

# The Limits of Theory: Pragmatic Challenges in Mobile Ad Hoc Systems

## (Invited Paper)

Franck Legendre, Martin May  
ETH Zurich  
8092 Zurich, Switzerland  
legendre,maym@tik.ee.ethz.ch

Vincent Lenders  
Armasuisse  
3000 Berne, Switzerland  
lenders@tik.ee.ethz.ch

Gunnar Karlsson  
KTH, Royal Inst. of Technology  
10044 Stockholm, Sweden  
gk@ee.kth.se

**Abstract**—The development of mobile ad hoc systems have considerably emphasized the need for a better understanding of the factors that influence the systems’ performance, i.e., mobility patterns, radio propagation, traffic characteristics and their interrelations. Depending on the context where the system is being used—whether at work, at home, or in some means of transportation—there will be various applications that provide benefit of using the system. These different applications in turn generate traffic with very heterogeneous characteristics. In addition, the radio propagation depends on the surrounding environment and the density of communicating nodes in the system and in other systems that compete for radio resources; hence the available data rates in the system are hard to ascertain. The uncertainty regarding both the traffic and the available data rates make it hard to evaluate the performance of the system. This uncertainty is due to the lack of deployed systems. We hence believe that research in mobile ad hoc networks has to couple theory and practice in a continuous feedback loop with experimentations running full-fledged applications. We show how this methodology has proved to be beneficial in order to overcome challenges and to determine main issues for future research in the context of our PodNet project. It has also allowed us to revisit assumptions and scepticism surrounding the feasibility of mobile ad hoc networks. PodNet is an architecture, instantiated in a prototype, that is dedicated to cooperative content distribution. It enables exchange of podcasts (for any kind of multimedia content) amongst mobile devices in a peer-to-peer fashion using IEEE 802.11 in ad hoc mode. PodNet has clear advantages over traditional content distribution approaches (e.g., newspapers, 3G) in terms of public availability, practicability, capacity, resilience to failures, and jamming.

### I. INTRODUCTION

At the early stage of wired networking research, teletraffic models and queuing theory were invaluable for performance evaluation, network optimization, and systems design. With the advent of infrastructure-based wireless communication systems (e.g., PCS, GSM), mobility models became used to include user mobility and spatial distributions of calls [14]. Since then, research has been ongoing in this area for new technologies such as GPRS and EDGE [20]. With self-organizing networks (e.g., ad hoc, delay-tolerant, sensor, and mesh networks), the difficulty in determining the inputs that affect the performance of the system—and the corresponding modelling of those inputs—has become greater. First, the

wide-area cellular systems rely on radios with ranges up to several kilometers, while the mobile ad hoc systems are based on low power radios of some tens of meters in range, occasionally reaching up to a few hundred meters in range. Since the radio propagation is influenced by the environment (e.g., indoors vs. outdoors, urban vs. rural areas), this means that it has to be modelled on a very fine spatial scale. As a result, it is extremely difficult to determine the physical data rates for a system in all of its operational conditions. Second, the system’s performance does not depend on the mobility of one node with respect to an infrastructure, as in the cellular case, but rather on many nodes’ mobility with respect to one another. Third, the dependencies of the inputs have to be considered: For instance, a high density of nodes might lead to low physical data rates due to channel contention but to higher connectivity at the network layer. Fourth, the networking support that is provided by the combined effects of the lower three protocol layers must then match the traffic from the applications that the users of the mobile ad hoc system want to run. Also the applications and the ensuing traffic characteristics depend on the mobility and the context of the nodes (no video when walking, for instance)

Assumptions on a system’s input are most often the result of simplifications or even speculations, which in the end might be proven wrong. For example, the de facto assumptions that mobility is a random walk constrained to a circle or a square, and that traffic generation at the packet level is an i.i.d. Poisson process, have eased the analytical tractability of modelling, but has shown to be wrong for many cases. Although still debated, human mobility has shown to be characterized by a power-law behavior for meeting and inter-meeting times [5]. Web and P2P traffic has been shown to exhibit a self-similar property [7] and a Zipf law for content popularity [28], [10]. In the work by Grossglauser and Tse [8] on the beneficial aspect of mobility on the capacity of ad hoc networks, the authors assumed i.i.d. random placement of nodes; but, as it has been observed in a study by Tuduca and Gross [23], this is typically not the case in real networks. In light of new insights on mobility, traffic, and their interrelation, we hence believe that mobile networking research must enter a new stage, with re-evaluation

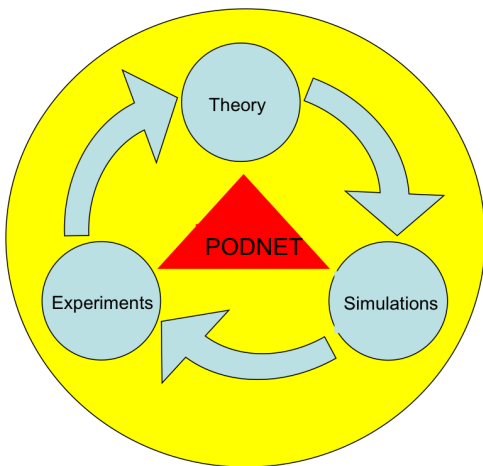


Fig. 1. The golden circle for research.

of all major assumptions that have been leveraged in the last decades.

With the mass production and consumer penetration of wireless devices (e.g., PDA, smart phones, sensors), it is now conceivable to deploy large-scale real-world testbeds. In fact, we need to couple theory and practice more closely in a continuous feedback loop, as shown in Fig. 1, to avoid theoretical research that unnecessarily reach dead-ends. There have been many successful experiments for discovering or better understanding characteristics of the performance input parameters, such as Haggie [5], ZebraNet [13], RoofNet [2], DieselNet [27], and CrowdadD [1]. Yet, many of these experimental works have been confined to the study of one input parameter at a time and not their interrelation.

In this paper, we present *PodNet*, a research project aimed at cooperative content dissemination (CCD). It integrates theoretical research along with experiments. PodNet expends the traditional podcasting concept for public (i.e., open and unrestricted) peer-to-peer delivery of contents amongst mobile nodes. PodNet allows exchanging podcasts or any kind of multimedia content from mobile devices to mobile devices in a peer-to-peer fashion using 802.11 in ad hoc mode. It is a potential candidate to complement infrastructure-based content dissemination techniques relying on EDGE or 3G for example. We soon discovered that a prototype was needed to validate theoretical results by performing mobility measurements and characterizing traffic patterns. To our knowledge, PodNet is one of the first approaches to study a self-organized system as a whole from the application with user-interactions to the physical layer, which will ease the study of interrelation between traffic, mobility, and radio channel conditions. The PodNet project has been in progress since 2005; although a lot has been achieved, we consider it still to be in its infancy given the number of issues that remain to be tackled.

The remainder of this paper is organized as follows. The PodNet architecture is overviewed in Section II. Section III presents different real-world scenarios where ad hoc content dissemination brings many advantages over current

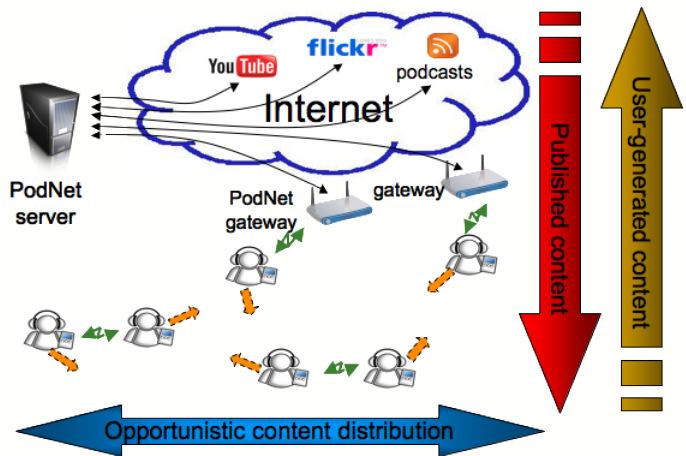


Fig. 2. The PodNet architecture.

infrastructure-based dissemination techniques. Section IV lists a number of challenges that have been overcome and that are still to be faced. Finally, Section V concludes our work.

## II. PODNET

### A. What is PodNet?

Wireless ad hoc podcasting [18], [19], as we propose in PodNet, expends the traditional podcasting concept for public (i.e., open and unrestricted) peer-to-peer delivery of contents amongst mobile nodes. Fig. 2 describes our architecture. Contents are provided in a wireless broadcasting area either through fixed *PodNet gateways* or by *mobile nodes* themselves. In the first case, a *PodNet server* fetches contents from servers across the Internet and forwards them to mobile nodes within the range of PodNet gateways; this is done in the regular podcasting mode. In the second case, a mobile node that has contents to share might provide data to another mobile node when they pass within radio range of one another. This is one of the new content distribution mode that we add to the existing one. The other content distribution mode we offer is the possibility for mobile devices to upload user-generated content to the PodNet server when passing within range of a PodNet gateway. Our ad hoc podcasting mode brings the following advantages. First, it provides nodes with contents when they are not connected to the Internet; second, it provides a new ad hoc broadcasting domain when also the sources of the data are mobile nodes (could be pictures or video/voice recordings from a mobile phone for instance).

The sharing of contents is based on a solicitation protocol by which a node asks a peer node for content (a peer node can either be a mobile device or a PodNet gateway). Hence, there is no flooding of contents in the broadcasting area. Contents are organized into *feed channels*, and nodes solicit *episodes* for one or more feed channels. The concept of feed channels allows a higher hit rate of the queries than if they were made for individual episodes of contents. The episodes of a particular feed channel will however reach a node in arbitrary order and the downloads might get interrupted when

Channel	Episode	Enclosure	Chunk 256 KB	Piece MTU size	
PIMRC	Keynote	Challenges of an Integrated Operator (MP3 1)	Chunk 1	Pieces 1-x	
		Between Commerce and Research (Video 1)	Chunk 2-z Chunk 1	Pieces x-y Pieces 1-x	
		Presentations	Video 2	Chunk 1	Pieces 1-x
	Panels	MP3 2		Chunk 2	Pieces x-y
				Chunk 1	Pieces 1-x
		The Internet of the Future (Pic 1) idem (Video 3)		Chunk 2	Pieces x-y
				Chunk 1	Pieces 1-x
	Social events	Dinner (PDF 1) Excursion (Video 4)		Chunk 1	Pieces 1-x
				Chunk 1	Pieces 1-x
				Chunk 2	Pieces x-y
Application Layer			Transport Layer		

TABLE I  
DATA STRUCTURE FOR EXAMPLE CONTENTS.

the nodes move away. We therefore introduce *chunks* (inspired by BitTorrent [12]), atomic data units that allow an incomplete episode download to be resumed with the same node or any other node that also has the episode. Table I gives examples of channels with their respective episodes and enclosures (i.e., sub-episodes). Enclosures are divided in chunks which are themselves divided into pieces to be transported in wireless frames.

Our wireless ad hoc podcasting system is simple in its protocols, and lightweight in the computational requirements. We allow only pair-wise node associations to increase the chance that downloads eventually succeed when the contact durations are short. As a result, we do not need an explicit multi-hop routing protocol. Dissemination of content is performed implicitly at the application layer, through content distribution by gateways and through node mobility. Note that the entire concept is centered around the availability of data items and the identity of the nodes that provide them are not important! This is in contrast with other instances of delay-tolerant networks where messages go from and to users: data can be replicated for increased availability, while users cannot.

### B. The PodNet Architecture

The PodNet system is composed of the following entities as depicted in Fig. 2:

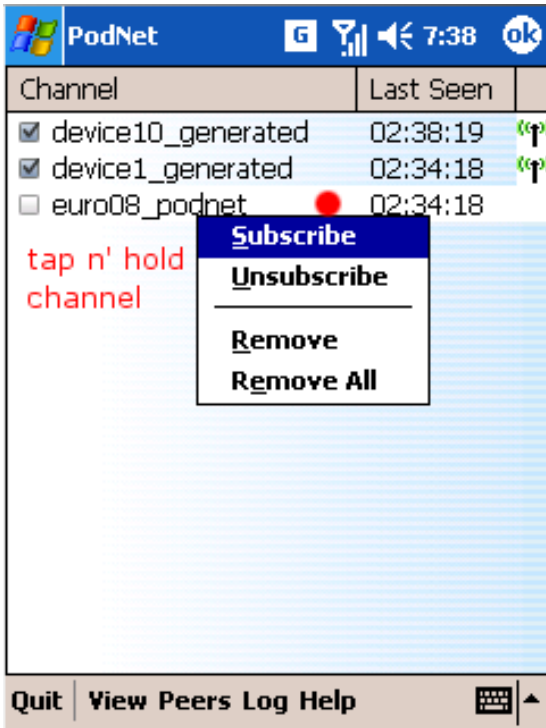
- A PodNet server
- PodNet gateways
- Mobile devices running PodNet

1) *PodNet server*: The PodNet server runs a web server and a content aggregator. The former is used as a front-end to display available contents (feed channels). The latter fetches contents from the Internet and also user-generated content sent by mobile devices through the PodNet gateways. It is

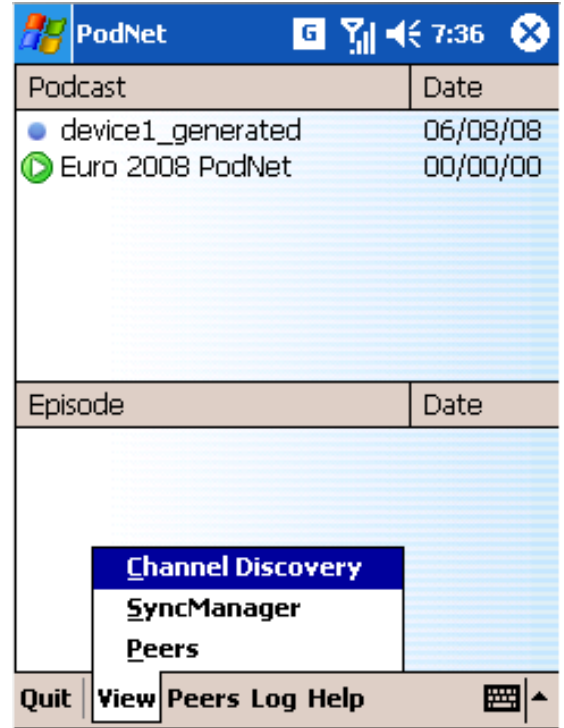
also responsible for transcoding content to be readable on any mobile device.

2) *PodNet gateways*: PodNet gateways provide a bridge between the PodNet server and mobile devices running PodNet. It allows interconnecting the wired Internet domain and the ad hoc podcasting domain seamlessly. These gateways can be installed in strategic locations (i.e., town centers, attractive places) or setup on-demand to support temporary events such as sport games, expositions or conferences. Gateways act similarly to the PodNet mobile devices, but may have mains power.

3) *Mobile devices*: Mobile devices running PodNet allow the distribution of content, either fetched from the Internet (through gateways) or from other devices. As described previously, these devices can also generate their own content (i.e., user-generated content) from their multimedia recording capabilities. These user-generated contents can then be distributed opportunistically between devices and to gateways that forward them to the PodNet server. Devices participating in the PodNet system must have a IEEE 802.11b/g interface configured in ad hoc mode with the same SSID (e.g., “PodNet”) and zero-configured at the IP layer. The PodNet software works automatically and provides a GUI that requires only a few interactions from the user: He or she has only to subscribe to feed channels of interest. The following explains in more detail how content is exchanged opportunistically between peers (i.e., either between two devices or one device and a gateway): the PodNet application sends UDP broadcast messages to discover peers. Among those in range, one peer is randomly selected and a TCP connection is created for data exchange. To speed up the synchronization process we employ a Bloom filter that summarizes the list of available contents [3], [4]. Figures 3 shows screenshots of the application running on Windows Mobile Edition. Note that our software is also ported to Symbian, Linux, and Mac OS X (iPhone to be released soon).



(a) View showing all discovered channels available locally and from peer devices. The user subscribes to the “euro\_2008\_podnet” channel.



(b) Main view showing all subscribed channels (or podcasts) and their episodes. No episodes are displayed here.

Fig. 3. Screenshots of the PodNet software.

### III. APPLICATIONS

There are many applications that benefit from PodNet or CCD in general. In the following, we classify these applications and discuss the advantages of the opportunistic communication in given application scenarios.

The intrinsic advantage of PodNet is the ability to disseminate content even without a permanent connectivity to all neighboring nodes. Besides extending network connectivity, many new applications may benefit from the cooperative nature of our content dissemination scheme.

With regard to the freshness of the available information, CCD is able to provide more up-to-date information to a user group than other existing solutions. For example in buses or trains where a connection to the Internet is not provided permanently, PodNet continues to distribute content as long as fresh content arrives in the area due to gateways and user arrivals. Thinking about the distribution of free dailies like 20 Minutes and Metro, devices that walk by gateways will download a snapshot of the latest news in electronic form rather than printed on paper. The devices will then be able to re-distribute the contents when disconnected from the Internet. The fact that different users will download different contents will lead to community effects, i.e., each community will collect and share its information but will be able at the same time to attract new members and will hence grow the community.

The same mechanism as the one of infostations [21], [26], [25] could be used in conferences or exhibitions. At some few locations fresh content is provided; the dissemination of the content is then handled by the users. The users are able to rate the content, add meta information and even contents, such as pictures, videos, and voice clips.

CCD increases the network capacity compared to non-cooperative content distribution schemes. For example in stadiums with thousands of spectators, the network capacity is bounded by the cell size to which the mobile nodes are associated. With CCD, the transmission range can be lower leading to a better spatial reuse of the wireless spectrum that results in a higher network capacity.

A final advantage lies in the concept itself. Since our system does not require any infrastructure support it is resilient to node failures. Consider a solar-powered gateway disseminating news reports that vanishes because night is falling and the batteries are exhausted. In this case, our system will still be able to rely on opportunistic communications to distribute content despite the gateway being down. Our system is also resilient to malicious activities such as jamming since content dissemination is spatially distributed. Potentially, the opportunistic content dissemination could be attractive for military applications or for rescue operations where existing infrastructures are powered down, damaged, or not available at all.

#### IV. PRACTICAL CHALLENGES

In the following, we detail practical challenges of CCD systems. These challenges are the outcome of problems we encountered during our research and implementation efforts with the PodNet system.

##### A. Data transfers

With IEEE 802.11, ad hoc communication between mobile devices is handled differently than communication between mobile nodes and an access point. The IEEE 802.11 standard foresees the *ad hoc* and the *infrastructure* association mode. Unfortunately, it is not possible to operate both modes with one chipset simultaneously which means that a node has to switch between the two modes if it wants to communicate with both worlds. The switching is typically time consuming and one does not want to do this frequently. Hence, a mobile device might miss connection opportunities because it resides in the wrong mode. [6] propose to virtualize the Wifi interface by switching between both modes but it is rather a patch than a clean slate solution.

Another issue is related to the way content is exchanged between the mobile users. The PodNet design is based on pairwise associations and content exchanges between the nodes in order to maximize the likelihood that a transfer may complete over links with short durations. However, given the broadcast nature of the wireless medium, we can imagine the benefits of having a broadcast distribution for cases in which multiple nodes in an area are trying to retrieve the same content. Broadcasting techniques are compared in [24] by Williams et alii. Since contacts between nodes are short lived, it is important that the physical layer provides a high data rate to ensure that one or more meaningful data units can be transferred to completion during a single contact. We would welcome research on short range radio design for ultra-high bit rates (several Gb/s).

##### B. Link setup time

A fundamental limit in mobile content distribution is the link setup time, i.e., the time between the moment nodes move into wireless range until they are actually able to exchange data. For example, in scenarios in which the nodes are moving in opposite directions, it is crucial to minimize this setup time in order to be able to exchange at least some data before the nodes move outside communication range and the link breaks. Unfortunately popular link layer technologies are not optimized for this type of communication. In particular, Bluetooth features setup times that are at least 10.24 seconds [9], the minimal inquiry time as specified by the standard. Other technologies such as IEEE 802.11 provide shorter link setup times, however, when considering the whole network stack including higher layer protocols at the network, transport, and the application layer, the setup time still remains an issue. In PodNet, to maintain the setup time as low as possible, we make use of different techniques such as using a Bloom filter to efficiently match the channels of interest of a node with the available channels at another node. We also make use

of cross-layer optimizations by for example including higher layer information in network layer beacons. Clearly, we are still at an early stage in this area and we need to rethink in the future the communication stack as a whole to optimize the setup time. We envision the ANA framework [15] as a possible venue to redesign PodNet.

##### C. Security and trust

Publishing content in PodNet may be allowed by anybody like in the Internet. However, to regulate the right to publish content or to prevent unwanted spam, we require authentication and authorization mechanisms in place. This imposes a particular challenge as the mechanisms must also work without connectivity to any authentication or authorization servers. Further, as content will typically be user-generated as well as distributed by the users in the system and the users might not have a priori trust relationships, advanced techniques are needed to establish trust in content and its integrity (e.g., [16]).

##### D. Cooperation incentives

CCD requires cooperation between the users and users must agree to share at least some of their local resources (e.g., energy, computing, communication and storage) to others. Incentives should be made available to achieve a critical amount of cooperation necessary for a decent performance. For example, a credit-based system could be enforced to reward users that abundantly share their resources. In [17], we showed empirically, using real mobility traces, that node collaboration drastically increases the performances of content dissemination while the per-device overhead (or load) is very low and remains evenly distributed. It confirms an earlier theoretical analysis by Helgason and Karlsson [11].

##### E. Software support

The software industry for mobile devices is heterogeneous in terms of architectures and relatively in its infancy compared to desktop platforms. We observe many different existing and emerging operating systems (Symbian, Windows Mobile, Linux, Mac Os, Android, etc) with different programming and runtime environments. Also, the systems are not as open as some manufacturers do not allow code from third parties or do not document well their programming interfaces. This imposes unnecessary challenges to third-party software developers that are implementing protocols and applications which should run ubiquitously and should be able to exploit the full capabilities of the mobile devices.

##### F. Performance and control

The performance of CCD depends on the density and mobility of the users. These two factors can typically not be controlled and are given by the operational environment. In order to improve the performance, we can however place dedicated gateways that help at distributing the content. Such gateways could be fixed or mobile, as for example in public transportation systems. The question is how to most effectively place the gateways and how to manage the content cache

of such gateways in order to maximize their utility. The performance is also affected by the transmission of beacons from the mobile devices: more frequent transmissions result in more contacts, at the cost of more energy spent on this service discovery. The optimal scheme is not determined. A further issue is distributed "garbage collection". Some content might be related to a particular location or time and it should not propagate beyond the unintended scope. Also injected spam needs to be removed from the system. In practical terms, it is however a challenge as there is no central authority that manages and removes content. Hence, this must be achieved in a distributed manner by coordination of the mobile users.

## V. CONCLUSION

In this paper we have described, PodNet, our architecture for cooperative content dissemination. By successfully combining theoretical works and experimentations in a continuous feedback loop, we have found a pragmatic methodology for our research. It has allowed us to determine central issues that developers of mobile ad hoc systems have face, and it has provided us with a fresh look at the case of content dissemination by revisiting some fundamental assumptions both theoretically and experimentally. We have identified several areas of further research: Ultra-high capacity and instantaneous hand-shaking is desired for the air interface; the system needs security and trust mechanisms, incentives to cooperation, and distributed garbage collection of stale data. The software support for mobile devices must be improved and become more open. We hope that PodNet will increase interest in mobile cooperative content distribution, and will pave the way for more ad hoc communication applications.

## ACKNOWLEDGEMENT

Authors would like to thank Prof. Serge Fdida for the system view of mobile systems we have borrowed from him and expressed in the Introduction.

## REFERENCES

- [1] CRAWDAD. <http://crawdad.cs.dartmouth.edu/>, 2006.
- [2] D. Aguayo, J. Bricket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proceedings of ACM SIGCOMM'04*, Portland, Oregon, USA, October 2004.
- [3] B. Bloom. Space/time tradeoffs in in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
- [4] A. Broder and M. Mitzenmacher. Network applications of bloom filters: A survey, 2002.
- [5] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. In *Proceedings of IEEE INFOCOM*, Barcelona, Spain, April 2006.
- [6] R. Chandra, P. Bahl, and V. Bahl. Multinet: Connecting to multiple ieee 802.11 networks using a single wireless card. In *Proceedings of IEEE Infocom 2004*, Hong Kong, Mar. 2004.
- [7] M. E. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: evidence and possible causes. *IEEE /ACM Transactions on Networking*, 5(6):835–846, 1997.
- [8] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. In *ACM Transactions on Networking*, 2002.
- [9] T. B. S. I. Group. Specification of the Bluetooth system: Part B Baseband Specification, December 1999.
- [10] K. P. Gummadi, R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. *SIGOPS Oper. Syst. Rev.*, 37(5):314–329, 2003.
- [11] O. R. Helgason and G. Karlsson. On the effect of cooperation in wireless content distribution. In *In Proceedings of IEEE/IFIP WONS*, Garmisch-Partenkirchen, Germany, Jan. 2008.
- [12] M. Izal, G. Urvoy-Keller, E. Biersack, P. Felber, A. A. Hamra, and L. Garcés-Erice. Dissecting BitTorrent: Five Months in a Torrent's Lifetime. In *Proceedings of Passive and Active Measurements Conference*, April 2004.
- [13] P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, and D. Rubenstein. Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. In *Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, San Jose, CA, USA, October 2002.
- [14] D. Lam, D. Cox, and J. Widom. Teletraffic modeling for personal communications services. *IEEE Communications Magazine*, 35(2):79–87, February 1997.
- [15] Autonomic Network Architecture, ANA Project. <http://www.ana-project.org/>, 2008.
- [16] V. Lenders, E. Koukoumidis, P. Zhang, and M. Martonosi. Location-based trust for mobile user-generated contents: Applications, challenges and implementations. In *ACM/Sigmobile HotMobile Workshop*, Napa Valley, CA, USA, February 2008.
- [17] F. Legendre, V. Lenders, M. May, and G. Karlsson. Narrowcasting: An empirical performance evaluation study. In *ACM MOBICOM CHANTS Workshop*, Los Angeles, USA, Sept. 2008.
- [18] V. Lenders, M. May, and G. Karlsson. Wireless Ad Hoc Podcasting. In *Proceedings of IEEE SECON*, San Diego, CA, June 2007.
- [19] M. May, C. Wacha, V. Lenders, and G. Karlsson. Wireless Opportunistic Podcasting: Implementation and Design Tradeoffs. In *Proceedings of the ACM Mobicom CHANTS Workshop*, Montreal, Canada, September 2007.
- [20] G. Nogueira, B. Baynat, and P. Eisenmann. An analytical model for the dimensioning of a gprs/edge network with a capacity constraint on group of cells. In *Proceedings of the ACM MOBICOM*, Cologne, Germany, August 2005.
- [21] M. Papadopouli and H. Schulzrinne. Seven degrees of separation in mobile ad hoc networks. In *Proceedings of the IEEE GLOBECOM*, 2000.
- [22] C. Rolland, J. Ridoux, and B. Baynat. Hierarchical models for web traffic on cdma-lxrtt networks. Technical Report TR-RP-LIP6-2006-06-11, Networks and Performances Analysis, LIP6, France, June 2006.
- [23] C. Tudeuce and T. Gross. A Mobility Model based on WLAN Traces and its Validation. In *Proceedings of INFOCOM*, Miami, Florida, USA, March 2005.
- [24] B. Williams and T. Camp. Comparison of broadcasting techniques for mobile ad hoc networks. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, 2002.
- [25] W. H. Yuen, R. D. Yates, and S. C. Mau. Exploiting data diversity and multiuser diversity in mobile infostation networks. In *Proceedings of the IEEE INFOCOM*, 2003.
- [26] W. H. Yuen, R. D. Yates, and S. C. Mau. Noncooperative content distribution in mobile infostation networks. In *Proceedings of the IEEE WCNC*, 2003.
- [27] X. Zhang, J. Kurose, B. N. Levine, D. Towsley, and H. Zhang. Study of a Bus-Based Disruption Tolerant Network: Mobility Modeling and Impact on Routing. In *Proc. ACM MOBICOM*, September 2007.
- [28] G. K. Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley, Cambridge, Massachusetts, USA, 1949.