User-Provided Networks: Consumer as Provider

Rute Sofia and Paulo Mendes, INESC Porto

ABSTRACT

This article describes and characterizes an emerging type of user-centric wireless network model, here named a user-provided network, where the end user is at the same time a consumer and a provider of Internet access. A discussion of challenges these new models face is also provided.

INTRODUCTION

Internet services and models have been going through a paradigm shift due to three main factors: widespread wireless technologies, an increasing variety of user-friendly and multimedia-enabled terminals, and the availability of open source tools for content generation. These factors have been empowering the end user with a new role, that of *micro-provider*,[1] and thus the end user is, at the same time, a provider and a consumer. Specifically addressing Internet access (*connectivity*), *user-provided networks* are the ultimate example of the application of micro-provider roles. Wireless fidelity (Wi-Fi) is shared freely and transparently among end users in a way that is technically and legally independent of access or infrastructure providers. Examples of such networks currently exist both commercially (e.g., FON [1], OpenSpark [2], Whisher [3]) and non-commercially (e.g., providing Internet access to a few wireless devices at home without having to set up a wireless router). These networks are here called *user-provided*, given that they "spread" by means of the end-user willingness to share connectivity. It is within this context that this article focuses, namely on the role and impact user-provided connectivity models may have on Internet architectures.

It should be highlighted that the recent wave of user-centric wireless models is a consequence of a paradigm shift in Internet services that started with the peer-to-peer (P2P) model: P2P was the first approach that, from a *cooperative* perspective, empowered end users as active Internet service providers. The networking perspective we cover in this article is a more recent paradigm, and due to its embryonic stage of development there is currently no clear understanding of either the technical or commercial potential user-provided networks may hold. This article offers a first exploration of the potential of these models by explaining what is currently available and how it may evolve.

The article is organized as follows. The next section goes over related work, while the following section provides a characterization of user-provided networks and their main properties. We then discuss challenges user-provided networks face. We then conclude.

RELATED WORK

AUTONOMOUS NETWORKS

In recent years, special focus has been put on systems and mechanisms that allow networks to self-organize and to automatically establish connectivity among involved entities, in order to accommodate future service needs. This is the core belief of BIONETS [4], an EU-funded project that gets inspiration from biological models. BIONETS aims to develop truly user-centric Internet models, allowing networks to naturally evolve and become autonomous, accommodating new services and societal needs.

HAGGLE [5] approaches the user-centric and autonomic perspective going beyond current network paradigms by exploring *application-driven* (opportunistic) message forwarding, as well as the impact of human communication on the network. SOCIALNETS [6] follows the same line of thought but aims to model networks according to human trust behavior.

The above projects relate to autonomic networks but do not take into consideration all the technical and commercial potential users playing the role of providers of Internet access may have.

INTERNET WHOLESALE MODELS

Telecommunication markets worldwide are witnessing a strong evolution toward full liberalization. The motivation for such movement is to foment competition between the different telecommunication players, network access providers (NAPs) and service providers (SPs). This should in turn trigger the development of new services and accelerate overall price reduction.

The traditional wholesale model is the *bundled model*, where normally an (incumbent) NAP also incorporates the role of SP. For this model, it is up to the NAP to keep the business relationship with an end user, thus restricting the choice of SPs.

[1] *Micro-provider is here employed in terms of both dimension and functionality provided, mostly related to Internet access.*

0163-6804/08/$25.00 © 2008 IEEE

This panorama changed in 1999 when, in the United States, the FCC regulated that incumbent NAPs would have to share their lines with regional competitors at a certain preset wholesale price, giving rise to a new model, the *unbundled model*. In the unbundled model the NAP simply owns the infrastructure, with SPs being separate entities: the end user must still keep subscription agreements related to both the access and the services being provided. The main advantage of this model is that the end user can freely choose between SPs.

The third and current wholesale model is the *SP-centric model*, where SPs have a direct relationship with end users concerning Internet access. The SP-centric model gives an end user the possibility to bypass NAP limitations and keep an agreement with an SP only. These SPs may have different requirements in terms of NAP; nonetheless, the way traffic is routed on the Internet does not take into consideration such requirements (e.g., multihoming). A potential solution [7] is to create a virtualized Internet architecture that allows SPs to lease physical resources on demand from different NAPs. Although this brings in more flexibility, it still restricts users' choice of specific SPs. A step further into the virtualization of the Internet architecture is the role of *virtual operator* of which the most common example today is a mobile network operator (MNO). A virtual operator is an entity that provides some form of service (e.g., connectivity coordination) but does not have its own infrastructure to provide the service.

User-provided networks go a step further concerning the notion of virtual operator, integrating the notion of micro-provider: the network is "spread" by means of the end-user willingness to share his/her subscribed Internet access and management is decentralized, or there is a central coordinator (virtual operator) in charge of management.

Today's Related Models

From a commercial perspective, FON and OpenSpark are commercial examples of user-provided networks. Albeit tremendously successful, their model faces some drawbacks as debated in the next section. Whisher goes a step further into user-centric models by placing all the functionality in software integrated into the end-user device. In contrast to the model embodied by FON and OpenSpark, Whisher allows dynamic dispersion of Internet access points, since sharing connectivity points are mobile.

Pocket switched networks (PSNs) [5] are described as mobile networks providing isolated connectivity graphs (*islands of connectivity*). PSN mechanisms are based on the pragmatic fact that today end users have, for a specific period of time and due to a specific schedule, intermittent connectivity between different islands of connectivity. Handling of intermittent connectivity is important in user-provided networks, since they also present connectivity islands.

Multihop networks are an instantiation of ad hoc networks where nodes communicate over multiple sequential links. Multihop networks normally integrate a high number of ad hoc devices, but of these only a small and fix amount behave as gateways to the Internet. Consequently, communication in large multihop networks is proven not to be efficient for real-time communications. User-provided networks support the idea that efficient sharing of available frequencies and bandwidth as well as integration of adequate incentives to cooperate will allow users to be better connected to the Internet.

User-provided networks support the idea that an efficient sharing of available frequencies and bandwidth as well as the integration of adequate incentives to cooperate will allow users to be better connected to the Internet.

Concept Overview

This section provides a characterization of user-provided networks, starting with models that are currently in use. To clarify the main differences from other autonomic networks, this section also provides a comparison of connectivity features for user-provided networks against ad hoc and multihop.

Figure 1 contains several examples of user-provided models that either are in use or can easily be deployed with current mechanisms, and are discussed in the next sections. We highlight that the equipment being represented as Wi-Fi access points can be any other such device (e.g., a WiMAX base station). A *micro-provider* (i.e., an individual or a community of individuals) holds one or more Internet access subscriptions and is willing to share connectivity.

Direct Sharing Model

The two topmost diagrams in Fig. 1 relate to models that are in use today. The *direct sharing* model (top left) relates to the model that is the basis of FON and OpenSpark, where Internet access is shared by means of an access point that provides radio connectivity publicly. In the figure A and B are two users with no trust relation whatsoever. Let us here assume that A owns the access point. Such an access point is configured in a way that allows A to securely access his/her subscribed services (e.g., private SSID, data protection by IEEE 802.1x/EAP) and at the same time provides open access (public SSID, free access) to any user belonging to a specific community. While A may be at home, B is roaming on the street and within range of the access point owned by A.

This model implies changes to the access point, and the growth of the network is directly related to the number and location of shared access points, but not on the microprovider's ability to move. The main advantage of this model is that even if A's equipment is shut down, connectivity can be shared as long as A's access point remains active.

User-Centric Relaying Model

The *user-centric relaying* model (top right, Fig. 1) moves the connectivity sharing capability into the end-user device: in the figure C is the element that is sharing connectivity, and A and B are profiting from such sharing. The represented access point may or may not directly belong to C. For instance, it can be a wireless router within C's home or simply a regular access point in a cafe where C bought prepaid Internet access.

In this scenario the end user becomes a micro-provider based on software functionality only. Relaying to the community is then performed by means of an ad hoc connection estab-

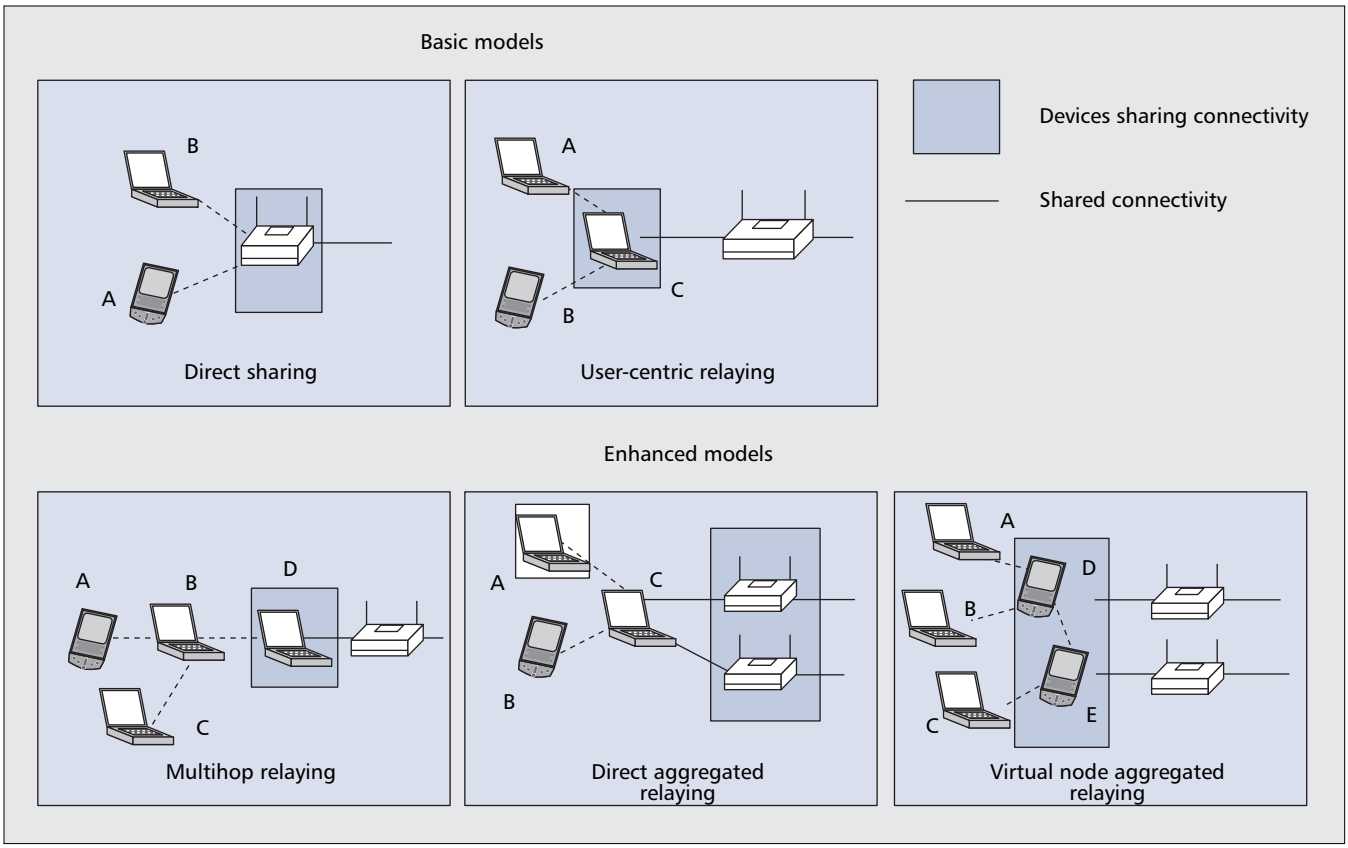

Figure 1. *Examples of user-provided connectivity models.*

lished between the interested elements. This model is the one followed by Whisher. In contrast to the direct sharing model, this one requires neither changes nor hardware acquisition, given that access nodes are not directly involved. To provide a specific example, assuming that C is in a coffee shop and using prepaid Internet access, C can still share this subscription. From a coverage and network spreading perspective, this model is more dynamic than the direct sharing model, given that the network changes according to micro-providers' mobility patterns.

Enhanced Models

Variations of the two basic models can easily be implemented today. For instance, a model that relies on multihop relaying (bottom left, Fig. 1) allows broader sharing, given that it can spread more than one hop away from the micro-provider. Variations may also be implemented in order to optimize connectivity sharing and relaying. In other words, assuming that within a scoped range a group of micro-providers is willing to share connectivity, variations can consider load balancing to specific access points (bottom center) or take advantage of aggregated backhaul capacity multiple micro-providers are offering (bottom right).

Characterization of User-Provided Networks

Although there are already commercial examples of user-provided networks, this concept is still in an embryonic stage. Hence, it is necessary to define the essential features to integrate into future models, assuming that user-provided networks may rely on any form of radio technology. We argue that user-provided networks rely on four basic properties: *connectivity sharing and relaying*, *cooperation*, *trust*, *and self-organization*, further elaborated on next.

Connectivity Sharing and Relaying — The first feature considered is the willingness of end users to openly share/relay connectivity, that is, the willingness of end users to become micro-providers within autonomic communities. In order to ensure widespread usage, connectivity relaying has to be deployed based on software or hardware that is available or easily deployable in end-user equipment, independent of radio technology, operating system, and device. Connectivity relay can be performed either actively or passively: the end user owning the device may or may not be aware that he/she is relaying connectivity. However, in either case it is necessary to ensure that the end user allows such relaying to happen. Furthermore, user-provided networks naturally follow human living patterns and thus have a higher probability of providing better coverage in densely populated areas. Connectivity islands will abound, and in order for user-provided models to be successful it is necessary to consider support for *intermittent connectivity*, in order to ensure reliable service even when nodes sharing their connectivity are switching across different Internet access locations.

Model	Connectivity sharing and relaying	Cooperation	Trust	Self-organization
User-provided	Self-organized Better in densely populated areas Active or passive	Incentive-based Reward-based Network of trust	Social interaction Reputation mechanisms	Human nomadic lifestyle Follows human mobility patterns
Ad hoc	None	Broadcast-based	None	Proximity of devices
Multihop	Routing-based	Routing-based	Routing functionality	Proximity Routing metrics

Table 1. *Comparison of main connectivity properties for user-provided, ad hoc, and multihop networks.*

Cooperation — A second main feature of user-provided networks is cooperation. Users must rely on some form of incentive to share connectivity, be it trust boundaries or some form of compensation (e.g., wider Internet connectivity). Cooperation mechanisms should ensure that there are both incentives and rewards for "good behavior," thus reducing the probability of malicious users entering the network. Cooperation aspects also relate to mechanisms that may apply in order to improve usage of shared resources. Concerning current models (discussed above), such an incentive is simply being able to take advantage of broader and free coverage. In the future incentives could relate to the amount of time (or throughput) that is shared by a specific end user: the more (service) the end user shares, the more the end user can benefit.

Trust — Trust is the third pillar, and has strong impact on cooperation as well as connectivity relaying. User-provided networks must consider social interaction and human interests as the basis for building trust. In addition, reputation mechanisms (e.g., similar to the eBay model) should be integrated. With these mechanisms as well as incentives to behave adequately, misbehavior will be reduced. *Good* behavior can be rewarded, for example, by providing more connectivity to end users that behave according to some pre-established criteria. It may also be necessary to implement adequate misbehavior detection and rule-out mechanisms to avoid network operation disruption.

Self-Organization — Self-organization, the fourth pillar, relates to the autonomic facet of user-provided networks: the ability to coordinate connectivity management in a decentralized manner. Moreover, it is not possible to predict how many users will be hanging from a specific micro-provider at a time; nor is it possible to predict the dynamic morphology of a specific user-provided network and how it will spread. Spreading can be characterized by the number of nodes sharing (or profiting from) connectivity, the way nodes interconnect (e.g., node degree), the average number of end-to-end hops, as well as the average bandwidth of which nodes can take advantage. The resulting topologies are therefore strongly dependent on human behavior, particularly *human mobility patterns* and *social interaction* models. Moreover, given that it is not possible to predict the growth and consequent network load associated with a specific micro-provider, techniques that allow the community to take advantage of aggregated backhaul capacity are crucial to allow adequate expansion of the network.

Comparison of Similar Approaches: Ad Hoc, Multihop, User-Provided

Table 1 provides a side-by-side comparison of the properties previously described concerning user-provided, ad hoc, and multihop networks.

Starting with the first property of user-provided networks, connectivity sharing and relaying is based on the willingness of some end users to share Internet connectivity, the willingness of some users to find micro-providers, and specific and individual user preferences (e.g., connectivity cost or throughput). Connectivity is therefore naturally better in densely populated areas, and it may spread by means of active hops (end users explicitly providing sharing) or passively (end users willing to relay connectivity that is shared by someone else). Ad hoc, on the other hand, does not integrate relaying. As for multihop, connectivity is based on a common routing mechanism. It is important to stress that while routing is mandatory in all devices participating in a multihop network, we assume that some or all of the devices in a user-provided network do not integrate any routing functionality.

Cooperation in user-provided networks follows human trust behavior. In contrast, pure ad hoc networks simply "cooperate" based on broadcasts, and multihop networks cooperate by implementing mandatory routing functionality. Cooperation heavily affects the way user-provided networks expand, which is not so for ad hoc.

Trust management heavily depends on human behavior for the case of user-provided networks: these will not work without adequate trust management. In contrast, there is no integrated trust mechanism in ad hoc networks, and multihop networks may integrate some form of trust based on routing functionality. Actually, the lack of incentives to forward third-party packets has always been a deployment limitation for multihop networks.

The fourth property, self-organization, is common to all autonomic systems, but in user-provided networks it completely depends on the nomadic lifestyle of users. In ad hoc networks such self-organization is based on proximity of devices, and for multihop networks it is based on routing metrics.

Being an embryonic concept, user-provided networks face several challenges. Short-term challenges relate to the need to implement some measures immediately in order to allow user-provided networks to expand adequately. Long-term, there are structural challenges that must be addressed.

CHALLENGES

Being an embryonic concept, user-provided networks face several challenges we categorize in terms of their impact scope. Short-term challenges relate to the need to implement some measures immediately in order to allow user-provided networks to expand adequately. From a long-term perspective, there are structural challenges that have to be addressed and whose main outcome is a paradigm shift in communication models related to the Internet.

SHORT-TERM CHALLENGES

End-User Accountability and Identification — The micro-provider's right to share its Internet access subscription has to be defined in terms of the established broadband subscription agreement. Such an agreement normally does not contain rules about Internet access sharing. However, it may include conditions restricting a customer's right to use his/her subscribed service upon the occurrence of infringements related to traffic volume or network load threshold. Given that the number of simultaneous users connected to a micro-provider is simply limited by the throughput provided, the micro-provider may, even without knowing and due to the community that is profiting from the micro-provider's subscribed access, be violating his/her agreement obligations. It is therefore necessary to consider ways to ensure that traffic can be adequately load balanced across micro-providers willing to share connectivity, to prevent network load surges that may endanger a microprovider's Internet access subscription.

Concerning end-user identification, currently they are normally identified within a community by, say, a set of credentials provided for a single registration. Connectivity is normally provided by means of a captive portal tool. Therefore, upon detection of malicious behavior the user will lose his/her access to the global community.

A far more complex problem is accountability of end users. In user-provided networks, from an operator perspective, only the micro-provider is accountable, particularly for models that evolve from the *user-centric relaying* model, given that the sharing functionality is completely beyond operator control. One could envision the entity that coordinates the identification of users within a specific community also being responsible for accounting. This is feasible as long as such an entity exists, as is common practice today.

Securing the Connection: Data Privacy — One of the major challenges in user-provided networks is the fact that, adding to the wireless media, traffic from users is to be relayed. Consequently, end users' misbehavior will be tragic for user-provided models. The attacks that may have more significant and negative impacts on user-provided networks are *incorrect traffic relaying attacks* as well as *impersonation* attacks. To fight back these threats, we argue that user-provided networks do not require tight security mechanisms but instead must integrate innovative trust management schemes, that follow human trust behavior. Under this category we place incentive and reputation schemes. Reputation mechanisms should be considered to monitor past actions of nodes and connectivity sharing points, information that may be used to decide on packet relaying and selection of connectivity shared points. Another approach to deal with security in user-provided networks is through the implementation of credit mechanisms; for example, connectivity credits can be offered to nodes sharing their Internet access points and nodes relaying data within connectivity areas.

Impact on Telecommunications Market Legislation — User-centric networking models have strong implications in terms of telecommunications market legislation. To give a concrete example, Finland seems to be the first country to legislate the positioning of user-provided basic models such as those embodied by OpenSpark and FON. In [8] these models are labeled "Web communities" where the functioning is based on mutual provision of networks and a micro-provider shares his/her access based on conditions defined by a network *coordinator* (e.g., another micro-provider). The Finnish Communications Regulatory Authority [8] already acknowledges that these networks are not covered by current wholesale models, given that the micro-provider may also be a simple end user profiting from shared connectivity, and there is no clear service or task splitting, in contrast to the regular Internet wholesale models.

This definition always assumes that there is a coordinator in charge of defining rules for a specific community. It does not foresee, however, the need to define models that are more autonomic, where the coordinator role may be distributed or networks simply self-organize according to incentive schemes to ensure smooth operation.

LONG-TERM CHALLENGES

Building Networks of Trust — User-provided networks are highly dependent on mechanisms capable of quickly developing a network of trust. This is a hard task to achieve, particularly given the nomadic lifestyle of end users and the difficulty in providing detailed accountability.

As already debated, there is the need to provide decentralized support that gives a community the possibility to grow dynamically. Consequently, *grassroots* trust mechanisms are the basis on which to build future user-provided networks. Web of trust (WOT) schemes are one possibility, given that they offer the means to rely on widespread cryptography tools (e.g., PGP) to develop networks of trust. Currently, trust within a specific user-provided network is confined to a specific community. In the future, trust models should not only consider community beliefs but actually depend on surroundings, as well as the level of confidentiality the user expects at a specific moment and for specific applications. Hence, the most adequate trust management models to consider are global ones, where each peer holds specific trust values and metrics other peers can access. In addition, systems that aim to fight back selfishness of peers (fight back the *tragedy of the commons*) have to be considered. The challenges faced by user-provided networks in terms of trust management

are very similar to those faced by virtual communities; consequently, mechanisms applicable to trust management in such communities should be analyzed as a starting point to build adequate networks of trust.

Internet Architectures — One of the key architectural guidelines of the Internet is the end-to-end principle [9], which describes a strong split between network and end systems. User-provided networks operate on the fringes of the Internet beyond the control of regular Internet stakeholders. The infrastructures they rely on are low-cost and open, capable of augmenting Internet coverage on the fly and in a viral way that follows human mobility patterns. Their elements are normally end-user devices; in contrast to today's networks, such (multimedia) devices will be part of the network: there will be no clear split between network and end systems anymore. Internet architectures and stakeholders will be affected in the same way micro-generation is affecting energy providers today.

To Route or to Relay — User-provided networks may hold elements that route traffic, but normally the devices simply relay connectivity, as explained before. While with relaying solutions connectivity has short range but the system is kept simple, with routing the connectivity range increases but the system needs to be organized based on topological information. Whether or not a solution is best suited to autonomic networks, particularly user-provided networks, needs further investigation. The option of whether to relay or route is strongly associated with the topology and the available cooperation mechanisms, both of which are expected to change dynamically and frequently in user-provided networks.

Summary and Conclusions

This article provides an overview of a new emerging type of user-centric networks, user-provided networks. It goes over their characterization, describing current examples and ideal properties that new models should incorporate in order to achieve broad and adequate coverage. Challenges associated with user-provided networks are categorized by their short- and long-term impact.

From a global perspective, it is clear that user-provided networks have tremendous potential, and are introducing a paradigm shift in Internet services and wholesale models, allowing the end user to be at the same time a consumer and provider. For instance, user-provided networks tend to evolve into more user-centric models, and such evolution is being acknowledged both by vendors and access operators, even though these Internet stakeholders do not yet know if this new phenomenon will be a threat or an interesting concept to explore. It is therefore essential to invest in research related to the challenges and advantages user-provided networks may bring to the Internet in the future.

From a global perspective, it is clear that user-provided networks have tremendous potential and are introducing a paradigm shift in Internet services and wholesale models, allowing the end-user to be at the same time a consumer and provider.

References

[1] FON Community; http://www.fon.com
[2] OpenSpark Community; http://open.sparknet.fi
[3] Whisher Community; http://www.whisher.com
[4] EU IST, "BIONETS: BIOlogically Inspired NETwork and Services," *Future and Emergent Tech. — Situated and Autonomic Commun.*, proj. ref. 027748, 2006–2010.
[5] EU IST, "HAGGLE — An Innovative Paradigm for Autonomic Opportunistic Communication," *Future and Emergent Tech. — Situated and Autonomic Commun*, proj. ref. 027918, 2006–2010.
[6] EU IST, "Social Nets: Social Networking for Pervasive Adaptation," *Future and Emergent Tech. — Situated and Autonomic Commun*, project ref. 217141, 2008–2011.
[7] N. Feamster, L. Gao, and J. Rexford, "How To Lease the Internet in Your Spare Time," *ACM SIGCOMM Comp. Commun. Rev.*, Jan. 2007.
[8] FICORA, "Application of the Communications Market Legislation to the Provision of Wireless Broadband Connections," *Ficora Memo.*, Aug. 2007.
[9] D. Clark, "The Design Philosophy of the DARPA Internet Protocols," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 18, no. 4, Aug. 1988, pp. 106–14.

Biographies

RUTE SOFIA [M] (rsofia@inescporto.pt) is a co-leader of the Internet Architectures and Networking (IAN) area, UTM, INESC Porto. From 2004 to 2007 she was affiliated with Siemens AG/Nokia-Siemens Networks GmbH & Co. KG., Munich, as a senior research scientist, working on topics such as global mobility and novel forwarding paradigms. She is a contributor to the Internet Engineering Task Force (IETF). Current research topics include global mobility management, routing, and cooperative access.

PAULO MENDES [M] (pmendes@inescporto.pt) is a co-leader of the IAN area, UTM, INESC Porto. From 2003 to 2007 he was affiliated with NTT DoCoMo Euro-labs, Munich, working on topics such as Mobile IPTV and augmented mesh routing. He is a contributor to the IETF. His major research interests are Internet architectures for mobile multihomed devices, delay-tolerant networks, cooperative networking, and mobility patterns.