against forgery. The work in [13], [9] considers a price-based approach.

Our approach is to use rules for influencing the behavior of the peers instead of prices. As we will argue, such rules can approximate arbitrarily closely the effects of prices, and are simpler to implement. One may think of these rules as being designed and enforced by a regulator whose goal is to improve the economic efficiency of the overall system by avoiding free riding.

The remainder of this paper is organized as follows. In section II, we motivate the PWC system and present its key design principles and high-level architecture. In sec-

tion III, we propose an economic model which motivates the use of rules to control the behavior of peers. Section IV describes in more detail the components of the PWC architecture and identifies certain key architectural and implementation issues that require further research (in some cases providing directions for the design of mechanisms to address them). Section V concludes the paper.

II. P2P WLAN CONSORTIA: MOTIVATION AND HIGH-LEVEL ARCHITECTURE

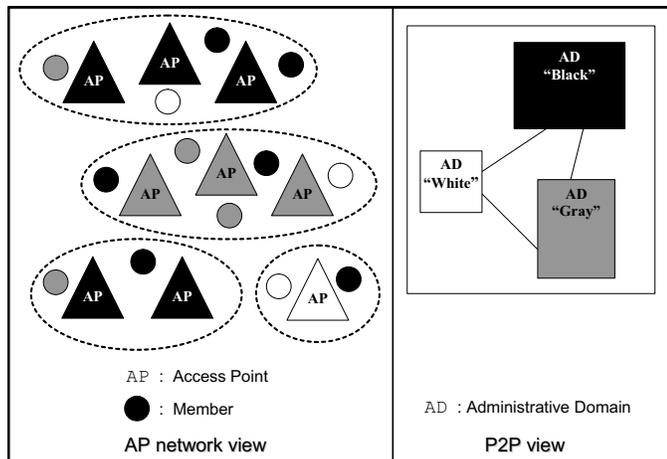


Fig. 1. A P2P WLAN Consortium consisting of three peers (ADs) and their AP networks and Members (indicated with the corresponding color)

A brief description of the main entities involved in the PWC follows, alongside the labels that will be used to reference them throughout this paper.

WLAN Administrative Domains (ADs): WLAN ADs constitute the peers in the PWC (see figure 1). We will use the terms ‘Peer’ and ‘AD’ interchangeably. We envisage ADs covering the full range of possible sizes. From a private residence with a home WLAN kit to a university campus with an internal network of WLAN cells and a company that offers WLAN access to employees, or even a WISP with a nationwide network.

WLAN Access Points (APs): The physical devices that offer local wireless coverage, which an AD deploys in order to cover specific geographical locations. The cells that APs form may or may not overlap.

AD Members (Users): The users interact with the PWC and consume network resources. These users should not be confused with the user-centered P2P notion of ‘Peers’. A PWC peer is both a provider and a consumer of resources but these two functions are well separated, with the AD providing resources to visiting users, while its own roaming members consume the resources provided by

other ADs. There is a number of ways a user can be associated with an AD: it could take place through a paying relationship (WISP case), a real-life family relationship (home WLAN), or some other arrangement (e.g. students that are registered users of their University’s AD).

User-Agents (UAs): The client devices and associated software components that users employ to consume AD resources. These devices would probably be portable and support standard Internet protocols and applications.

A. The Case for P2P WLAN Consortia

The main novelty of the PWC is its P2P nature. We claim that when coupled with a flexible set of system rules regarding reciprocity, a PWC would be a more efficient solution than others because of:

1. **Scalability:** a PWC can achieve wide coverage, as opposed to hotspot-only coverage that WISPs offer today, since global infrastructure costs can be effectively shared among (potentially millions of) ADs and the system can be build over time, with independent and small investment decisions.
2. **Decentralization:** the PWC is designed around complete AD autonomy and AD independence from central authorities, a fact that can make the PWC more socially-acceptable and economically efficient.
3. **Flexibility and low complexity:** the PWC replaces Telecom-style (or ISP) peering agreements (roaming contracts) among providing peers with more flexible arrangements. In traditional peering agreements, peers accept to serve all of each other’s roaming customers, creating unbalanced situations when the roaming traffic is not symmetric. In our proposal, peers have control over the amount of resources they release to roaming customers as a result of the peering agreement. This ‘managed’ peering with its extra flexibility allows peers to benefit more and hence creates more motivation for participation.
4. **Economic efficiency:** the PWC would work as a regulated market (e.g., the regulator could specify the rules) instead of a free market, where certain operators might acquire strong market power (and for example raise prices for services above the socially optimal because of their strong market position). The problem of tuning the appropriate parameters faced by the regulator becomes simpler as the number of peers grows and peers belong to a small number of types.

B. High-Level Architecture

Figure 2 shows two administrative domains (AD1, AD2) of a PWC. In white, we represent the support modules that would exist in any typical WLAN AD, even if it wasn’t participating in a PWC. These modules include the

the P2P software running on each agent. By ‘rules’ we mean P2P community-wide constraints on peer behavior, that may replace price mechanisms or supplement them. For instance, rules may express constraints on the relation between the rate of resource availability and resource requests made by a peer, or constrain the behavior of peers who wish to join a particular group, or constrain prices that can be charged.

The Model

We propose a mathematical model which allows us to evaluate the effect of various parameters in the economic performance of the overall system. It also justifies the use of rules as a substitute to prices.

Let $b_i(\{Q_j\}, \{r_{ij}\})$ be the rate of benefit obtained by peer i (an AD) when the rate of service requests (roaming members of this AD) directed to peer j is r_{ij} , and these are served by peer j with quality Q_j . In our case, think of Q_j as expressing the specific success probability that user agents face while requesting access from peer j (or packet delay in the case that service requests are always accepted but assigned to congested resources). In our model, peer i controls the rate of service requests he makes to its peers, being less than some maximum rate \bar{r}_{ij} , by not allowing all his roaming customers to request service at their remote location (here \bar{r}_{ij} is the rate of roaming customers of peer i which enter the area covered by peer j). In most cases, there may be substitution between elements of the vector $\{r_{ij}\}$, captured in the definition of the function b_i . This models the case of access services of different peers spanning the same geographic location.

The rate of cost incurred by peer i is denoted by $c_i(Q_i, \sum_k r_{ki})$, where Q_i is the quality level maintained for serving the requests of its peers, when these arrive with a total rate $\sum_k r_{ki}$ (to simplify notation, we implicitly assume that the sum is for all $k \neq i$). This cost is a function of the resources allocated by peer i for serving the other peers with quality Q_i . In many cases this amount of resources is only known privately to peer i . This motivates us to write it as a function of the *observable* (by third parties) quantities r_{ki} and Q_i . Then the net utility of peer i is $b_i(\{Q_j\}, \{r_{ij}\}) - c_i(Q_i, \sum_k r_{ki})$, and the social welfare is

$$SW = \sum_i [b_i(\{Q_j\}, \{r_{ij}\}) - c_i(Q_i, \sum_k r_{ki})]. \quad (1)$$

Here, the benefit function $b_i(\cdot, \cdot) \geq 0$ is assumed to be increasing and strictly concave in both of its arguments (expressing saturation effects), and the ‘cost’ function $c_i(\cdot, \cdot) \geq 0$ to be increasing and strictly convex in both its arguments, for all i . To reflect the fact that a peer

gains no benefit when the success probability is zero, let $b_i(0, \cdot) = 0$.

Our goal (think of us as the social planner) is to seek prices under which the maximum is achieved in (1). Then, we would post these prices and induce each peer i to operate at the socially optimal levels of Q_i^* and $\{r_{ij}^*\}$. Taking derivatives of SW with respect to Q_i and r_{ki} and setting these equal to zero we obtain

$$\sum_{j \neq i} \frac{\partial b_j}{\partial Q_i} - \frac{\partial c_i}{\partial Q_i} = 0, \quad (2)$$

and

$$\frac{\partial b_j}{\partial r_{ji}} - \frac{\partial c_i}{\partial r} = 0. \quad (3)$$

The above equations suggest to use prices (different for each peer i)

$$p_i^Q = \sum_{j \neq i} \frac{\partial b_j}{\partial Q_i}, \quad p_i^r = \frac{\partial c_i}{\partial r}, \quad (4)$$

where the derivatives are computed at the optimum of (1). Using these prices, peer i offering quality level Q_i is rewarded a negative charge (receives) $p_i^Q Q_i$, and incurs a positive charge (pays) $\sum_{j \neq i} r_{ij} p_j^r$.

The resulting prices p_i^Q and p_j^r motivate¹ the use of a rule for peer i of the form

$$Q_i \geq \sum_{j \neq i} \alpha_{ij} r_{ij} + \beta_i, \quad (5)$$

where the vectors of weights $\{\alpha_{ij}, \beta_i\}$ are defined for each peer i based on the optimal prices. More specifically it can be easily proven that for

$$\alpha_{ij} = \frac{p_j^r}{p_i^Q},$$

$$\beta_i = Q_i^* - \frac{\sum_{j \neq i} p_j^r r_{ij}^*}{p_i^Q}$$

the above rules lead the system to the optimal equilibrium.

Observe that since optimal prices need to be personalized, optimal rules must also be personalized in the general case where peers are not symmetric. In practice (for instance when there is a large number of peers, each peer being ‘small’), uniform prices and hence uniform rules are expected to perform adequately. Or, if there is a small set

¹There is a more general result which suggests a close correspondence between optimal prices and optimal rules. Hence, the argument for using rules is not of an informational advantage but of simplifying implementation.

of peer classes, we may use rules specially targetted to a particular class, and have peers self-select the class they want to belong.

In such a model the rule involves the quantities Q_i and r_{ij} . Both quantities may not be easy to measure accurately since these involve parties with conflicting interests. For instance, in the case of Q_i representing the probability a service request to be accepted by peer i (estimated as the ratio of successful service requests to the total number of requests in a given time interval), how would one measure such a quantity? Asking peers about their perceived Q_i offered by peer i may not lead to truthful answers since they may have the incentive to downgrade peer i in order to force him to raise his performance. One may consider the possibility of using reputation for motivating peers to answer truthfully. For instance, after collecting answers from all peers, the peers with answers being outliers in the above statistical sample may be punished and have their reputation lowered. To offer the right incentives we could make (5) depend on the reputation of the peer. For instance, we could use

$$q_i Q_i \geq \sum_{j \neq i} \alpha_{ij} r_{ij} + \beta_i, \quad (6)$$

where q_i is the reputation of peer i .

In this simple reputation model we assumed that peer i offers the same quality Q_i to all peers, i.e., he can not discriminate against individual peers (for instance, by using cryptography to hide the particular identity of the roaming customers). The model can easily be extended to allow for such a service differentiation, where peers can offer different blocking probability to members of different ADs.

IV. IMPLEMENTATION ISSUES

A. Rule-Design Issues

Even though we presented above a framework for economic analysis of a PWC, there is a need to select and specify completely the appropriate functions for the economic models, which will lead to the determination of the range of the coefficients for the rules governing the PWC.

Picking the right benefit and cost functions that are simple enough but reflect realistically the AD's economic preferences, is an important task in this work. Once such functions are chosen, then the optimal prices and the corresponding rules will be determined. Then, using simulation at this more abstract level, we can study the effect of the rules on the economic efficiency of the system, together with sensitivity and stability issues (note that decisions about resource provisioning and consumption are made asynchronously by each AD, and roaming traffic is in general probabilistic). As a last step, we can study by

using a more detailed simulation the behavior of a real PWC with a large number of ADs using rules with parameters and decision functions obtained by the analysis of the more abstract model.

The benefit and cost functions in the economic model should take into account a number of features that characterize and differentiate ADs. Some of these may be objective. In the case of PWC, there are two objective differentiating characteristics of peers: (1) their capacity² and (2) their 'footprint'. Peers with high capacity can serve a large number of customers with lower cost and better QoS. The number and geographical location of the access points that an AD shares in a PWC (the AD's footprint), affects the demand that this AD faces, and as a result the value it generates to the system. Peers that offer WLAN connectivity in remote areas offer less amount of their resources than other peers, since they serve less requests. Nevertheless, they generate greater (per request) value and they contribute significantly to the 'ubiquitous access' target of the system.

More generally, the design of appropriate rules, easily and provably enforceable, and robust is a key open issue, discussed further below.

B. Decision Support System for ADs

The question of whether it is (economically) beneficial for an AD to join a PWC and more importantly the levels of resources that the AD should make available to the PWC might be a very complex one. Also, such issues should be answered frequently due to changing load conditions (both of local and roaming traffic). A *Decision Support System* would be beneficial here. One could imagine that such a system could be incorporated into the peer software, permitting automatic operation. Many components of this system will implement parts of the economic model validated by simulation, as discussed earlier.

C. Generic Rule Structures

We have discussed providing incentives to contribute by imposing community-wide rules. These rules represent a fundamental statement of the nature of the P2P community (here, the PWC). They must be made explicit to the peers and we expect different communities to adopt different structures of rules to attract peers. To enable this competitive playing field, we aim to provide developers with a framework in which it is easier to create the software that instantiates these rules. That support comprises three elements of functionality, a generic service provision

²the wireless access bandwidth or bandwidth to the Internet depending on which is the bottleneck

middleware, a system of local policy implementation, and a distributed rule enforcement system.

C.1 Generic Service Provision

We are building a framework that supports the ability of peers to offer services to one another. This provides a structure within which we can describe, create, and enforce rules in terms of services consumed and provided. The generic framework includes facilities for identification and authentication of peers within the community and processes to constrain who can join and who must leave. It also provides mechanisms by which providers can obtain signed receipts for services provided.

C.2 Local Policy Execution

Local policy execution functionality allows peers to define the rules and local policies that they wish to apply to their participation in the communities. Upon joining a community, peers will become aware of the rules that they are obliged to obey. Peers will be able (and expected) to simply import these rules into this execution functionality which will then attempt to ensure they conform (or at least warn them of impending breach). In this way the community will operate efficiently.

In the case of the PWC, each peer i knows the rate at which its members request service from other peers. Thus he can compute $\sum_{j \neq i} \alpha_{ij} r_{ij}$ and thus deduce Q_i , the proportion of access requests he must satisfy in order to meet the rules.

However, this functionality does little to prevent peers from acting against the rules. In our scenarios we do not rely on peers to be using special purpose computing platforms that can be trusted by others. In these circumstances, it will remain possible for peers to arbitrarily modify the behavior of the software running on their local device. We cannot and do not attempt to remove the autonomy of peers in this way. Instead the system imposes distributed constraint structures so that peers will have an incentive to conform to the community rules. We call the implementation of these structures rule enforcement.

C.3 Rules Enforcement

Rules enforcement has two aspects: distributed accounting to assess consumption and contribution and then the application of incentives i.e. punishments and rewards.

Distributed Accounting

When one peer receives service from another then our generic service provision middleware builds in protocols that require the exchange of one or more signed receipts. These then provide unforgeable evidence of consumption

that can be provided to the community. Any peer can deduce the rate of consumption of any other, by inspecting and aggregating these receipts.

In general this computation requires access to every single peer to deduce consumption for any other particular peer. However, this is not as serious a problem as it sounds. This computation need only be undertaken when it is required to check the conformance of a peer to the rules. In a community where most peers obey the rules most of the time, such checking can be done rarely and still have an adequate deterrent effect. Should it be necessary to compute this information more often, then statistical approaches can be used. A reasonable approximation may be calculated by sampling peers, particularly where stratification of the peers is possible.

Under a more optimistic system model in which we assume some large proportion of peers behaving correctly then it is also possible to implement systems in which these receipts are transmitted around the network and aggregated. This can allow much more efficient and regular interrogation of the status of peers and hence more thorough investigation of rule-breaking.

In the case of the PWC, this form of distributed accounting allows any member of the community to assess $\sum_{j \neq i} \alpha_{ij} r_{ij}$ for peer i . However, to discover breaches of the rules will also require them to estimate Q_i and this is more difficult. This is because no peer can provide unrepudiable evidence to a third party that another peer has *refused* him service. It is never possible to distinguish a refusal to provide service from a failure in the requesting process (or a malicious false claim that a request was made). For such measures the distributed accounting system must fall back on more statistical techniques as developed to assess the reputation of peers.

Distributed accounting provides the functionality for peers to gather aggregated opinion from the other members of the community. In this way a peer can deduce an average of the claimed values of Q_i that other peers have received from i . We would like to use this measure as an estimate of the real Q_i .

This approach essentially amounts to an online reputation system. There is much research on such systems, and this includes discussions of various theoretical difficulties such as difficulties in eliciting honest reports. However, in practice, they seem to work reasonably well [16].

Reward and Punishment

We have described how distributed accounting allows rule-breaking to be discovered by the community. What sanctions can the community apply? The answer to this question is strongly tied to the nature of identity in the sys-

tem. In some communities there may be a strong notion of identity. The PWC could well operate in this way; an AD has a physical presence which defines their identity and is visible to other peers e.g. when their members roam to it. This physical presence is expensive to change and so ties peers strongly to a long term identity. Sanctions can easily be applied to this identity. The simplest is to bar the peer from the community, for a period of time or for ever. If any form of legal contract exists then it may well be possible for the sanctions to extend from the online community into the real world through law enforcement.

In most online communities, however, the notion of identity is weak. Typically, we expect peers to be identified by pseudonyms which cannot be forged but can be abandoned and replaced. If there is also a significant cost of entry then that provides an incentive to peers not to change identity and hence expulsion is still a meaningful sanction.

Communities with cheap pseudonyms need some other way to build trust. This is usually achieved by treating newcomers badly [10]. Again these structures allow expulsion to be a meaningful sanction.

D. 'Disconnected' Operation

It is highly desirable for the scheme to be operational even when no communication is possible between the server of the AD a roaming user is visiting and the server of its home AD. This would be possible if the User Agent (UA) of the member of an AD who visits another AD could 'carry' with it all the required credentials and could adequately prove its AD's identity and its good standing in the PWC (e.g. reputation, ability to 'pay', etc.).

V. CONCLUSION

We have introduced the concept of a Peer-to-Peer Wireless LAN Consortium and motivated its existence in economic terms. We supported our view that rules based control is more appropriate in this setting than a price based, free market approach and discussed the relationship between prices and rules. Then we discussed a number of important issues that need to be investigated and resolved in order to design such practical and efficient consortia.

REFERENCES

- [1] Cometa Networks. <http://www.cometanetworks.com/>.
- [2] Kazaa. <http://www.kazaa.com/>.
- [3] MNET. <http://mnet.sourceforge.net/>.
- [4] Pass-One. <http://www.pass-one.com/>.
- [5] E. Adar and B.A. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000.
- [6] S. Buchegger and J. Y. L. Boudec. Nodes bearing grudges: Towards routing security, fairness, and robustness in mobile ad hoc networks. In *Proceedings of 10th Euromicro Workshop on Parallel, Distributed, and Network-based Processing*, 2002.
- [7] L. Buttyan and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *ACM/Kluwer Mobile Networks and Applications (MONET)*, 8(5), October 2003.
- [8] Rajkumar Buyya, David Abramson, Jonathan Giddy, and Heinz Stockinger. Economic models for resource management and scheduling in grid computing. *Special Issue on Grid Computing Environments, The Journal of Concurrency and Computation: Practice and Experience (CCPE)*, May 2002.
- [9] J. Crowcroft, R. Gibbens, F. Kelly, and S. Ostring. Modelling incentives for collaboration in mobile ad hoc networks. In *Proceedings of WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, March 2003.
- [10] E. Friedman and P. Resnick. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.
- [11] S. Fluhrer, I. Mantin, and A. Shamir. Weaknesses in the Key Scheduling Algorithm of RC4. In *8th Annual Workshop on Selected Areas in Cryptography*, August 2001.
- [12] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for Sharing in Peer-to-Peer Networks. In *Proceedings of WELCOM'01*, November 2001.
- [13] O. Heri, S.-C. Mau, and N. B. Mandayam. Pricing for enabling forwarding in self-configuring ad hoc networks. 2003.
- [14] IEEE Std 802.11. Wireless LAN Medium Access Control (MAC) and Physical Layer Specifications. ISO/IEC 8802-11:1999(E), 1999.
- [15] P. Michiardi and R. Molva. A game theoretical approach to evaluate cooperation enforcement mechanisms in mobile ad hoc networks. In *Proceedings of WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, March 2003.
- [16] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12), 2000.
- [17] S. Saroiu, P.K. Gummadi, and S.D. Gribble. A Measurement Study of Peer-to-Peer File Sharing Systems. In *Proceedings of Multimedia Conferencing and Networking, San Jose*, January 2002.
- [18] V. Srinivasan, P. Nuggehalli, C. F. Chiasserini, and R. R. Rao. Cooperation in wireless ad hoc networks. In *Proceedings of Infocom 2003*.
- [19] S. Zhong, J. Chen, and Y.-R. Yang. Sprite: A simple, cheat-proof, credit-based system for mobile ad-hoc networks. In *Proceedings of Infocom 2003*.