

A Case for Research with and on Community Networks

Bart Braem,
Chris Blondia
University of Antwerp - iMinds

Christoph Barz,
Henning Rogge
Fraunhofer FKIE

Felix Freitag,
Leandro Navarro
Universitat Politècnica de Catalunya

Joseph Bonicioli,
Stavros Papathanasiou
AWMN, Greece

Pau Escrich,
Roger Baig Viñas
Guifi, Spain

Aaron L. Kaplan
Funkfeuer, Austria

Axel Neumann,
Ivan Vilata i Balaguer
Pangea, Spain

Blaine Tatum,
Malcolm Matson
OPLAN, United Kingdom

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.

The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

Community Networks are large scale, self-organized and decentralized networks, built and operated by citizens for citizens. In this paper, we make a case for research on and with community networks, while explaining the relation to Community-Lab. The latter is an open, distributed infrastructure for researchers to experiment with community networks. The goal of Community-Lab is to advance research and empower society by understanding and removing obstacles for these networks and services.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations—*network management*

Keywords

Community Networks, Testbed, Community-Lab

1. INTRODUCTION

Recent technological developments have pushed forward the internet and its possibilities, leading to a seemingly omnipresent internet. However, providing sustainable, cost-effective and high quality internet connection, with coverage for all citizens is still a challenging problem. Often this stems from economic causes, as internet provision in a metropolitan area is usually more economically attractive than providing access in rural areas.

“Community networking”, also known as “bottom-up networking”, is an emerging model for the Future Internet, where communities of citizens build, operate and own open IP-based networks. Hundreds of community networks operate across the globe, in rural and urban, rich and poor areas. These networks are usually run by non-profit organizations and can cooperate with local stakeholders to develop

community services, including local networking, voice connections and Internet access.

In the remainder of this paper, we first explain the main concepts behind community networks and give some example networks. Next, we describe research challenges in community networking and explain their relation with Community-Lab. Then we explain the infrastructure behind Community-Lab, followed by an overview of our open data efforts with community networks.

2. COMMUNITY NETWORKS

From a technical point of view, community networks are large-scale, distributed and decentralized systems composed of many nodes, links, content and services. They are extremely dynamic and diverse, as they are built in a decentralized manner, mixing wireless and wired links with diverse routing schemes with a diverse range of services and applications. The network is governed by an open peering agreement like [1], which avoids barriers for the participation in the network. Governance, knowledge and ownership of the network are open. Therefore these networks are not just decentralized but also self-owned and self-managed by community members, self-growing in links, capacity and services provided.

These large, decentralized, dynamic and heterogeneous structures raise challenges which can be of interest to researchers, both as a source of inspiration and as a field to apply research to. This paper wants to show interesting challenges present in community networks. It is almost impossible to compile a list of all existing community networks, as some of them are highly informal and only locally known. To illustrate the geographical and technical diversity of community networks, in what follows a list of examples is given.

In Southern Europe, Guifi.net [2] in Spain is a very large network consisting of more than 20,000 nodes and more than 24,000 links. For a detailed analysis of the Guifi.net topol-

ogy please refer to [3]. AWMN, the Athens Wireless Metropolitan Network in Greece [4], comprises more than 2,500 nodes.

Community networks expand over neighborhoods as a collaborative effort of its inhabitants, by establishing new nodes or groups of nodes linked to other nearby nodes. Nodes connect using affordable and accessible wireless IEEE 802.11 a/b/n technology, using equipment from various manufacturers, with diverse dynamic routing protocols running on different zones of each network. Operation is done in the unlicensed ISM frequency bands at 2.4GHz and 5GHz. Most networks use wireless technology although fibre links are increasingly used. In Central Europe, FunkFeuer in Austria [5] is composed of multiple smaller networks in cities like Graz and Wien, and is linked with the community network wlan slovenija in Slovenia [6]. In the South, wlan slovenija is connected to AWMN. In Western Europe, Wireless Antwerpen [7] covers a large part of Belgium and has recently connected with Wireless Leiden [8] in the Netherlands.

Networks exist almost everywhere on earth. E.g., in the USA, WasabiNet [9] is running in St. Louis, Missouri while The Personal Telco Project [10] is deployed in Portland, Oregon. In Latin America, numerous networks exist and cooperate, e.g. Bogota Mesh [11] and Monte Video Libre [12]. In Melbourne, Australia, Melbourne Wireless [13] is a quickly growing community network.

3. RESEARCH CHALLENGES

While the networks listed above are operational, numerous of challenges still lie ahead. Some of these challenges closely resemble “classical” network challenges, while others are very specific to the community networks. In this section we try to give an overview of challenges. However, given the broad range of community network technologies, this list should definitively not be considered as exhaustive.

At the physical layer, community networks often use wireless networks because of their lower costs when trying to build large-scale networks. However, the absence of cabling requires extensive wireless planning. Especially when forming a large scale mesh network in a dense urban area, channel allocation becomes very challenging to achieve correctly. Moreover, when deploying IEEE 802.11 technology over long distances, some networks have links spanning more than 20 kilometers, the MAC protocols have to be optimized or radically changed to keep functioning.

The characteristics of heterogeneity, required network neutrality, openness and size of these networks are a great challenge to routing protocols and its implementation on low-cost devices. When networks roll out their own fiber connections, as is the case in e.g. Guifi in Spain or Funkfeuer in Austria, they have to terminate the fiber. In this case, the networks can exploit the possibilities offered by optical technology normally only available in enterprise networks. When combining wireless networks with broadband fiber connections, assigning proper routing metrics becomes challenging. Networks often combine multiple link types to achieve a larger geographical penetration, and often also combine this with multiple internet uplinks. While this strategy can offer more bandwidth to end users, it requires intelligent tuning of routing protocols to properly exploit the available bandwidth given a preferred load distribution. Fiber connections can also benefit the network graph properties as long distance links can complement the relatively short dis-

tances covered by wireless links, reducing its diameter and increasing resilience to faults.

To reduce costs and democratize their construction, community networks are often built with simple, low cost off-the-shelf hardware. The nodes are usually running an open source distribution, such as Linux (Openwrt) or FreeBSD. In general this leads to a need for robustness, in the node itself and in the network. E.g., the node may need to reconfigure itself to cope with hardware problems, and should be easily upgraded remotely. For the network, routing protocols have to be robust and flexible to cope with failure, while preferably incorporating cross-layer optimizations to improve performance. This is especially important in the context of wireless links, where metrics that take into account the signal quality tend to heavily fluctuate.

A broad range of application services is used in these community networks, such as VOIP, content distribution, on-demand and live media streaming, instant messaging, remote backups and updates, file storage and file sharing. These services face enormous challenges due to the limited capacity of servers and links, the structure and diameter of the network graph, or fluctuations in the network due to load, faults, and diverse changes. Operating in this large and constantly changing environment requires the deployment of distributed service infrastructures that exploit locality, react to environmental changes and rely on cross-layer optimizations.

From a privacy point of view, community networks pose unusual challenges. Users should be able to cooperate in the network, while maintaining the privacy of their data and the data they relay. This leads to different threat models and a new notion of trust between users. Network traffic has to be organized in a fair way, respecting network neutrality. Turning these requirements into policies is challenging, verifying this automatically is even harder. This correlates to network optimizations like the placement of proxies, sometimes required to reduce the load on the community network.

Standardization is also an important challenge in the community networks context. This is not only limited to protocol design. Little standardization effort has been spent in the management of community networks, as they are very diverse by nature. However standardization is a prerequisite of further growth and sustainability of this operational model, as it helps interoperability when the networks grow and possibly merge or federate.

These challenges serve only as examples, as numerous other well-known network problems arise in community networks and could also be easily tested in the networks. E.g. bufferbloat [14], routing protocol scalability and self-configuration can strongly impact community network performance.

4. COMMUNITY-LAB

To allow more research with community networks, the Community-Lab [15] testbed is being developed and operated by the European FP7 CONFINE project (Community Networks Testbed for the Future Internet, FP7 Integrated Project 2011-2015) [16], federating existing community networks. The goal of the testbed is supporting experimentally-driven research on community-owned open local IP networks. To achieve this goal, the CONFINE testbed integrates with and extends three existing community networks: Guifi.net (Catalonia, Spain), FunkFeuer (Vi-

enna and Graz, Austria) and AWMN (Athens Wireless Metropolitan Network, Athens, Greece). Each of the three networks consists of 500 to 20,000 nodes, a larger number of links and even more end-users. In the CONFINE project, Fraunhofer (Germany), iMinds (Belgium) and UPC (Spain) contribute with research, while Pangea (Spain) operates the testbed and the OPLAN Foundation (United Kingdom) coordinates dissemination.

Simulations or small controlled laboratory experiments can produce results that are hard to translate to real scenarios [17]. Similarly to Planetlab in academic networks, Community-Lab aims at providing researchers a realistic environment for extensive experimentation, performed in real community networks. The testbed is operational, with about 50 nodes deployed in the different networks, a number which will grow over the next three years.

While Community-Lab is focused on research with community networks, as outlined previously, there are also a number of challenges to solve with research into community networks. We believe that this can ultimately remove obstacles for community networks in terms of scalability and sustainability. We expect that this research will directly impact the quality of community networks by improving current shortcomings and providing a sustainable model of community networking. As such, we believe in community networks as one viable model in a future internet.

4.1 Testbed Architecture

Community-Lab is an open, distributed infrastructure where researchers can deploy experimental services, perform experiments or access open data traces. The testbed is inspired by PlanetLab but differs in two main ways. First, the domain is community networking: a large number of modest devices, with limited and more fragile network links and all infrastructure is less managed than commercial or academic networks. Second, Community-Lab is designed to allow experiments ranging from the link layer to routing, transport and application layer, or even social experiments. These differences result in a high-level architecture equivalent to that of PlanetLab, with design differences and features as described below.

Community-Lab consists of at least one *portal* or *controller* and a set of *nodes* that are embedded in different community networks as depicted in figure 1. Each node consists of *slivers* which are grouped at a higher level in *slices*. As such, a slice is defined as *a set of resources spread over several nodes in a testbed which allows researchers to run experiments over it*. A sliver is defined as *the partition of the resources of a node assigned to a specific slice*. The purpose of the Community-Lab controller is to manage and control the testbed, i.e. manage its users, nodes, slices and slivers. It provides an aggregation point where node owners can register testbed nodes, experimenters can select nodes to create slices of resources using a web or REST API, and nodes can retrieve information to create local slivers. A *Community-Lab node* (see figure 2b) consists of two or three devices: the *community device*, the *research device* and an optional *recovery device* to force the research device to reboot in case of malfunction. These devices are connected by a wired local network, with the community device acting as a gateway.

The goal of this separation of the *Community-Lab node* into components is twofold. First, it allows the community

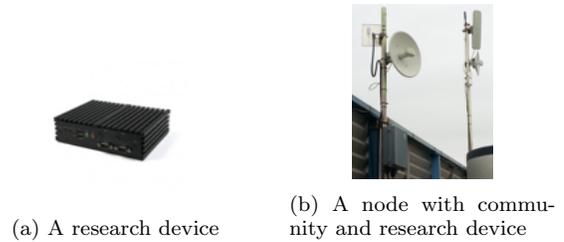


Figure 2: Community-Lab node components

device to adapt to the particular requirements of the community network, whereas the research device presents an homogeneous software stack to ease the tasks of the researcher. Second, this separation tries to preserve the stability and setup of the community network by keeping the experiments separated from the community device. As such, the effects of the experiments can be controlled and relatively powerful *research devices* (e.g. Geode LX or x86 Atom boards with at least 1 GB RAM, see figure 2a) can be connected to less powerful community devices. Still, research tools usually expect more powerful devices so caution is required [18].

Limits are also imposed by the number of Wi-Fi cards that can be plugged into the *research devices* for experimental purposes and the also heat the system generates. The latter is mainly important for outdoor installations where the board has to be put in a sealed outdoor encasement. Here the use of radio-to-router communication protocols like DLEP [19] allows for a separation of the radio interfaces from the research device without losing access to cross layer information for routing decisions. Thus, *research devices* can be deployed in a sheltered environment with adequate cooling systems while being connected to small waterproof embedded computers, most of them with a single radio device connected directly to the antenna and acting as simple bridges.

It should be noted that a Community-Lab node may either be isolated from others or within what we call a *Community-Lab cloud*. An isolated Community-Lab node only interacts with community nodes, and as a consequence it will be typically the target of application level experiments that want to analyze their performance in a community network environment. In contrast, a Community-Lab node belonging to a Community-Lab cloud can directly interact with other Community-Lab nodes and might have additional wireless interfaces to permit researchers to experiment on lower network layers.

Details of the *Community-Lab node* software architecture are depicted in figure 3. This design intends to facilitate the addition of a research device to any community device, with minimal or no CONFINE-specific changes to the community device configuration. At the same time, it provides researchers with familiar Linux environments with root access and rich connectivity options including simple NAT access (like that of most home computers), public community network addresses (to serve incoming connections), full L2 access to VLAN-tagged networks (to implement routing experiments), community network traffic sniffing (with anonymization and only where admitted by the community networks because of the privacy aspects), and raw L1 access (under certain conditions).

The *research device* runs a custom firmware (based on

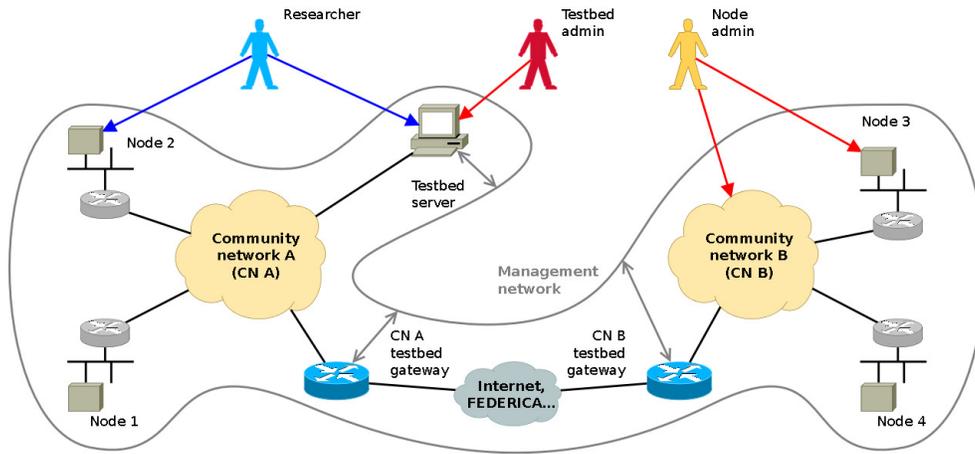


Figure 1: Architecture of the Confine testbed.

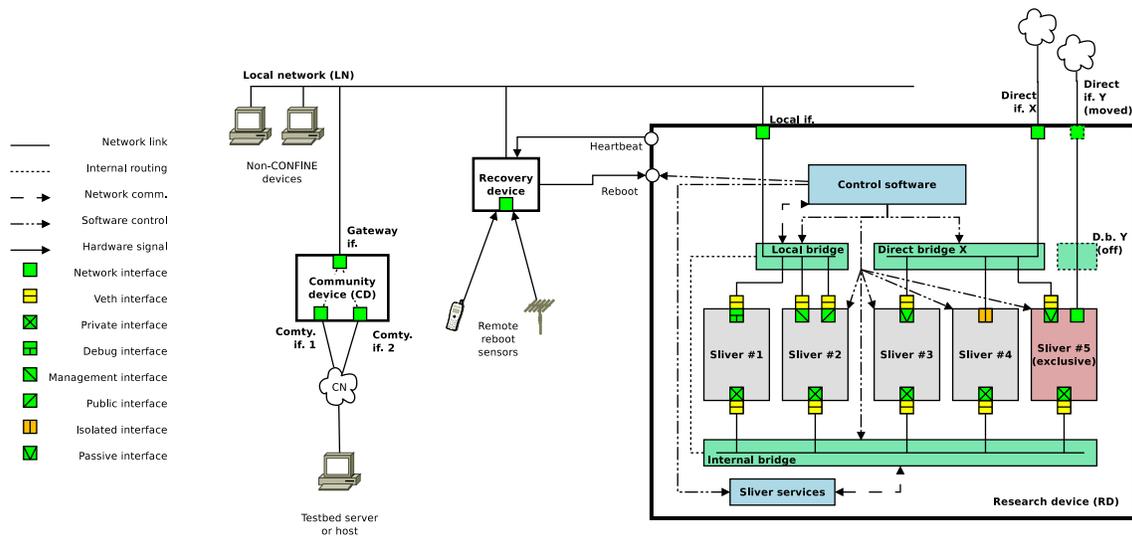


Figure 3: Architecture of a Community-Lab node.

OpenWrt) provided by CONFINE which allows simultaneously running several slivers implemented as Linux containers (LXC). Slivers have limited access to the device's resources and to one another, thus ensuring slice and network isolation. This is guaranteed by the control software via tools like ebtables, arptables, iptables, tc, Open vSwitch etc.

The *research device* implements an internal bridge. The internal address of the *research device* in this bridge is the same in all testbed nodes and it belongs to a private network address space which does not collide with community or local networks' addresses. The *research device* offers some basic sliver services on the internal address, including NAT gateway access to the community network. The *research device* also implements a local bridge which connects to the local network through a wired interface (the local interface). The bridge is used for simple network layer access to the community network through the *community device's* gateway address, and the local address of its *research device* in the bridge is fixed and used for testbed management and remote administration. For easy setup and local adminis-

tration, the local interface may also have a recovery address that is easily predictable, or the same in all testbed nodes and that belongs to a private network which does not collide with community or local network addresses (nor those of the internal bridge). A debug address which is also private and easily predictable can be used to access different research devices in the same local network for debugging purposes. The *research device* may have additional direct interfaces, each one connected to its own direct bridge. These interfaces may be connected to the community network at the link layer and used for experiments below the network level. All the aforementioned bridges are managed by the control software in order to ensure network isolation between slices (i.e. between slivers running in the same device) and to preserve the stability and privacy of the community network. The connectivity of a sliver is determined by its network interfaces, which are requested by the researcher at sliver definition time and depend on the interfaces provided by the *research device* and their features. Default routing configuration is explicitly controlled by the researcher to avoid traffic unexpectedly flowing through unwanted interfaces.

For additional details about the general Community-Lab node architecture and network connectivity please refer to [20].

4.2 Deployment status

At the moment of writing, the testbed is operational with more than 50 nodes in several locations in Europe. The project plans this number to grow over the next three years. The community network testbeds currently consist of nodes from FunkFeuer in Vienna and Graz in Austria, from AWMN in Greece in the Areas of Attica, Athens, and Tesseloniki, and from Guifi.net in Barcelona. All testbeds are connected via the *FEDERICA* (Federated E-infrastructure Dedicated to European Researchers) [21] infrastructure. In addition to the Community-Lab testbed, CONFINE maintains two additional academic testbeds for experimental purposes, connected to the Community-Lab testbed over FEDERICA. In Belgium the academic testbed is maintained by iMinds, in Germany it is maintained by Fraunhofer FKIE. The current state of the different testbeds including information on node hardware and connectivity between the experimental radio interfaces of the research nodes can be found in the CONFINE wiki [22].

5. CONFINE OPEN DATA SETS

The primary goal of the CONFINE project is the development and maintenance of Community-Lab, a testbed for experimentation with Community Networks. However, only restricting the project to the testbed would allow for limited outside participation. The project hopes to be able to share testbed access with a number of partners outside the open call, which is not feasible at least on short term. More devices would be needed, which comes with strong financial implications.

Therefore, the project also wants to generate open data sets. This will allow for outside participation, with a strong focus on community networks and to a lesser extent the testbed. Actually, the open data efforts will be focused more on the Future Internet context of CONFINE, rather than the testbed itself. As Community Networks are certainly a viable candidate as the access technology of the Future Internet, this adds an interesting perspective to the project. In what follows, we will explain the idea of open data and outline data sets CONFINE wishes to publish.

5.1 Open Data

Open data is described on Wikipedia as “*the idea that certain data should be freely available to everyone to use and republish as they wish, without restrictions from copyright, patents or other mechanisms of control*”. In the scope of CONFINE, we consider open data to be data from the project which is freely and easily available to everyone interested. To this end, the project wants to make it straightforward to obtain, to interpret and to analyze the data.

To simplify obtaining the data, the project has set up <http://opendata.confine-project.eu/> using the Comprehensive Knowledge Archive Network (CKAN) [23] software. This central catalog points to open data available from the different CONFINE partners. With CKAN, the datasets can be easily tagged and commented on. The project also wants to make sure all data can be easily downloaded in a single step. E.g., we want to avoid requiring separate downloads of weekly data sets. To this end, some of the data set

generators are modified to generate larger sets.

To simplify interpreting data, the project wants to add a description and explanation to each dataset available in the data catalog. This will require documenting the different fields of each data set, which we hope will result in more usage of the data as it avoids costly and tedious deciphering of the data formats. Finally, to simplify analysis of the data, the project will point to existing usage of the data sets. By handing examples to the potential users of the open data sets, the barrier to usage can be lowered.

5.2 CONFINE Data Sets

The selection and preparation of the open data sets is currently work in progress. However, we want to give three examples of data sets we will try to publish, to give an idea of our open data set strategy implementation.

First, the project wants to release data sets from the *node databases* of the community networks in the project, while taking care to exclude personally identifiable information. While multiple implementations of such a database exist in the different networks, they always serve as an inventory of the deployed infrastructure in the community networks. Usually it consists of the nodes, the links and the locations of deployed hardware. Sometimes, more information on e.g. the offered services and IP ranges is also available. This data set could be of interest to researchers or other community networks who want to learn more about deployment patterns and hardware typically deployed in large scale mesh networks like community networks. Also, it provides opportunities to correlate this open data with experiment results. Some networks in the CONFINE project also have historical data about growth of the network. This data set can be published to learn from growth patterns in community networks. Also, this allows for interesting visualization of both temporal and geographical data. While not immediately scientifically applicable, visualizations like this do serve as a communication medium for the size and impact of community networks.

Other networks have been recording the network topology (from a routing point of view), at a fixed frequency and during multiple years. This can provide insights in the evolving topologies, possibly correlated with the growth patterns described above. The project also wants to extend this topology monitoring to all networks, to offer a broader insight on network topologies with different routing protocols. Researchers might be specifically interested in correlating the network topology which is stored in the node database with their experiments.

These are just three examples, the project is preparing more data sets as part of the experimental research the project partners are performing using the testbed in collaboration with the different community networks involved in the project.

6. OPEN CALL FOR PARTICIPATION

The testbed is ready for experiments, growing in functionality, tools and number of research devices spread across the participating community networks. An open call for participation was published in September 2012, resulting in the reception of 36 applications. Five were selected by an independent review panel for funding with 50,000 EUR to cover the preparation and performance of experiments. Each of these 5 applications represent an external research group

with previous promising research results. The researchers will take advantage of the Community-Lab testbed to advance their research with new experiments running for one year.

In September 2013, with a more mature and larger testbed, a second open call for participation will be announced to allow the selection and support with project funding of a larger set of new experiments from external participants. You are welcome to apply!

7. CONCLUSIONS

With this paper, we wanted to make a case for research with and on community networks. Community-Lab is an excellent example of how community networks and researchers can cooperate to drive forward research, with the help of community networks. In this work we also described the architecture of our testbed, as an example of research with community networks. In addition, the open data sets published by the CONFINE project also serves as a starting point for research on community networks. While certainly different from well known commercial or private networks, we hope community networks can serve as a source of inspiration for these networks and in the long term as a source of inspiration for the Future Internet.

Acknowledgments

This work is supported by the CONFINE Integrated Project 288535. Special thanks to the participants in the CONFINE project, and to everyone involved in community networks!

8. REFERENCES

- [1] Pico peering agreement. <http://picopeer.net/>.
- [2] Guifi.net. <http://guifi.net/>.
- [3] D. Vega, L. Cerda-Alabern, L. Navarro, and R. Meseguer. Topology patterns of a community network: Guifi.net. In *WiMob 2012*, oct. 2012.
- [4] AWMN. <http://www.awmn.net/>.
- [5] Funkfeuer. <http://funkfeuer.at/>.
- [6] wlan slovenija. <http://wlan-si.net/>.
- [7] Wireless Antwerpen. <http://wirelessantwerpen.be/>.
- [8] Wireless Leiden. <http://wirelessleiden.nl/>.
- [9] Wasabinet. <http://gowasabi.net/>.
- [10] Personal telco project. <http://personaltelco.net/>.
- [11] Bogota mesh. <http://www.bogota-mesh.org/>.
- [12] Monte video libre. <http://montevideolibre.org/>.
- [13] Melbourne wireless. <http://melbourne.wireless.org.au/>.
- [14] Jim Gettys and Kathleen Nichols. Bufferbloat: Dark buffers in the internet. *ACM Queue*, 9(11), 2011.
- [15] Community lab. <http://community-lab.net/>.
- [16] CONFINE: Community networks testbed for the future internet. <http://confine-project.eu/>.
- [17] Stuart Kurkowski et al. Manet simulation studies: The incredibles. *ACM SIGMOBILE*, 9, 2005.
- [18] Nathan Samson, Glenn Daneels, Bart Braem, and Chris Blondia. Overhead analysis of embedded wireless testbeds. In *CNBuB 2012*, October 2012.
- [19] C. Barz and H. Rogge. Improved community network node design using a DLEP based radio to router interface. In *WiMob 2012*, oct. 2012.
- [20] A. Neumann, I. Vilata, X. Leon, P.E. Garcia, L. Navarro, and E. Lopez. Community-lab: Architecture of a community networking testbed for the future internet. In *WiMob 2012*, oct. 2012.
- [21] FEDERICA. <http://fp7-federica.eu/>.
- [22] CONFINE wiki. <http://wiki.confine-project.eu/>.
- [23] CKAN. <http://ckan.org/>.