Ontology-driven Framework for Community Networking Management

João Paulo Barraca, Rui L. Aguiar
Instituto de Telecomunicações – Aveiro, Universidade de Aveiro, Portugal
jpbarraca@av.it.pt, ruilaa@ua.pt

***Abstract*—Wireless technologies present many benefits when used to deploy metropolitan networks. Resilience, cost, adaptability and self-organization are some of the benefits provided by Wireless Networks. Especially when mesh technologies are used, wireless networks are adequate for the deployment of (user-centric) community networks. The increased participation of users and the social relations existing among them makes desirable the merging of the mesh and social aspect of these networks. If proper interfaces are deployed, users are able to map existing (or to be developed) social structures such as communities into the mesh network. This paper focuses on the problem of creating a community driven mesh network, and in the representation and maintenance of user relations, trust and functional parameters. We present a novel ontology designed for the management of such environments.**

***Index Terms*—Social Networks, Mesh, Community Network, Ontology**

INTRODUCTION

Wireless Mesh networks (WMN) are being deployed in metropolitan areas, providing broadband access in a cost efficient manner. Although different scenarios exist for this concept, with the ongoing commoditization of multiple wireless technologies (such as WiFi and Wimax) the user will have an increasing role on the creation of such meshed networks. Instead of simply having scenarios where a building has a high bandwidth connection which is shared by a group of people (such as the ones living or working at it), new scenarios are appearing where users gain the ability of buying a router which can connect (directly) to the high bandwidth backhaul. Advanced routers can further extend the connectivity range of the network, allowing other users to connect by using them as relays. In the future also terminals (laptops, PDAs and even cell phones) will be enhanced with such advanced wireless mesh technologies. Although the range (as well as capacity and reliability) of terminals or, in general, SOHO (small office, home office) equipment is usually small, frequently they are effective in providing connectivity to several nearby apartments. Moreover, due to the low cost of such equipments the ratio between the number of equipments and area will be very high in urban areas, providing some basic (yet efficient enough) redundancy mechanisms.

In current mesh networks (and in general the Internet), the network is transparent to almost everything running on it, including user preferences. Users act as clients consuming a connection service provided by a network provider. The increasingly popular P2P models and social oriented applications foster new direct interactions between users, but these are framed only inside isolated applications: the network remains an uncontrolled pipe delivering packets independently of its origin, user or application.

For scenarios of high density mesh networks, where part of the equipment belongs to end-users, it would be natural to provide for a certain level of control to the equipment owners (e.g. the common bandwidth controls of P2P applications). This level of control is desirable due to the differences in user behaviour [1]. Each individual has a set of rules he wishes to apply in order to favour the shared applications he uses. Ultimately, in a meshed network without shared interests between the users, no communication agreement would be reached between users on the same network segment, resulting in the total absence of communication - which would negate the whole concept of user-centric community mesh networks. Thus user based network management comes at a price, and a balance is required to be established between user preferences and the overall community setup. Furthermore, since users are not typically tech savvy, managing all possible aspects of such a network infrastructure (even if a simple community) is too complex for many (most) persons.

The management interface between users and such community networks should thus make use of concepts, which are native to most (if not all) users. At this point, social concepts such as friends, neighbours and communities appear as possible natural candidates. Networks and (social/personal) clusters can be mapped into communities, friends to peers close-knit and neighbours (as in being in the same neighbourhood) to neighbours (as in being radio range). For any automated management system to operate on this environment, these concepts need to be formally represented.

In this paper we will present the challenges found in order to create an ontology representing the interactions in a social aware mesh network. With this work we further close the gap between the community social interactions and the network management of the associated community network. The final objective is to translate the possible social interactions between users into relevant network management functions.

This work has been partially supported by the European Commission project IST-WIP under contract 27402 and the Fundação para a Ciência e Tecnologia of the Portuguese Ministério da Ciência, Tecnologia e Ensino Superior under the fellowship SFRH/BD/32548/2006. The work presented reflects the vision shared by the authors but not necessarily the vision of the IST-WIP Project in its full extent.

978-1-4244-2036-0/08/$25.00 ©2008 IEEE

For such, we propose an ontology-based framework, which we briefly illustrate in a target metropolitan scenario.

In Section II we describe some of the most relevant work in the areas of wireless networking and the study of social communities. Section III overviews the definition of community and the deployment mechanisms inherent to these kinds of groups. The design challenges motivating the construction of our ontology framework are described on section IV. Section V describes relevant points of our unified ontology., while Section VI describes the network management framework currently driving our work. Finally in Section VII the conclusions of this work are presented, together with the future steps we plan to pursue.

I. Related Work

Much work has been developed in wireless mesh architectures and related concepts. Some proposals consider networks to be composed by homogeneous equipments using only a single technology connectivity medium (802.11) [2]. Others propose the use of two or three tier architecture providing more bandwidth [3]. Both types of proposals aim at creating a network providing a bulk communication pipe. QoS techniques, as proposed in [4], can supply some differentiation by applying the using traffic classes. Still user preferences have no direct impact on network management, only on the resulting performance metrics as a side effect of individual user applications. Moreover network configuration is supposed to be universally shared by all participants with little or none self-configuration capabilities [5]. Such concept may not be the most desirable due to the heterogeneous nature of devices and the environment. What is called "dynamic operation" on mesh networks is frequently a matter of actions on spectrum allocation, discovery, address self-configuration and dynamic routing mechanisms with no interaction between them.

The works presented in [6], [7], [8], [9], [10], [11] discuss some social mechanisms and characteristics of communities and were used as basic references for this work. In particular, the author of [6] presents key conclusions about interactions found on neighbourhood communities. Although the scenario focused does not deal with wireless mesh networks, it provides much insight related to neighbourhood and interest communities. Some of the most relevant conclusions are; 1) neighbourhood interaction is location based and contained within the neighbourhood, but not necessarily about the place. 2) Occasional encounters between members accumulate to form strong bonds between individuals and even form clusters. 3) Individual interactions change with time. First there is the need to interact more and participate in shared activities. As the individuals get more mature, their interests focus on management of the shared environment. 4) Only some individuals value cooperation and tools to increase the interactions between members of the community.

These conclusions present important opportunities and/or challenges for which should be addressed in community networking. Next we discus the key elements of such communities and the challenges (both social and technical) we identify as potentially important.

II. Communities

An important aspect requiring definition when designing a system supporting community based operation is the definition of the term community, and the identification of its basic elements. Community is a buzzword vastly used by companies to represent their user base or by the academia to represent a group of users which are closely-knit. In social terms, the definition of community is more problematic. One of the most accepted definition [8] points to the existence of universal aspects such as *membership*, *influence*, *integration and fulfilment of needs*, and *shared emotional connection* on a community. Also, basic properties such as *tolerance* [9], *reciprocity* [10] and *trust* [11] are frequently raised as important to communal operation. Interestingly the location factor is sometimes ignored, in particular since the advent of the Internet, as communal characteristics can be observed even if members are not in close physical proximity. Nevertheless, location is a factor relevant in our scenarios due to the limitations and characteristics of wireless connections. More important, clustering of nodes improves location of local services, increasing both the efficiency and scalability of the network. This is a basic characteristic of wireless mesh networks, making location a relevant element to consider.

When applied to wireless mesh networks, communities represent both a communication and an interaction contexts. That is, a common set of rules, permissions and behaviours, for users, services and the network layer itself. Rules may define the wireless channel or the maximum bandwidth to use. Permissions represent the eligibility of a given client to use a particular service or resource (e.g. trust based). Behaviour relates to the expected actions on particular events and may restrict the applications used, services supported or the action to take to punish others.

Figure 1 depicts several wireless communities on a typical city section as envisioned in a metropolitan scenario. Some communities have a local scope and only reach one or two buildings (composed by neighbours). Here, the location is vital to the definition of the community. Users not living on the neighbourhood will be unable to join the community and strangers will only be able to access public services. Services available on these communities, the knowledge it contains and the management functions will be highly related to the place. In terms of wireless networking, a neighbourhood is comprised of nodes sharing similar traffic and interference patterns.

Other communities connect users with similar interests across the existing wireless infrastructure. These are very similar to neighbourhood communities in terms of structure and management functions but the services available and its purpose are different. Also, these users may have network conditions very diverse making integration of the community context with the network a more complex task.

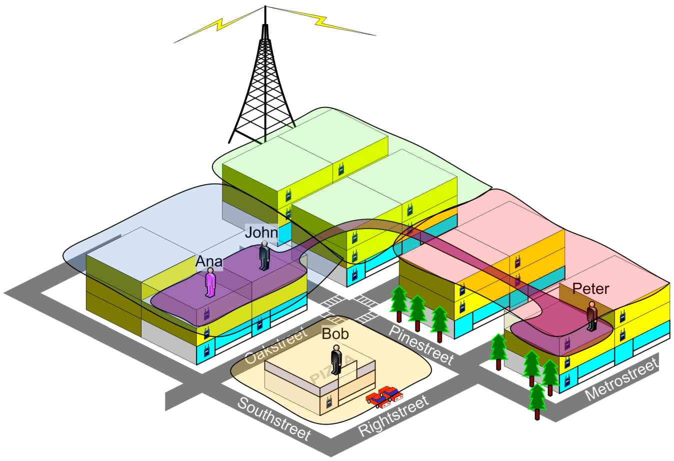

Figure 1 - Multiple communities in a typical city environment. Color blobs used to represent the communities

III. Design Challenges

There are many challenges to the integration of wireless mesh and social networks. Some are technology dependent and will be naturally solved as the technologies evolve, while others are related to operational aspects. Many other challenges appear at the social level, which, in this case, must also be considered.

Wireless technologies suffer from several limitations affecting the creation of a wireless mesh. Low resilience to interference, low bandwidth and low security (when compared to wired technologies) are problems well understood with many proposed solutions to minimize them. Interference robustness can be achieved by deploying proper monitoring mechanisms, and enhancing equipments with more advanced spectrum management functions. This aspect is actually benefited by self-organization inside neighbourhood communities, as closely located equipments will tend to be affected by similar interference sources. Security mechanisms should also be deployed to mitigate the problems created by a shared medium. However these mechanisms must be designed taking in consideration the trade-off between security, robustness and resources consumption. This is required so that the bandwidth available to user data is maximised, while providing a high level of security. Careful deployment of the wireless equipments, directional antennas and MIMO techniques, can increase the available bandwidth as well as connectivity range. However, in a wireless mesh network where users are able to add new wireless routers to the infrastructure, this process if far from being controlled.

Even if all technology driven challenges are understood, several other functional challenges arise. These are related to how the proposed framework architecture can be mapped to software and how it can operate in a way that is, at the same time, robust, pervasive, dynamically adaptable, and lightweight. Since community wireless mesh networks have no centralized policing system, reputation and trust has to be subjective. Also, it must be robust to unfaithful accusations and Sybil attacks.

Although technology issues related to ad-hoc and wireless mesh networks have suffered high scrutiny from the academia, key challenges still arise because of the users' behaviour. As the community base networks should take in consideration purpose, user preferences, location and expectations, an efficient, yet lightweight management platform, able to consider all these aspects must be developed.

A social challenge, which must be correctly addressed, is the motivation to participate in community. Early adopters can feel attracted to participate solely due to the innovation effect, and will surely help in the start-up process. The remaining potential users will only join a community if they have any need or interest which can be fulfilled by the existing community members. The same applies to influence, and other basic community elements. Keeping the number of users above the critical mass is vital and ultimately will depend on the members' cooperation and services available. While we cannot guess the services in use in a particular community, we surely can shape cooperation through the mechanisms deployed. In particular, as shown in [12], without punishment mechanisms, member cooperation rapidly ceases to exist. These results may seem a too dramatic but they reflect the complex nature of community operation.

All these challenges, as well as the basic community elements in section III, need to be taken in consideration when designing any management systems for community networks.

IV. Ontology for Community Based Mesh Networks

Previously we identified some of the requirements to social interaction as well as wireless mesh networks. For providing an automated management for community networks, some information models are required. Based on the discussions on the previous sections, we take the first steps in creating an ontology representing the structure of a community based wireless mesh network. In this section we will describe some of the particular aspects of the ontology, which will be further clarified when discussing the accompanying management framework in Section VI. The result of our ontology will be a formal data and relation model, which will serve as basis for the design and implementation of a community management framework.

An ontology promoting community based operation of a network should be designed so that it supports the identified elementary aspects and basic properties of a community environment. This way, network management will be more natural to the users, potentially also enhancing the interactions between them. From the previous analysis on the natural community processes, several functional requirements were derived. *Membership* can be stated as whether a user has access to community contents and services or not, thus imposing some mechanism to attribute the member status. *Influence* relates to the capacity of contributing to others decisions thus resulting in the capacity to shape the path of the community. This element is close to the element of integration and fulfilment of needs, as influence may increase the level of integration. These aspects combined impose the existence of a *distributed decision mechanisms* capable of giving members the capacity to influence others behaviour. The shared emotional context can only be directly applied to users and is the well known sense of belonging present in most groups. In our case this context is generalized by representing it is represented as a common communication and knowledge context.

The basic properties of *tolerance*, *reciprocity* and *trust* further impose the need for an efficient reputation system, tightly coupled with the other community processes. So, we consider the existence of a *reputation system* enabling users and services to estimate others' behaviour. This system is cross-layer, calculating reputation results at the different layers of the stack by receiving transaction proofs. Eventually also the user is able to provide information about the result of a transaction, thus rating the peer.

We also assume the existence of a *role based management* system tightly integrated with the reputation system. With this approach we envision to achieve a better mapping of the social mechanisms into the network management infrastructure. According to this assumption, roles are delegated to peers based on the services they provide, their reputation, location and equipment being used.

In terms of the data model, the community itself is a container for entities united under a single doctrine. *Doctrine* represents the collection of all knowledge shared by the members of the community. It should be noticed that there must be at least one entity and one doctrine in every community. In reality a community must have more than one participant; however, in its initial stage, upon its creation, this may become untrue.

We represented these concepts using OWL, given its adequacy for this representation. *Code* 1 illustrates the representation format used, with the Community concept. Naturally we had to define the identifier, comment, purpose and location fields, to store some additional information about the community – and not all of this and may be available to foreign users. This information may be either broadcasted to the local network or published in a known directory server. Using tags or keywords on the purpose field is highly recommended as it helps user systems to rapidly find the most appropriate community to join if several are available.

Entities (*Code 2)* are a composed concept grouping one or more users, equipments and services. Such a broad concept is required to correctly represent the way we users participate in a community. When users interact on a community, participation is not independent of other factors. Participation is always conditioned by the services provided and the equipments used (also location and many other factors). The combination of all these aspects is what other members of the community will see. In other words, an entity is a persona (or avatar) grouping a user, its resources and services available to the community. While users communicate with other users, or access specific resources, this enables trust and reputation to be also maintained on an entity basis. So, entity reputation reflects the reputation (weighted average) of the member users, equipments and services. In this aspect an entity can range from a user and its laptop to an entire corporation or operator.

```
<owl:Class rdf:about="#Community">

  <owl:DatatypeProperty    rdf:ID="#identifier" />
  <owl:DatatypeProperty    rdf:ID="#purpose" />
  <owl:DatatypeProperty    rdf:ID="#comment" />
  <owl:DatatypeProperty    rdf:ID="#location" />

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasEntity"/>
        <owl:minCardinality
          rdf:datatype="&xsd;int">1</owl:minCardinality>
    </owl:Restriction>
  </rdfs:subClassOf>

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasDoctrine"/>
        <owl:cardinality
          rdf:datatype="&xsd;int">1</owl:cardinality>
    </owl:Restriction>
  </rdfs:subClassOf>

...

</owl:Class>
```

Code 1 - Example definition of the community concept using OWL

```
<owl:Class rdf:about="#Entity">

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasUser"/>
        <owl:minCardinality
          rdf:datatype="&xsd;int">1</owl:minCardinality>
    </owl:Restriction>
  </rdfs:subClassOf>

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasEquipment"/>
        <owl:minCardinality
          rdf:datatype="&xsd;int">1</owl:minCardinality>
    </owl:Restriction>
  </rdfs:subClassOf>

...
</owl:Class>
```

Code 2 - Example definition of the Entity concept using OWL

```
<owl:Class rdf:about="#AccessPoint">
  <rdfs:subClassOf rdf:resource="#Equipment"/>

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasPower"/>
      <owl:hasValue rdf:resource="#powerAC"/>
    </owl:Restriction>
  </rdfs:subClassOf>

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasService"/>
      <owl:hasValue rdf:resource="#serviceForward"/>
  </owl:Restriction>
  </rdfs:subClassOf>

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasInterface"/>
      <owl:hasValue rdf:resource="#interfaceWIFI"/>
    </owl:Restriction>
  </rdfs:subClassOf>

  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasInterface"/>
      <owl:hasValue rdf:resource="#interfaceEthernet"/>
    </owl:Restriction>
  </rdfs:subClassOf>

...

</owl:Class>
```

Code 3 - Example an instance of the Equipment class describing an Access Point in OWL.

It should be noticed that assessing the reputation of a given service or resource is still possible as reputation should be kept by transaction. Also, there is no possible interaction capable of affecting entity reputations directly as entities do not interact due to its abstract nature.

Equipments are composed by several hardware devices (either real or virtual) with particular characteristics. The description of all devices constitutes the equipment profile and allows others nodes to infer the capabilities of a particular equipment. Later it can be used if some role presents a hardware requirement as a restriction. For instance, membership of a community created to provide backhaul wireless communications could require all equipments to have at least one WiMAX interface.

One example instantiation of the Equipment class is the one of an AccessPoint (*Code 3*), which by obvious reasons, must only be applied to equipments, and only to those with, at least one Wifi and one Ethernet interfaces. Also, it requires the device to not rely on batteries and to support the forwarding service.

As there is no method to actually verify the correctness of the entire hardware profile presented, users will be able to forge any profile they want in order to achieve particular roles. This can potentially be avoided using some trusted hardware platform, which is becoming increasingly available.

The knowledge space, represented as the doctrine, is used by the entities to define communication guidelines and to self organize. These guidelines take the form of roles containing rules and attributes. Peers request a role and, if the requirements are met, gain the designated permissions but must also comply with the expected behaviour. Rules and attributes define the reaction to determined events and the configurations to apply in each circumstance. In this situation, membership is attributed by delegating the Member role. All objects, such as identity information and public keys are also part of this same knowledge repository, ultimately represented by the doctrine.

V. Community Management Framework

Due to the complexity of managing such environments, automated systems are required, and thus formal representations need to be developed. All the objects (services, users, peers, roles, delegations…) instantiated for management purposes need to be represented in a language such as XML, having a structure that has to be defined by an ontology, to support the potential complexity of existing relationships. Objects are particular to a given community and are permanently available to every participant, according to its permissions. Together with the architecture proposed, we envision to provide the required functionalities enabling communal behaviour. This architecture follows a modular approach with several interfaces and a network abstraction layer. The software components composing the meshed community architecture are depicted in *Figure 2*.

A. Community Management

In social terms, the Community Management module manages the common community context as represented by the doctrine. Initially it takes user provided or a default configuration and bootstraps the initial doctrine. The doctrine then evolves by adding and changing delegations independently of this module. Other important function is support for merge and split events. Split events occur whenever a group of members of the community is unable to contact the remaining members. Community operation suffers little from this event as the sub-communities may have the capability of managing their members. The doctrines of the two partitions evolve independently, eventually forming two different communities with completely different doctrines.

Merging occurs when the disconnected nodes finally join the community and the several partitions finally merge. Other possibility is when two different communities whish to merge under a single doctrine. There are several issues with the synchronization of the two sub-doctrines. According to the ontology defined, all the objects created must be inspected and duplicates must be removed. Moreover, some objects may define conflicting parameters, which will require elimination of one.

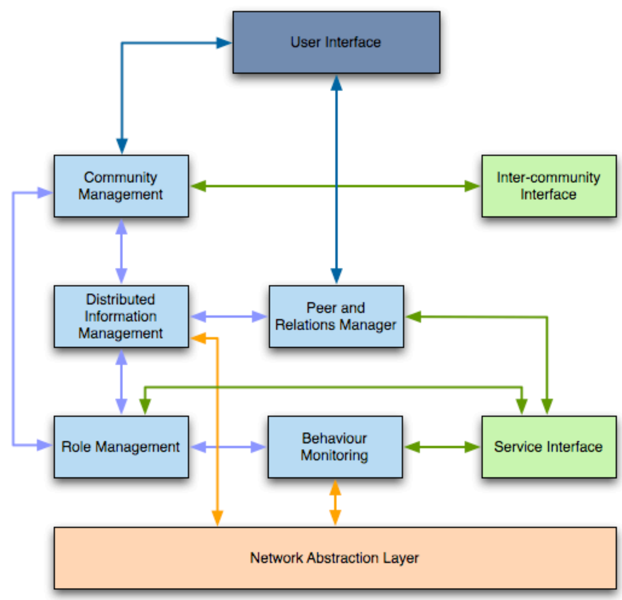

Figure 2 - Software architecture of the community framework

B. Peer and Relations Management

The essential function is to actively monitor a number of peers on the network, such as neighbours and known peers. Information such as their availability, uptime and round trip delay is present on a peer profile. One important particularity is that the profile is also available to others to consult, as it is stored on the *Distributed Information Management* module. This way, any peer on the system can easily assess what others think of a particular peer.

From all peers known, the one with best performance metrics (preferable a physical neighbour) is defined as the attachment peer. Others are kept as backup. By keeping in touch with these peers, even if a group is not contiguously connected to the rest of the community, communication will still be maintained using the delegated peers to route messages to the remaining nodes.

Besides performance metrics, the reputation of a peer is maintained on the profile. Also, this module is responsible for calculating the reputation based on result tokens sent by other modules of the community framework. Every time a transaction occurs, the result is analysed (by the service, user or *Behaviour Monitoring* module) against the rules active, and a reputation value is derived. One example can be the transfer of a music file, leading to an increase in the peers' reputation in case the file is the expected one. If it is not or is corrupted, the reputation will be decreased. For many mechanisms, at network and application layers, reputation values can be derived automatically, while for higher layers, users are required to provide some feedback on the result. This is required because it is difficult for modules, monitoring packet transfers, to assess the correctness of a high level transaction.

C. Role Management

Communities are typically not composed by members with equal rights and responsibilities. This happens due to scalability issues related to the resources spent taking decisions and managing the community information. Even if considering small communities such as private forums or any other social community, there is frequently the need for some delegation of responsibilities on a restricted number of participants. Other aspect is that as the capabilities of each equipment, radio interference experienced and location are different between the participants, some may potentially adopt different roles because they are more suited to them. In the case of a wireless mesh network some participants may operate as routers, others provide gateway functionalities to the Internet, some others provide storage space, or music streaming. Following this approach, all participants of a community are entitled to one or more roles defining their behaviour, configuration parameters and permissions. The delegation of a role will depend on the characteristics of the participant (services provided, resources) and is handled by the Role Management module.

D. Behaviour Monitoring

Peers are expected to behave according to the roles they have. The result is they must act in respect to a set of rules benefiting the community functionality, instead of their own. Roles are dynamic, enabling the community doctrine to react to environment changes if a peer requests it. However, not all peers will benefit from the active doctrine as it focuses on maximizing operation of the all, not the one.

Under these conditions, it is natural to assume some peers will be tempted to manipulate the doctrine (by proposing changes) in order to maximise their utility function. If this becomes impossible, because other peers refuse the proposed configuration, they may simply not follow the configurations imposed by its roles. The *Behaviour Monitoring* module observes the interactions between peers and checks if they seem to be violating the role parameters. In this case, both interactions with the local peer and with peers in the neighbourhood can be monitored. However, due to the long known effects of the wireless medium, and the additional processing capabilities required, promiscuous monitoring should not be used on low capacity, or battery dependent peers. The information collected is provided to *the Peer and Relations Management* module, resulting in a change on a peers' reputation.

E. Interfaces

The several interface modules (User, Community, Services and Network) provide means of the community modules to communicate with the several layers it impacts.

The User interface presents an interface where users can configure the initial parameters of the doctrine, input their preferences, apply for roles and provide feedback from previous transactions. This feedback will then be taken in consideration when calculating the reputation of a peer.

Services communicate with the community framework in order to read configuration parameters, verify permissions and provide feedback from the interactions. Like in the case of the User interface, the feedback provided is used to update the reputation of a peer.

Several technologies can be used in the communication between peers, like Wimax, Wifi or Bluetooth. In order to abstract the underlying communication medium to the community framework, a Network Abstraction Layer provides a common interface to all the technologies. This module also supports encapsulation and decapsulation of data packets. Through this operation it is possible to create overlay networks and route packets between nodes not on the same physical network.

F. Distributed Information Management

One important aspect of a community is the distribution of information between the participating members. All members must have an updated version of the doctrine so that they can apply its configurations or monitor the behaviour of the remaining peers. However, every time a member joins the community, new delegations must be issued and distributed to all other nodes. The overhead of replicating the doctrine on all members is overkill and certainly would limit the scalability of the architecture.

The *Distributed Information Management* module implements a P2P based repository of objects permanently available to others. Caching and automatic replication are used to improve the performance of the repository and reduce query latency. Each object is represented in the distributed storage as a main object and several replicas. All peers can read the replicas, which are cached automatically as a form of redundancy, but only the master node for a particular object can write it. The determination of the master node is found using a DHT system, which in its current stage is based on the structured DHT implementation named Chimera [13].

VI. Conclusion

We presented the challenges encountered when uniting the local and network layers in the scenario of a metropolitan wireless mesh network. Based on this work an ontology was described together with a software architecture making use of the proposed ontology. With this work, we envision that a more user-friendly management paradigm for wireless mesh networks will be obtained. Also, by closing the gap between social and wireless networking, we create the roots for a novel interaction model, where social ties are present on the fabric of the network. Future work will rely on further increasing the capabilities of our ontology, increase the understanding on community interactions, and identify its benefits (and short comings) to the management of wireless mesh networks. Furthermore we plan to finish the ongoing prototype implementation of the framework and evaluate it.

References

[1] Panayotis Antoniadis, Benedicte Le Grand, Anna Satsiou, Leandros Tassiulas, Rui Aguiar, João Paulo Barraca, Susana Sargento, *Community Building over Neighbourhood Wireless Mesh Networks*, IEEE Technology and Society Magazine, special issue on the Potentials and Limits of Cooperation in Wireless Communications: Toward Fourth Generation (4G) Wireless, March 2008 (to be published).

[2] Audeh M., *Metropolitan-scale Wi-Fi mesh networks*, IEEE Computer, Vol 37, Issue 12, pp 119-121, ISSN: 0018-9162, Dec 2004,

[3] Raheleh B. Dilmaghani, B.S. Manoj, Babak Jafarian, Ramesh R. Rao, *Performance Evaluation of RescueMesh: A Metro-scale Hybrid Wireless Network*, Proceedings of WiMesh Workshop, IEEE Workshop on Wireless Mesh Networks, 2005

[4] Jiang H., Zhuang W, Shen X, Adbrabou A., Wang P, *Differentiated services for wireless mesh backbone*, IEEE Communications Magazine, Vol 44, Issue 7, pp 113-119

[5] Akyildiz I.F., Xudong Wang, *A survey on wireless mesh networks,* IEEE Communications Magazine, Vol 43, Issue 9, pp 23-30, ISSN: 0163-6804

[6] Marcus Foth, *Facilitating Social Networking in Inner-City Neighborhoods*, IEEE Computer, Vol. 39, Issue 9, pp 44-50.

[7] Samuel Bowles, Herber Gintis *Social capital and community governance,* Economic Journal, Royal Economic Society, vol 112(483), pp 419-436.

[8] McMillan D.W., Chavis D.M., *Sense of Community: A definition and Theory*, American Journal of Community Psychology, Vol 14, Issue 1, pp 6-23

[9] Michael Walzer, *On Toleration,* New Haven: Yale University Press

[10] Robert D. Putnam, *bowling Alone: The Collapse and Revival of the American Community*, New York: Simon & Schuster, ISBN: 0-7432-0306-6

[11] James Coleman, *Foundations of Social Theory,* Belknap Press, ISBN-978-0674312265

[12] Ernst Fehr, Simon Gachter, *Cooperation and Punishment*, American Economic Review, Vol 90, Issue 4, pp 980-994

[13] Chimera, A Structure Peer-to-Peer Overlay, http://current.cs.ucsb.edu/projects/chimera, as in January 2008.