Community Building over Neighborhood Wireless Mesh Networks

PANAYOTIS ANTONIADIS, BENEDICTE LE GRAND, ANNA SATSIOU, LEANDROS TASSIULAS, RUI L. AGUIAR, JOÃO PAULO BARRACA, AND SUSANA SARGENTO

The fourth generation (4G) network paradigm has long been sought. A user-centric vision for such "always best connected" next-generation 4G networks is neighborhood Wireless Mesh Networks (WMNs). In this context, the vision for the formation of WMNs reflects the trade-off between the immediate self-interest of the user, and the user's need for social contacts. Notably, users would be required to pool their resources in order to support the creation and operation of the underlying communication network (participating at all physical, access, and network layers), but also for the service provision on top of it. We argue that the design of *communities* suitable for this environment will encourage users to participate, enable trustworthy network creation, and provide a social layer, which

Digital Object Identifier 10.1109/MTS.2008.918950

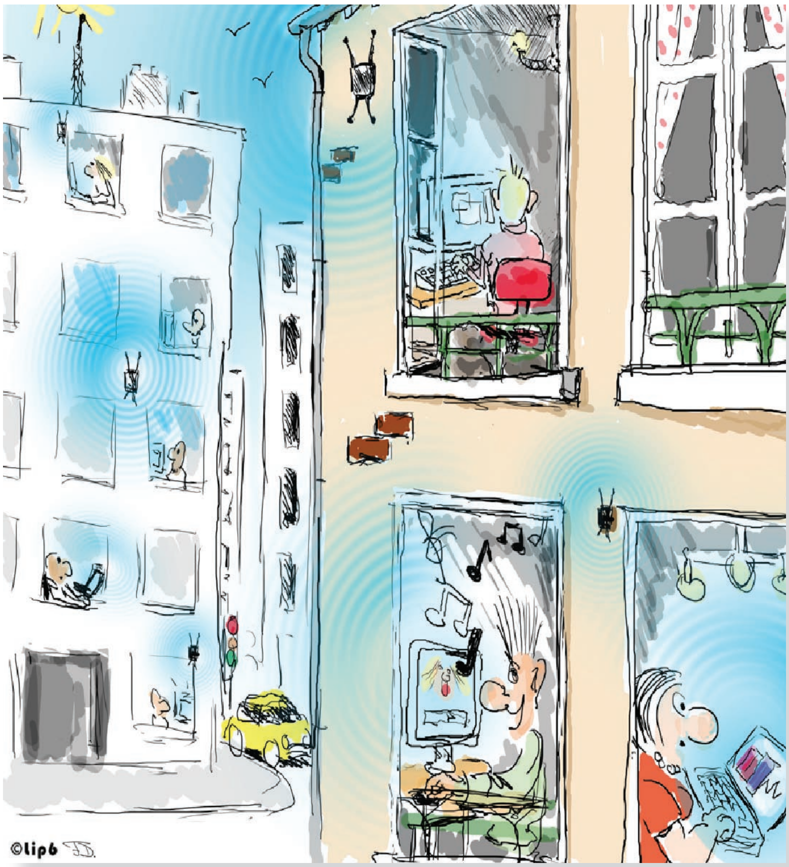

DRAWING CREDIT: ©FRÉDÉRIC DELVALLE—LIP6. USED WITH PERMISSION.

can be exploited in order to design cross-layer incentive mechanisms that will further encourage users to share their resources and cooperate at lower layers.

The Technology – User Rules!

4G networks will provide a converged environment for technology and service provision, and wireless communications are an essential aspect of this environment. In recent years different initiatives have tried to exploit the advances of Wi-Fi technology to offer ubiquitous Internet access, mostly in the form of free or commercial wireless hotspots. However, the cost, which notably is mainly due to fixed and recurring costs of the wired infrastructure, along with Wi-Fi's relatively small transmission range, make it difficult to cover wide areas using only hotspots. Wireless Mesh Networks (WMNs) provide attractive means to reduce costs, since only some of the wireless routers that form the network must have a direct connection to the Internet. Both university projects (e.g., M.I.T.'s Roofnet,[1] Rice University's TAP[2]) and town initiatives (e.g., Athens Wireless, Paris Sansfil) have started along these lines, building city-wide WMNs with directional antennas (wireless backhaul networks) offering added value services and/or free Internet access. Other municipalities (e.g., Philadelphia, San Francisco) have also started the deployment of "Wi-Fi blankets" (Municipality Wi-Fi), but there is growing criticism as to what extent this is a feasible scenario for the moment, both economically and technically.[3] Nevertheless, such wireless infrastructures can provide connectivity with significantly less cost than wired solutions in many cases (e.g., in remote and secluded areas).

This cost issue is not a major factor when exploiting user-owned infrastructure, i.e., the already available home users' wireless access points. Interestingly, both research and commercial efforts exist to enable peer-to-peer communities to share Internet access through their wireless access points for mobile use [1], [2]. Moreover, the existence of numerous wireless access points in urban areas has brought closer the wider vision of Nicholas Negroponte, five years ago, of a "Wi-Fi 'lily pads and frogs' broadband system built by people for the people" [3], or neighborhood self-organizing WMNs, according to Microsoft [4]. That is, in addition to providing a wireless (one-hop) link towards the Internet, users can form wireless mesh networks with their own wireless access points exploiting the large amount of unutilized connectivity, which could further facilitate people's access to the Internet. However, this is not WMN's only potential value: among many other advantages, WMNs also can provide additional network capacity (e.g., for content distribution or games), enable the sharing of other resources such as storage (e.g., for backup services) and content (file sharing or caching), and offer services to mobile users and/or peer-to-peer applications. Additionally, WMNs could also strengthen the *social capital*[4] between people living in the same neighborhood (and close the gap between virtual and physical communities) by supporting a large variety of – possibly novel – social and collaborative applications.

Finally, WMN deployment can be seen as an intermediate step towards more general *ad-hoc*, user-centric (including mobile users) environments. A user's physical neighborhood offers a friendlier environment for being introduced to this type of communication (e.g., in most cases there are already many trusted users from real-life peers). Moreover, investment on this collaborative communication can be considered worthwhile, since their peers will remain for a long time in their vicinity. So, neighborhood WMNs could play an additional role towards the 4G vision [6] by motivating people to experience and exploit this type of communication.

However, interest in forming a neighborhood wireless network may depend on the expansion of the wired Internet. If the wired Internet is available everywhere and non-congested, users might not have an obvious incentive to participate in a wireless architecture, because even "neighborhood-oriented" services and communities could be supported through the Internet.[5] Nevertheless, even in this extreme case, neighborhood WMNs have a specific value, obviously related to privacy and trust issues, but other factors are relevant as well as we present in the following.

An important differentiating feature of neighborhood WMNs compared to the wired Internet is the need for users to actually contribute to their creation and operation. This means that incentives – explicit or implicit – should be in place in order for users to participate and share their resources. Interestingly, the existence of such incentives would then provide the means for supporting valuable community-oriented applications, which would increase the value of the network itself and thus further encourage users to participate and cooperate.

[1]http://pdos.csail.mit.edu/roofnet/design/.

[2]http://taps.rice.edu/.

[3]See, for example, http://news.com.com/Cities+deploying+Wi-.

[4]Social capital is a core term, which is widely used and variably defined in social sciences. We use it in our work to reflect the overall value generated in a community due to the social activity and networks formed (see [5] and references therein for a discussion on how the Internet could affect social capital).

[5] E.g., http://www.local2me.com/ or http://www.i-neighbors.org/.

One way to define an incentive behavior is to consider it as a system rule, whose goal is to influence participating agents to behave.

Need for Cooperation at Different Layers

In our vision, a user participating in a neighborhood WMN is a *node* at different layers of the system architecture (physical, access, network, application, and social), as depicted in Fig. 1. At the physical and network layers, terminals are usually assumed to comply with a predetermined protocol that prevents them from antagonizing and pursuing their own benefit in an autonomic fashion. In certain cases though, the interaction has a game theoretic nature and desirable behaviors emerge in the context of an antagonistic framework [7], [8]. Power control is such an example. If energy consumption is not an issue (as in our case), then it is hard to discourage the terminals from using maximum transmission power as they seek better link quality and more bandwidth, in which case they may cause diminished bit rates for everybody else due to excessive interference.

At the network layer there is an even clearer need for to cooperation, and for incentives to promote that cooperation. In a WMN, users' available bandwidth will be reduced if they forward packets belonging to other nodes. This is a traditional problem in mobile *ad-hoc* networks (facing stringent resource constraints), which has led to significant research work regarding the provision of suitable and trustable incentives for cooperation (see [9]-[11] among many others).

At another level, users participating in a peer-to-peer application (such as file sharing, backup services, or grid computing) need to contribute additional resources (e.g., content, storage, CPU cycles, etc.). But, ideally, they would prefer "free riding" on the contributions of their peers by consuming available resources and services without contributing themselves [12]. Similar behavior (called "lurking") arises also in the case of on-line communities where users' active participation is of critical importance.

Moreover, additional incentive issues arise when users should contribute themselves to the implementation of the required management functionality at different layers (e.g., build routing tables or reply to service queries). Clearly, the need to take all these different decisions and potential selfish strategies into account makes the design of suitable incentive mechanisms a very challenging task, which would lead to a conceptually complicated game theory problem. Finally, relying on the limited skills of normal users for tampering the software is not safe, since hacked versions could be made easily available in the Internet. This means that, even if a suitable protocol was implemented in the corresponding software (e.g., in the case of power control), incentives to steer users to follow this protocol would still be necessary.

Incentive Mechanisms and the Human Perspective

One way to define an incentive mechanism is to consider it as a system rule, whose goal is to influence participating agents to behave in a certain manner, by rewarding (or punishing) them according to their actions. For example, in a traditional market, a *price* is a monetary reward for production and a punishment (a charge) for consumption. The designer's task is then to decide on the mechanism to compute and set "prices appropriately" in order to reach a specific goal.

The two most common objectives considered in economics are social welfare maximization (also called economic efficiency) and fairness. The social welfare maximization approach considers two user metrics (namely *utility* and *cost* for consuming and contributing resources, respectively), which

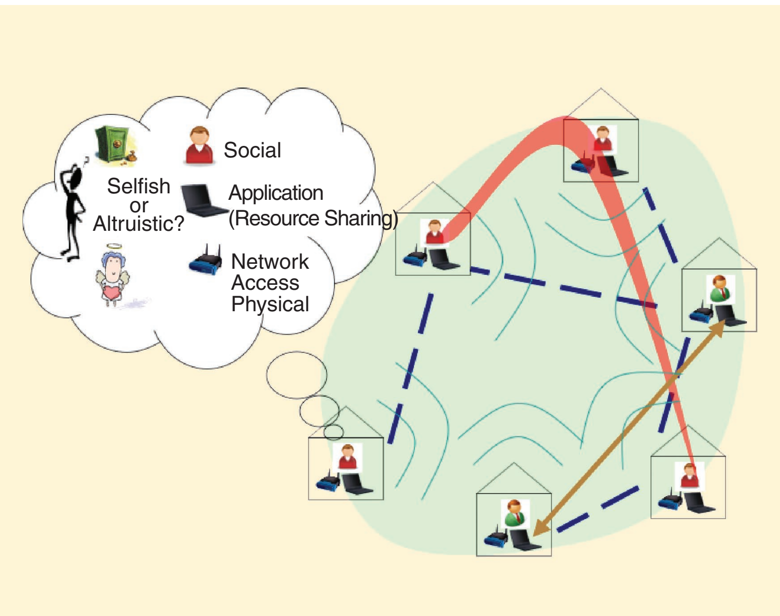

Fig. 1. The different layers of required user cooperation.

Table I
Incentive Mechanisms for Resource Provisioning

Mechanism	Incentives	Users' decisions
Pricing	Charges/payments for consumption/contibution	Level of consumption and/or contribution
Entry Fees	Fixed contribution, fair resource allocation	Participate or not
Reciprocity	Consumption=contribution	Level of consumption
Reputation	Resource allocation based on past behaviour	Level of contribution - Quality of Experience

are in general private information. It aims at maximizing the total utility minus the total cost, assuming participation is *rational* (i.e., agents seek to maximize their own net benefit). On the other hand, the fairness approach treats all agents equally either by principle, or by acknowledging the inability to convey more information.

Although the choice between these two objectives is controversial across several disciplines such as political philosophy, sociology, and economics, the selection of a certain approach is not always due to economic considerations. For example, the complexity of computing optimal prices in many economic problems and/or obtaining the required information, the difficulty of implementing micro-payments, and the mental burden that they may require, are some reasons why pricing mechanisms proposed in the literature for addressing many of the aforementioned problems (e.g., [9], [11]) are not deployed in practice despite their nice theoretical properties.

As a result, alternative "fair" solutions are usually favored, such as simple fixed contributions [2] or reciprocity [12]. For example, reciprocity dictates that all users should contribute the same amount of resources they consume. But although this is a theoretically simple incentive mechanism, its enforcement is not trivial in a distributed environment, since it requires the existence of some kind of virtual currency, except in few cases where a direct exchange of resources is possible (e.g., BitTorrent). Nevertheless, it still puts a significant mental burden on users, and can discourage altruistic behavior, which, according to current practice, seems to play an important role in the context of peer-to-peer (p2p) applications.

Acknowledging the above issues, reputation mechanisms [14], originally introduced in distributed marketplaces (such as eBay), have been considered promising incentive mechanisms in the context of all types of p2p applications providing a more qualitative (than quantitative) way to reward/punish good/bad behavior. More specifically, a user's reputation could be seen as a way to aggregate her past behavior into a single value. This value is in general a function of the individual users' *ratings* based on the corresponding user's observed behavior. Then rewarding users with high reputations (e.g., giving them priority) and/or punishing those with low reputation (e.g., denying service to them) would ideally provide them with the suitable incentives to maintain high values of reputation and thus behave correctly. An important issue is the reliability of the reputation values: these values have to be computed correctly (i.e., based on truthful ratings). This is a challenging problem when users may easily create a new identity/pseudonym and when information regarding the effort exerted by a user as a function of the outcome of a transaction[6] is hidden. This fact and the freedom in defining rewards and punishments increase the difficulty in formally evaluating the outcome of a specific mechanism. This may explain why such a plethora of reputation mechanisms have been proposed in the literature [15].

The mechanisms discussed so far (and summarized in Table I) assume that humans behave rationally, which is actually highly debatable. There are many cases where people seem to actually contribute "for free" (e.g., in p2p file sharing applications some peers provide a huge amount of content although no explicit incentive mechanisms exist), or at least, their reward is not apparent if considering only typical economic assets. The main motivation for such an "altruistic" behavior could be inherent or rely on subtler motivations related to long-term rewards [16] or immaterial ones such as the feelings of self-esteem, happiness, affection, etc. The latter require a social environment in order to be generated. The community spirit, the feeling of influence and/or importance, social status, and shared emotional connections are some of the aspects of social life that can play a positive role towards this end (see [17] for a well established theory on the *Sense Of Community* concept).

[6]E.g., in an *ad-hoc* network, when a packet does not reach its destination it is not straightforward for the sender to decide which intermediate node was responsible for this failure.

> Users' available bandwith will be reduced if they forward packets belonging to other nodes.

However, unlike human societies, in on-line environments, social interactions are restricted by the human-computer interface. Recently, the term *social software* has become established to describe the functionality implemented for supporting the operation of an on-line community. Many successful on-line communities owe their success to clever details incorporated in their software to reward cooperative behavior (see [18] and references therein). The way people interact with each other and create relationships, how they represent themselves, the feedback they receive concerning their popularity and activity, the elasticity in deciding which part of their activity is public or private are some of the crucial design choices in this context.

We propose the provision of such social incentive mechanisms encoded in the definition of suitable community rules and the corresponding social software in order to encourage people to share resources and services. The stronger the community ties, the less strict the corresponding incentive mechanisms need to be at the resource level. This is the cornerstone of our neighborhood WMN concept. But even when explicit rules are required, the existence of a community will create a suitable environment for providing additional social rewards related explicitly to resource sharing (see "Cooperation and Resource Provisioning" section), fundamentally cross-layer incentive mechanisms.

Interestingly, new sociological and psychological aspects arise due to the technologically enhanced social image of a user participating in our neighborhood communities. And specific research should be carried out in order to understand in depth how people's behavior would be affected in this new social and technological environment. Such a study is out of the scope of this article. In our work we identify the need to include the social layer in the provision of incentives at the lower ones and describe the required interactions between them. Further multi-disciplinary work is required in order to design specific mechanisms. In the following, before analyzing in more depth our cross-layer approach, we propose some basic characteristics of community building over neighborhood WMNs (prioritizing them according to their independence from the future evolution of telecommunications) and a corresponding network creation procedure, which is critical for the bootstrapping of the system.

User Participation

The most important differentiating characteristic of wireless neighborhood communities, when compared to Internet communities, is the *de facto* physical proximity between participants. This property may be exploited by applications; in the case of social applications, for example, physical proximity adds value to potential acquaintances thanks to the ability to transfer them into the real world as well, but also thanks to the increased level of trust and intimacy between people living in the same neighborhood.

However, trust is not obviously granted even for people living in the same neighborhood and privacy issues still exist. We could say that *privacy* is related to the reluctance to disclose personal information to others, and lack of *trust* to the reluctance to interact with others (either because of disagreements with their actions or because of the perceived level of their commitment and contribution). Interestingly, in the context of on-line communities, two different concerns exist regarding privacy: personal information stored in central databases of companies vs. personal information traveling "in the air." For the former, the self-organizing and distributed nature of the neighborhood WMN approach provides an attractive solution since it does not require central databases, while for the latter it is always possible to use secured communications with trusted people.

Concerning trust, neighborhood WMNs provide a realistic environment for the deployment of reputation-based mechanisms since relations are long-lived (compared to the highly dynamic nature of *ad-hoc* networks where it is not easy to build trust –"there is not enough time"). The potentially available location information could further assist towards strengthening the notion of identity, and thus make reputation mechanisms less vulnerable to whitewashing and Sybil attacks [19], [20]. Additionally, there is an existing social context, which should be exploited towards this end (e.g., in our environment people could know each other from real life or physical meetings arranged through the on-line neighborhood communities). This trust information would provide the means for users to discriminate between different community members. They could, for example, enforce different encryption policies in different sub-communities in order to protect their privacy (e.g., encrypting the content exchanged with trusted members), but also minimize their personal risk (e.g., by forbidding encrypted traffic of non-trusted users to pass their internet connection and possibly participate without their will in illegal actions). We elaborate more on these issues in the next section.

Note also that self-organization, besides its practical benefits, could also offer an important psychological advantage compared to centrally-managed solutions, since it could be seen as a "natural" extension of everyday social interactions, retaining the feeling of independence. In addition to its psychological effect, this independence could also be materialized in the communities design itself. In particular, users have the ability to configure the community forming and management functionality to suit their own needs. For example, they are free to implement effective and cost-efficient solutions to various challenging networking problems such as security support, spectrum management, or mobility. In this latter case, they could enjoy free mobile access in their neighborhood (e.g. for Internet or other services).

In addition to these general characteristics, a number of specific services and/or applications offered by existing on-line communities in the Internet could provide additional motivations for participation. Moreover, local services could be devised exploiting the *de facto* physical proximity between users, such as the management of real life neighborhood activities (e.g., decision making, organizational activities, announcements, etc.), games, socializing, and security support (e.g., supervision of both the digital and the geographic neighborhood).

Finally, we should stress that building communities over a neighborhood WMN could be seen also as a goal for 4G networking (rather than the means), since the value they provide may be significant (especially nowadays when virtual interactions have started threatening the physical ones). The fact that users themselves will have to support their communities should be exploited in order to strengthen their relationships and community spirit, rather than be seen only as a vision that could be reached through incentive mechanisms.

Network Creation

The first step towards community building in the neighborhood is the creation of the underlying wireless network. That is, the nodes should create *links* with their neighboring nodes and define their "next-hops" to forward queries or to reach specific destinations. Of course, in this context the notion of a link is artificial due to the broadcast nature of the medium; the creation of these links actually corresponds to an agreement between involved parties to use the same media access rules and participate in the same network. These agreements would be part of the community *doctrine* [21], together with the definitions of the different possible member roles, the membership rules and requirements, the incentive mechanisms, and the privacy or security parameters.

Ideally, routing tables should be created automatically according to network performance criteria. However, users might sometimes be reluctant to allow any link to be formed mainly because of associated costs and/or trust and privacy issues. In our approach, trust is built at the community level, and the system will automatically handle the required interactions at the network layer to set the "already trusted" community members as a network on itself. In this case, when network creation is treated as an inherent part of a community's activity, the creation of a link corresponds to a community "join" and is subject to the doctrine set by each community (e.g., certain communities could define geographical boundaries, or a minimum amount of resource contribution, or require an invitation from a member). Additionally, such links could belong to different types (e.g., "testing," "trusted") and have a differentiated treatment (see Fig. 2). Different treatments could correspond to different levels of trust (enforced through the exchange of cryptographic keys),

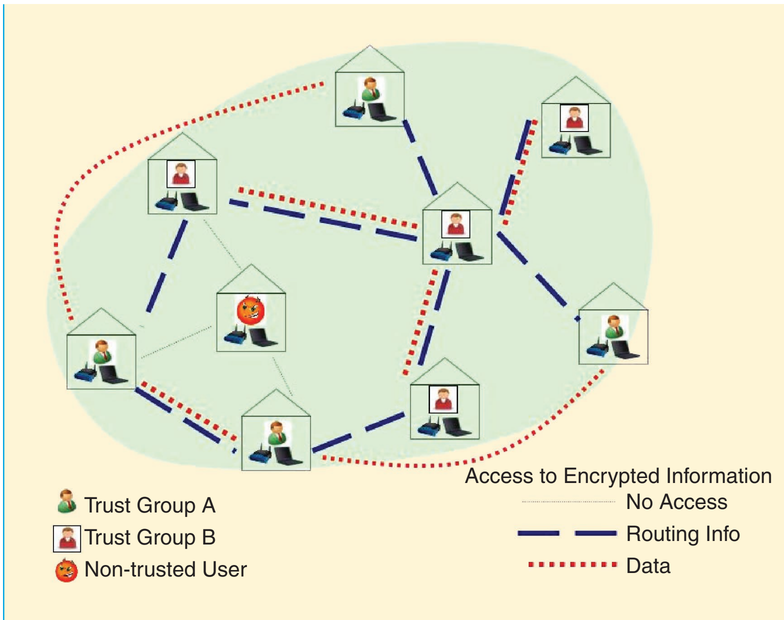

Fig. 2. Different types of links depending on the level of trust.

> **Different treatments could correspond to different levels of trust.**

different behavior in case of congestion (i.e., priority given to the more trusted ones) or in terms of propagation of service queries.

It will be beneficial for users to be part of trusted links, and the cross-layer approach to incentives will also promote appropriate behavior at higher levels. Moreover, this approach would ensure the maximum possible connectivity without making sacrifices in terms of trust, and could additionally provide the suitable incentives for users to behave according to the community doctrine at the social level. Thus, under this community-aware network creation, a *zero-level* community would be implicitly created with some standard functionality. On top of this, the community will be allowed to freely create their sub-communities with their own doctrines (in the same way *flickr* and *myspace*, for example, enable users to create their own groups with their specific participation and operation rules). This approach provides significant flexibility to users concerning the trade-off between privacy and accountability and that between trust/security and connectivity. However, there is always a compromise to be made towards this end, which in our case is related to the default trust level assigned to unknown zero-level community members.

The overall created network would finally be a WMN. This could potentially lead to distributed or hierarchical solutions (based on the concept of super-peers [22]), depending on the environment. In some cases (e.g., wireless access points provided by a municipality), the use of some centrally managed elements could even be assumed. Hierarchical solutions are particularly attractive, since they simplify community management and although there is a cost for super-peers, many users could agree to play this role (as current practice shows). This is in many cases due to altruistic motivations, but also to more practical benefits such as higher resilience to network events, faster information retrieval, access to better services and resources. In any case, providing additional explicit social incentives (e.g., visibility or acknowledgement) to users becoming super-peers would further ensure their existence. Additionally, one could also exploit the natural structure created inside communities and define meaningful dependencies between the roles formed as a consequence of social activities and interactions, and those assigned to users for the network (and/or application) management functionality.

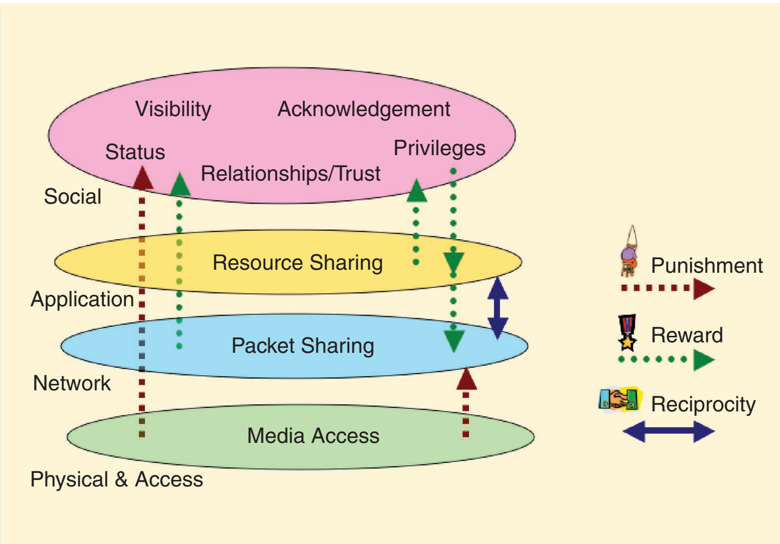

Fig. 3. Cross-layer incentive mechanisms.

Cooperation and Resource Provisioning

We will now provide some insights on how to link the social activity of users together with the required cooperation and resource provisioning at the lower layers (the main interactions between the different layers are summarized in Fig. 3). The new dimensions in our context compared to existing work for online communities [18] are the following: a) users contribute in the creation of a communication network with their own devices, b) their natural ability to self-organize, and c) the *de facto* proximity between them. The characteristics of this new environment can be exploited and incorporated in the system in order to support the envisioned neighborhood communities. Our goal here is not to propose specific solutions...but to define a framework for future research on the very interesting social, economic, and technological aspects that appear in this context and, most importantly, on their interdependencies.

Physical and Access Layer
For this layer, there are ongoing research efforts to build efficient cooperative protocols [7], [8]. Based on these protocols one should devise ways to detect deviations from the encoded behavior. Then, when such a deviation is detected by the neighboring nodes it should be considered misbehavior and the user should be punished through an appropriate change of the status of the corresponding links as described above.

In addition, specific social punishments could be used to further discourage misbehavior. Such punishments can harm explicitly one's social image (e.g., by means of a public "black list" or warning messages). Alternatively, they could affect the privileges and/or position of a user in the community (e.g., limit her ability to browse the complete graph of her social network).

Although physical proximity and long-lived relationships imply that users need to rely on each other, and punishments as "node isolation" will be quite effective, at the same time they can discourage participation. However, since cooperative behavior is encoded in the corresponding protocols, deviation from it cannot be but intentional (e.g., downloading a hacked version of the protocol). In contrast, for higher layer resources, rewards are used instead of punishments (see also Fig. 3).

Network Layer
At the network layer, the main service is packet forwarding. Reciprocity mechanisms seem to require users to serve the same amount of traffic as the one they insert in the network. However, users have incomplete information concerning the nodes which forward their packets more than one hop away, which makes the enforcement of such a mechanism difficult to implement. Another problem with this approach is that nodes at the edge of the network will not be able to accumulate the necessary "credit," in order to satisfy their needs. A cross-layer approach can give chances to these edge nodes to provide other kinds of services in exchange, belonging possibly to different layers and ideally not depending on the network to be provided (e.g., original content, local services such as Internet access to mobile users, etc.). So, a user's contributions at the application layer can provide him with network services, like packet forwarding or connectivity (e.g., when mobile) or the opposite.

> **Freedom should be given to communities to self organize.**

An alternative approach is to give social incentives; in this case, it is important for such incentives to be positive rather than negative since we wish to encourage resource sharing rather than detect and punish misbehavior. This choice of perspective could play an important psychological role [23]. In this sense, users should be rewarded (instead of punished) according to their contribution in terms of social image and position in the community. But their quality as members of a social community is not necessarily correlated with their resource contributions at the network layer. This means that users should have a distinct status or reputation related to their behavior at the network layer. To this end, we propose the introduction of a new type of social relationship "my network neighbor" (and "my neighborhood community"), which will be valid only between physically connected users, and required for network setup. Then, additional incentives concerning the per-usage contributions could be applied in the context of these communities.

Application Layer
At this layer, depending on the application, users could be required to contribute different types of resources and/or services (from storage space or CPU cycles to Internet access or expertise). Thus, as in the case of the network layer, the ability of users to exchange different types of resources would increase the efficiency of the system [24]. Additionally, social incentives are again an attractive means to reward resource provision, treating it as a positive act.

The only difference for both approaches would be that now the interactions are not limited by the network topology (but they depend on it). Moreover, specific solutions are even more dependent on the specific environment and the corresponding resources entailed. As a result, freedom should be given to communities to self-organize and configure the corresponding mechanisms to suit their own needs.

Social layer
As proposed above, the social layer could play the role of a punishing/rewarding framework for encouraging the desirable behavior of users at the lower layers. However, additional incentive issues arise for the behavior of users at the social layer itself. To this end, all standard mechanisms discussed in the context of social software [18] apply. But, additionally, one should exploit the specific environment, especially for the design of the zero-level community, in order to further increase

their effectiveness. For example, the ability of users to physically meet should be taken into account and possibly formalized in the software design (e.g., provide the means to organize physical meetings, account for and reward the participation of users in these meetings, etc.). This ability should also be exploited to promote socializing, which is an important value generator in this context (e.g., socializing games exploiting the physical network topology such as random walks and visits in the neighborhood).

Additionally, the fact that people share their own resources, contributing to the community's operation, should be incorporated into the community's interface and social interactions. This fact could further strengthen the community spirit creating feelings of solidarity and independence towards the wired Internet infrastructure.

Communities and WMNs

Communities are an increasing guiding line for future networks. Their social nature reflects human communication needs, and they seem particularly appropriate as a basis for future user-centric communications, based on wireless technologies. In this context, communities span from the physical layer to the higher social layers, reflecting human interests, and creating self-organizing wireless mesh networks (WMNs).

Social relations will encourage users to participate in the formation of the envisioned neighborhood WMNs, and will generate the necessary information to ensure that the underlying network is formed among trusted and interested users. To this end, cross-layer incentive mechanisms are nevertheless required, using either reward or punishment strategies at different levels, reflecting the current user (node) behavior towards the community.

Author Information

Panayotis Antoniadis and Benedicte Le Grand are with the Université Pierre et Marie Curie - LIP6/CNRS, France. Anna Satsiou and Leandros Tassiulas are with the University of Thessaly & Informatics and Telematics Institute, Greece. Rui Aguiar, João Paulo Barraca, and Susana Sargento are with the Instituto de Telecomunicações, Universidade de Aveiro, Portugal.

Acknowledgment

This work was supported in part by the European Commission IST project WIP (no. 27402) and the national projects RNRT AIRNET (no. 01205) and PENED (no. 01ED540) PENED. We thank Marcelo Dias De Amorim for his help during the definition of our problem, and Frédéric Delvalle for the open page's original drawing. We are also grateful to the anonymous reviewers for their accurate and inspirational comments and suggestions.

References

[1] FON; http://en.fon.com/

[2] E.C. Efstathiou, P.A. Frangoudis and G.C. Polyzos, "Stimulating participation in wireless community metworks," in *Proc. IEEE INFOCOM 2006*, Spain, 2006.

[3] N. Negroponte, "Being wireless," *Wired Magazine*, vol. 10, no. 10, 2002.

[4] Microsoft Research, "Self-organizing neighborhood wireless mesh networks," Microsoft Corporation, 2008; http://research.microsoft.com/mesh/.

[5] A. Quan-Haase and B. Wellman, "How does the Internet affect social capital," in *Social Capital and Information Technology*, M. Huysman and V. Wulf, Eds. Cambridge, MA: M.I.T. Press, 2004, pp. 113-132.

[6] S. Frattasi, H. Fathi, A. Gimmler, F.H.P. Fitzek, and R. Prasad, "Designing socially robust 4G wireless services," *IEEE Technology & Society Mag.*, vol. 25, no. 2, pp. 51-64, 2006.

[7] A. Giannoulis, K.P. Tsoukatos, and L. Tassiulas. "Lightweight cross–layer control algorithms for fairness and energy efficiency in CDMA ad–hoc networks," in *Proc. WiOpt'06*, U.S.A., 2006.

[8] J. Huang, R.A. Berry, and M.L. Honig, "A game theoretic analysis of distributed power control for spread spectrum ad-hoc networks," in *Proc. IEEE Int. Symp. Information Theory*, Australia, 2005.

[9] J.P. Barraca, S. Sargento and R.L. Aguiar, "The polynomial-assisted ad-hoc charging protocol," in *Proc. 10th IEEE Symp. on Computers and Communications*, Spain, 2005.

[10] S. Buchegger and J.-Y. Le Boudec, "Performances analysis of the CONFIDANT protocol," in *Proc. of the 3rd ACM Int. Symp. on Mobile Ad-Hoc Networking and Computing (MobiHoc'02)*, Switzerland, 2002.

[11] M. Jakobsson, J.-P. Hubaux, and L. Buttyan, "A micro-payment scheme encouraging collaboration in multi-hop cellular networks, in *Proc. 7th Int. Conf. Financial Cryptography,* Guadeloupe, 2003.

[12] M. Feldman and J. Chuang. "Overcoming free-riding behaviour in peer-to-peer systems," in *ACM SIGecom Exchanges*, vol. 5, no. 4, pp. 41-50, 2005.

[13] P. Antoniadis, C. Courcoubetis, and R. Mason, "Comparing economic incentives in peer-to-peer networks," in *Proc. Special Issue on Network Economics*, *Computer Networks.* Elsevier, 2004.

[14] C. Dellarocas, "Reputation mechanisms," in *Handbook on Economics and Information Systems,* T. Hendershott, Ed., Elsevier, to be published.

[15] J. Mundinger and J.-Y. Le Boudec, "Reputation in self-organized communication systems and beyond," in Proc. Inter-Perf '06, Italy, 2006.

[16] R. Trivers, "The evolution of reciprocal altruism," *Quart. Rev. Biology*, vol. 46, pp. 35-57, 1971.

[17] D.W. McMillian and D.M. Chavis, "Sense of community: A definition and theory," *J. Community Psychology*, vol. 14, no. 1, pp. 6-23, 1986.

[18] T. Erickson, "Sustaining community – Incentive mechanisms in online systems," Final Report of the Group 2005 Workshop, U.S.A., 2005.

[19] E. Friedman and P. Resnick, "The social cost of cheap psedonyms," *Economics and Management Strategy*, vol. 10, no. 2, pp. 173-199, 2001.

[20] J. Douceur. The Sybil attack, in *Proc. 1st Int. Workshop on Peer-to-Peer Systems,* U.S.A., 2002.

[21] S.L. Keoh, E. Lupu, and M. Sloman. "PEACE: A policy-based establishment of ad-hoc communities," presented at Annual Computer Security Applications Conf. (ACSAC), Arizona, U.S.A., 2004.

[22] B. Yang and H. Garcia-Molina, "Designing a super-peer network," in *Proc. the 19th Int. Conf. on Data Engineering,* Bangalore, India, 2003.

[23] J. Andreoni, "Warm-glow versus cold-prickle: The effects of positive and negative framing on cooperation in experiments," in *Quart. J. Economics*, vol. 110, no. 1, pp. 1-21, 1995.

[24] A. Satsiou and L.Tassiulas, "A trust-based exchange framework for multiple services in P2P systems," in Proc. 7th IEEE Int. Conf. on P2P Computing, Ireland, 2007.