



Towards a Usenet-like Discussion System for Users of Disconnected MANETs

Julien Haillet, Frédéric Guidec

► **To cite this version:**

Julien Haillet, Frédéric Guidec. Towards a Usenet-like Discussion System for Users of Disconnected MANETs. WON'08, Mar 2008, Okinawa, Japan. IEEE CS, pp.1678-1683, 2008. <hal-00341698>

HAL Id: hal-00341698

<https://hal.archives-ouvertes.fr/hal-00341698>

Submitted on 25 Nov 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards a Usenet-like Discussion System for Users of Disconnected MANETs

Julien Haillot and Frédéric Guidec
University of South Brittany, France

Abstract

Applications that rely on the publish-subscribe model (such as those dedicated to news or event distribution) require a communication model whereby information can flow towards any interested receiver rather than towards set destinations. Content-based communication is a style of communication that perfectly fits the needs of such applications. Yet supporting this kind of communication in partially or intermittently connected mobile ad hoc networks (MANETs) is still a challenge, for end-to-end content-driven routing structures can hardly be constructed and maintained in such environments.

In this paper we provide an overview of a protocol designed to support opportunistic, delay-tolerant, content-based communication in a disconnected MANET. We also describe our project of using this protocol to implement a peer-to-peer, Usenet-inspired discussion system.

1. Introduction

The Internet has grown based on a technology enabling the deployment of a world-wide, resilient communication infrastructure. This infrastructure in turn has fostered the development and general acceptance of application services that rely on the traditional client-server model. In order to benefit from these services, users (or clients) need to connect to remote servers through long-distance, dependable communication links.

Mobile ad hoc communication now proposes an entirely different vision, in which human beings can be equipped with handheld devices capable of communicating directly with one another, based on proximate transmissions, and without relying on any kind of infrastructure network.

Consider the mobile ad hoc network (MANET) shown in Figure 1. This network is composed exclusively of user-carried devices, such as laptops, personal digital assistants (PDAs), tablet PCs, and smart phones. Its topology is therefore highly unstable and liable to change unpredictably as users get close to each other, or move away for a while.




Figure 1. Illustration of a disconnected MANET composed of handheld devices

In such an environment, there is no guarantee that end-to-end connectivity ever exists between any pair of handheld devices. Indeed, when devices are very sparsely or irregularly distributed, the whole network appears as a collection of distinct “islands”. Communication between devices that belong to the same island is possible, but no temporaneous communication is possible between hosts that reside on distinct islands (such as between devices E and H in Fig. 1).

In such conditions, a method must be devised in order to allow communication between non-connected parts of the network. Delay-tolerant networking is an approach that can help with that respect [1]. In a delay-tolerant network, a message can be stored temporarily on a host, in order to be forwarded later by this host when circumstances permit. If this network is a MANET, then mobility becomes an advantage, as it makes it possible for messages to propagate network-wide, using mobile hosts as *carriers* that can move between remote fragments of the network. Finally, since contacts between user-carried devices are usually hardly predictable, they must be exploited opportunistically for exchanging messages between these devices.

Opportunistic, delay-tolerant networking can thus help support communication in a disconnected MANET. Yet applications that require end-to-end, temporaneous connectivity between clients and servers can hardly be deployed in such an environment. Even asynchronous application services such as email or discussion groups must be revised if

they are to be used in a disconnected MANET.

In this paper we address the dual problems of supporting communication in a disconnected MANET, and of implementing a realistic use-case application in such a challenged environment. We actually focus on supporting content-based communication in a disconnected MANET. Content-based communication is a style of communication that fits perfectly the needs of applications dedicated to information sharing. With this style of communication, the flow of information is interest-driven rather than destination-driven [2]. Senders simply send (or “publish”) information in the network without addressing it to any specific destination. Receivers specify the kind of information they are interested in (or “subscribe” to this kind of information), without regard to any specific source. In Section 2 we provide an overview of a protocol we designed in order to support this kind of communication in a disconnected MANET. In Section 3 we discuss our project of using this protocol to implement a peer-to-peer discussion service similar in principle to the legacy Usenet system. Like the Usenet system, this service should allow users equipped with handheld devices to post and selectively receive topical articles. Unlike the Usenet system, though, the service we consider should not rely on a distinct population of news servers for storing and providing articles on demand. Instead it should rely exclusively on the collaboration of handheld devices, so that articles disseminate by being stored, transported, and opportunistically exchanged by these devices. In Section 4 we discuss implementation details about this discussion system we plan to develop. Related work is mentioned in Section 5, and Section 6 concludes this paper.

2. A protocol for content-based networking in disconnected MANETs

In fully connected wired networks, content-based communication can be achieved by constructing an underlying communication system, whose role is to forward each piece of information from its sender to all interested hosts [2]. This system is usually organised as a logical, content-driven routing infrastructure, which itself is implemented as an overlay network that covers the whole physical point-to-point network. This approach can hardly be applied in a disconnected MANET, since in such a network the absence of end-to-end connectivity between distinct islands precludes building any network-wide overlay. Consequently, the protocol we designed does not attempt to build any structure to support routing decisions. Instead it leverages on the principle of opportunistic, delay-tolerant networking. The general idea is to let neighbouring hosts opportunistically exchange pieces of information they are missing, based on their respective advertised profiles. A host that behaves as a receiver of a particular kind of information is expected

```
<descriptor
  id="254d3g64z36cd"
  date="Fri Oct 12 09:52:11 CEST 2007"
  from="Fred"
  newsgroup="comp.networking"
  keywords="mobile,ad hoc,delay-tolerant,\
            opportunistic,gossip-based" />
<profile
  newsgroup="comp.(networking|communication)"
  keywords="sensor|vehicular|opportunistic" />
```

Figure 2. Examples of a document’s descriptor and of a host’s interest profile

to serve as a mobile carrier for this information, so while roaming the network it can help transport this information towards remote parts of this network.

Our protocol actually supports the content-driven dissemination of data units we refer to as “documents”. A document is a structured piece of information a mobile host can publish in the network, so it can disseminate and ultimately be received by any host that shows interest for this particular kind of information. A document is composed of two parts: its descriptor, and its content (or payload). The descriptor can be perceived as a collection of attributes, which can provide any kind of information about the corresponding document, such as its identifier, its origin, its topic, a list of keywords, the type of its content, etc. It can typically be formatted in XML, as shown in Figure 2.

As a general rule, we assume that the size of a document is far greater than that of a descriptor, which is itself significantly greater than that of a document identifier. The protocol leverages systematically on this contrast in order to minimise the amounts of data exchanged by neighbouring hosts.

Each mobile host is characterised by its “interest profile”, that is, a description of the kinds of documents it is interested in, and thus—implicitly—the kinds of documents for which it is willing to serve as a mobile carrier. For practical reasons, we define this profile as a predicate applicable to document descriptors. By applying this predicate to a descriptor, a host can decide if it must store the corresponding document in its cache or not. An example of a host’s interest profile is shown in Figure 2. This profile can match any document whose descriptor refers to either the *comp.networking* newsgroup, or the *comp.communication* newsgroup, and that contains any of the specified keywords. (This profile therefore notably matches the descriptor shown in Fig. 2.)

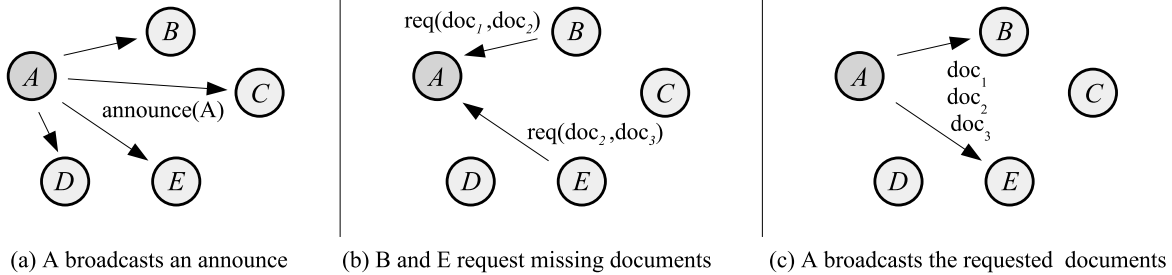


Figure 3. Illustration of a transmission cycle around host A

The content-driven dissemination of documents in the network relies on a simple model, whereby transient contacts between mobile hosts are exploited opportunistically for exchanging documents between these hosts, according to their respective interest profiles. Each host periodically broadcasts an announce that combines a description of its own interest profile, and a catalog that contains the descriptors of locally cached documents that can be of interest to its neighbours (see Fig. 3.a). By broadcasting its own interest profile, a host lets its current neighbours know what kind of documents it is interested in. Conversely, by receiving similar information from its neighbours, each host can maintain an accurate vision of the kinds of documents they are interested in. Thus whenever a host needs to broadcast an announce, the catalog it contains can be adjusted in order to fit specifically the interest profiles of its neighbours. With this approach, the cost of broadcasting an announce periodically is kept at a minimum, as the size of the catalog a host broadcasts periodically is continuously adjusted in order to fit exactly the interest profiles of its current neighbours.

When a host discovers that one of its neighbours can provide a document it is interested in (that is, a document that matches its own interest profile and that is not already available in its own cache), it can request a copy of this document from this neighbour (Fig. 3.b). Upon receiving a request for a document that is available in its local cache, a host sends this document on the wireless channel (Fig. 3.c). Note that when a document is requested by several neighbours of the announcer, it is sent in broadcast mode, rather than being sent in unicast mode to each requester successively. This approach reduces the amount of duplicate transmissions on the wireless medium and, as a side-effect, it helps preserve the power-budget of mobile hosts.

This protocol has been fully implemented in Java, and embedded in DoDWAN¹, a middleware platform we designed in order to evaluate and demonstrate the potential of opportunistic networking in challenged networking environments such as disconnected MANETs, vehicular net-

works, or sensor networks. In its current implementation, DoDWAN relies on the transmission facilities of a standard IP protocol stack. Document descriptors are formatted in XML (as shown in Fig. 2), and all transmissions are performed by embedding control messages (i.e., announces and requests), descriptors, and documents in UDP datagrams. When a host needs to publish a large document (that cannot fit in a single UDP datagram), this document is segmented, and the resulting segments are encapsulated in a number of “fragment documents”. These fragments are then published separately, so they can follow different routes in the MANET before being collected and reassembled only by hosts that are interested by the information they contain.

More details about the protocol can be found in [3], that also presents some of the results we obtained when evaluating its performance with a simulator, as well as in real conditions.

3. Towards a Usenet-like discussion service for disconnected MANETs

In this section we discuss the implementation of a peer-to-peer discussion system that should make use of our content-based communication protocol, and provide users with a service that is inspired from the legacy Usenet system. We first give an overview of the Usenet system, highlighting its major characteristics, before considering the implementation of a system capable of providing similar services in an infrastructure-less wireless network.

Overview of the legacy Usenet architecture. Usenet is an effective, non-centralised discussion system that allows users to post topical articles to so-called newsgroups. It relies on a large, constantly changing set of interconnected servers, which are associated pair-wise in order to form a loose overlay over the physical Internet. Each server hosts a number of newsgroups, meaning it can store articles pertaining to these newsgroups, and provide these articles on demand to peer servers and to user-level client programs.

¹DoDWAN stands for *Document Dissemination in Wireless Ad hoc Networks*.

End-users must use a client program (a.k.a. *newsreader*) in order to access articles stored on a server. With such a program, a user can subscribe to some of the newsgroups hosted by the server, post new articles to these newsgroups, and fetch recent articles from the server. Peer servers synchronise regularly in order to exchange recent articles.

Interaction between peer servers, and between servers and clients, is supported by the Network News Transfer Protocol (NNTP). This protocol itself relies on TCP, which provides reliable, session-based data transport in IP networks.

A Usenet-like discussion system for disconnected MANETs. The Usenet system has been designed in order to rely on the resilient communication links provided by the Internet. Moreover, it leverages on a clear decoupling between clients and servers: users are assumed to run client programs, that must access remote servers in order to post or fetch new articles.

Our objective is to design a discussion system we refer to as OppDNA (*Opportunistic Dissemination of News Articles*) in the remainder of this paper. The OppDNA system should closely mimic the Usenet system, while running exclusively on a collection of handheld mobile devices (such as laptops, PDAs, or smart-phones) interacting only based on proximate, ad hoc transmissions. With this system, a user should be able to post and receive topical articles, in a manner similar to that allowed by the Usenet system. Moreover, we would like to take benefit from the legacy software implemented in standard newsreaders, and from the fact that users are already familiar with such programs. Users should thus be allowed to rely on standard newsreaders (which are, basically, NNTP clients) in order to access newsgroups, even though the underlying system that maintains these newsgroups is liable to behave differently from that of the Usenet architecture.

In a disconnected MANET, no device can be considered as being stable enough, accessible enough, and resourceful enough to play the role of a dedicated news server. Consequently, the newsreader running on a user's handheld device cannot be associated with a stable, remote server. The solution we propose consists in having each device run a local pseudo-server, with which the user's newsreader can interact through NNTP sessions whenever the user wishes to either post or fetch new articles.

In the standard Usenet architecture, servers are administered explicitly by dedicated personnel. A server's administrator can create local newsgroups when needed. He is also responsible for deciding which newsgroups the server must host, with which peer servers (or "feeds") it should synchronise, and how often this synchronisation should be performed. In the OppDNA system, the owner of a handheld device must somehow be allowed to "administer" the OppDNA program running on this device. Since a hand-

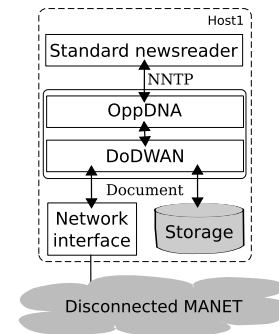


Figure 4. Architecture of an OppDNA instance

held device is usually severely resource-constrained, it is not reasonable to consider that a user should allow his device to host *all* newsgroups. Instead the list of locally hosted newsgroups should be defined cautiously by the user, considering how much storage resource can be devoted to this activity on the device, and selecting primarily newsgroups he is liable to subscribe to through his newsreader. Moreover, although Usenet administrators can rely on newsgroup lists (available world-wide) to select the newsgroups they wish to serve on their servers, this approach is not applicable in the OppDNA system. No device in a disconnected MANET can serve as a central repository for maintaining the list of all newsgroups. Instead a *discovery* mechanism is required for OppDNA users to learn about newsgroups.

4. Overview of the OppDNA system

General architecture. Figure 4 shows the general architecture of an OppDNA program. This program must run in the background on each mobile host that is expected to participate in the discussion system. It is organised in two main layers. The upper layer is in charge of managing newsgroups, and interact with the user's newsreader through a standard NNTP session. The lower layer manages the interactions of the OppDNA program with instances of the same program (or *peers*) that run on other mobile devices. In practise, whenever two mobile devices—and therefore two OppDNA peers—get in radio contact, they should use this opportunity to exchange articles pertaining to the newsgroups each peer is interested in. The content-based information dissemination protocol implemented in the DoD-WAN middleware (as described in Section 2) is perfectly suited to support such exchanges.

Thus, whenever a user posts an article to a particular newsgroup, this article is simply passed from the newsreader to the associated OppDNA program. The article is first converted into a format the DoD-WAN middleware can handle, and the resulting document is then deposited in the cache

maintained by DoDWAN. Thence the article will disseminate among mobile devices by being passed opportunistically to any device whose advertised profile shows that it is interested in the newsgroup the article has been posted to.

Symmetrically, whenever the DoDWAN middleware obtains from a peer host an article that matches its “interest profile”—that is, an article that pertains to a newsgroup it has been ordered to host locally—this article is simply deposited in the local cache, and the upper layer of the OppDNA program is notified of this reception, so the article can be made available for the user’s newsreader.

“Administration” of the OppDNA program. The interest profile that determines what kinds of documents the DoDWAN middleware should strive to collect from neighbouring devices must be set according to the user’s preferences.

In the Usenet system, the newsreader can register the user’s subscriptions, and fetch articles accordingly from the server. Yet it cannot instruct the server to subscribe to any newsgroup, as the server’s subscriptions are controlled directly by its administrator. In the OppDNA system a method must be devised to allow a user to configure the local OppDNA program. Several alternatives can be considered, such as giving this program a user-friendly textual or graphical front-end. The approach we favour consists in using a dedicated “administration” newsgroup, and posting specially formatted articles to this newsgroup in order to pass instructions to the OppDNA program. Thus, assume a user wishes to create a new newsgroup called *opportunistic.networking*. This user simply needs to post to the *oppdna.admin* newsgroup an article with a *Subject* line of the form:

Subject: ANNOUNCE opportunistic.networking

The body of this article should contain a free-form description of the topics covered in this newsgroup.

By default the OppDNA program should be configured so as to receive all articles pertaining to the *oppdna.admin* newsgroup. A user that wishes to get informed about existing newsgroups therefore simply needs to read articles posted to *oppdna.admin*. Thus, assume that a user learns about the *opportunistic.networking* newsgroup, and decides that his handheld device should receive the corresponding articles (and therefore participate in their distribution in the network). This user can instruct the local OppDNA program to host this newsgroup, by “posting” to *oppdna.admin* an article with a *Subject* line of the form:

Subject: SUBSCRIBE opportunistic.networking

When such an article is received by the upper layer of the OppDNA program, it is not passed to the DoDWAN middleware. Instead the interest profile that rules the behaviour of the DoDWAN middleware is modified so that, from now on, it will strive to collect articles that pertain to the *opportunistic.networking* newsgroup.

The user can of course instruct the OppDNA program to stop hosting a particular newsgroup by posting an `UNSUBSCRIBE` command to the *oppdna.admin* newsgroup. The interest profile advertised by DoDWAN is then modified accordingly, and its cache is purged from all articles pertaining to this newsgroup.

5. Related work

A number of protocols have been designed in the last few years in order to support either *destination-driven* routing, plain broadcasting, or group-based multicasting in disconnected MANETs [4, 5, 6, 7, 8]. In contrast the support of *content-based* communication in such networks has not justified much research so far. A notable exception with that respect is [9], which describes an approach whereby a content-driven multi-hop routing structure (limited to a given horizon) is built around each host. A utility-based function is used to select the best carriers and/or forwarders for each kind of message, and mobile carriers are meant to transport messages between non-connected parts of the network. Our protocol relies on a slightly different approach. Instead of attempting to construct and maintain a routing structure, it relies on the principle of letting neighbouring hosts exchange documents they are missing, based on their respective advertised profiles. From this viewpoint it can be considered as being a member of the large family of gossip-based protocols [10, 11].

Complementarily to papers addressing the development of communication protocols for disconnected MANETs, a few recent papers also discuss the design of applications capable of exploiting these protocols [12, 13, 14]. Indeed, the very characteristics of opportunistic, delay-tolerant communication imply that these applications must be able to tolerate asynchronous—and possibly fail-prone—transmissions. Electronic mail (or E-mail) and news (or event) distribution are obvious candidates, as they can tolerate conditions in which temporaneous end-to-end connectivity between users is not guaranteed. Thus, several papers propose to rely on destination-driven, delay-tolerant routing protocols for implementing a messaging service capable of operating in a disconnected MANETs. [13] and [15] propose to link mobile hosts roaming a disconnected MANET with remote SMTP servers located in the wired Internet, using DTN-enabled proxies or gateways as a means to bridge the gap between the disconnected and connected parts of the network.

Our project of implementing a discussion system for users of disconnected MANETs somehow differs from the abovementioned works. First, instead of relying on a destination-driven routing protocol for transporting messages towards a set destination, we rely on a content-driven dissemination protocol for transporting topical arti-

cles towards all interested receivers. Second, the system we describe relies exclusively on the collaboration of mobile devices without ever resorting to remote Internet-based servers (although a hybrid system allowing opportunistic synchronisations with NNTP servers is an option we consider for future work). Interestingly, [14] proposes to use a quite similar approach in order to allow mobile hosts to exchange typed content (such as small audio or video files), using topical channels as a means to selectively publish—and subscribe for—different kinds of content.

6. Conclusion

In this paper we have provided an overview of a protocol we designed in order to support content-based delay-tolerant information dissemination in disconnected MANETs. This protocol opportunistically exploits transient contacts between mobile hosts for exchanging documents between these hosts, according to their respective interest profiles. We also presented our project of using this protocol to develop OppDNA (*Opportunistic Dissemination of News Articles*), a peer-to-peer discussion system that should be similar in principle to the legacy Usenet system, without relying on any server. OppDNA is meant to be used by users equipped with handheld devices. It should rely exclusively on the cooperation of these devices to disseminate newsgroup articles in the network. In the future we plan to give users of the OppDNA system a better control of the behaviour of their OppDNA-enabled handheld device, notably by setting priority levels on the newsgroups they host on this device. We would also like to give users the possibility to control the spatial and temporal scope of dissemination for the articles they post. Finally, we consider linking the OppDNA with the Usenet system, allowing OppDNA peers to post articles to—and fetch articles from—standard NNTP servers whenever they get a chance of connecting to the Internet.

Acknowledgements. This work is supported by the French *Agence Nationale de la Recherche* under contract ANR-05-SSIA-0002-01.

References

- [1] K. Fall, “A Delay-Tolerant Network Architecture for Challenged Internets,” in *Proceedings of ACM SIGCOMM’03*, Aug. 2003.
- [2] A. Carzaniga and A. L. Wolf, “Content-based Networking: a New Communication Infrastructure,” in *NSF Workshop on an Infrastructure for Mobile and Wireless Systems*, no. 2538 in LNCS, pp. 59–68, Springer-Verlag, Oct. 2001.
- [3] F. Guidic and Y. Mahéo, “Opportunistic Content-Based Dissemination in Disconnected Mobile Ad Hoc Networks,” in *Int. Conf. on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM 2007)*, pp. 49–54, IEEE Press, Nov. 2007.
- [4] Z. Zhang, “Routing in Intermittently Connected Mobile Ad Hoc Networks and Delay Tolerant Networks: Overview and Challenges,” *IEEE Communications Surveys and Tutorials*, vol. 8, pp. 24–37, Jan. 2006.
- [5] L. Pelusi, A. Passarella, and M. Conti, “Opportunistic Networking: Data Forwarding in Disconnected Mobile Ad Hoc Networks,” *IEEE Communications Magazine*, vol. 44, pp. 134–141, Nov. 2006.
- [6] A. Vahdat and D. Becker, “Epidemic Routing for Partially Connected Ad Hoc Networks,” Tech. Rep. CS-2000-06, Duke University, Apr. 2000.
- [7] J. Ott, D. Kutscher, and C. Dwertmann, “Integrating DTN and MANET Routing,” in *Proceedings of the 2006 SIGCOMM Workshop on Challenged Networks (CHANTS’06)*, pp. 221–228, ACM Press, Sept. 2006.
- [8] G. Karlsson, V. Lenders, and M. May, “Delay-Tolerant Broadcasting,” in *Proceedings of the 2006 SIGCOMM Workshop on Challenged Networks (CHANTS’06)*, pp. 197–204, ACM, Sept. 2006.
- [9] P. Costa, M. Musolesi, C. Mascolo, and G. P. Picco, “Adaptive Content-Based Routing for Delay-Tolerant Mobile Ad Hoc Networks,” Tech. Rep. RN-06-08, UCL Department of Computer Science, Aug. 2006.
- [10] A. Datta, S. Quarteroni, and K. Aberer, “Autonomous Gossiping: a Self-Organizing Epidemic Algorithm for Selective Information Dissemination in Mobile Ad-Hoc Networks,” in *IC-SNW’04 (International Conference on Semantics of a Networked World)*, no. 3226 in LNCS, pp. 126–143, June 2004.
- [11] S. Baehni, C. S. Chhabra, and R. Guerraoui, “Selective Information Dissemination in a Mobile Environment,” in *Proceedings of the 6th Middleware Conference*, Dec. 2005. Extended version.
- [12] J. Scott, P. Hui, J. Crowcroft, and C. Diot, “Haggle: a Networking Architecture Designed Around Mobile Users,” in *Proceedings of the 2006 IFIP Conference on Wireless on Demand Network Systems and Services (IFIP WONS 2006)*, Jan. 2006.
- [13] T. Hyyryläinen, T. Kärkkäinen, C. Luo, V. Jaspertas, J. Karvo, and J. Ott, “Opportunistic Email Distribution and Access in Challenged Heterogeneous Environments,” in *Proceedings of the 2007 SIGCOMM Workshop on Challenged Networks (CHANTS’07)*, pp. 97–100, ACM Press, Aug. 2007.
- [14] V. Lenders, G. Karlsson, and M. May, “Wireless Ad Hoc Podcasting,” in *4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON ’07)*, pp. 273–283, IEEE CNF, June 2007.
- [15] E. Yoneki and J. Crowcroft, “Haggle Deliverable 6.1: User-Centred Design of Haggle Applications,” Feb. 2007. Haggle project, http://www.haggleproject.org/images/b/b3/Haggle_D6.1.pdf.