User-centrism in wireless networking

A dissertation presented
by

Pantelis A. Frangoudis

to
The Department of Informatics
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of

Computer Science

Athens University of Economics and Business
Athens, Greece
April 2012

Thesis advisor
Prof. George C. Polyzos

Author
Pantelis A. Frangoudis

User-centrism in wireless networking

Abstract

This dissertation delves into various aspects of user-centrism in today's wireless networking landscape. We reconsider the traditional operator/provider-centric view of wireless networking by proposing user-centric solutions to problems along three research dimensions: (i) wireless access, (ii) provision of communication services, and (iii) information collection to be used, among others, for network management purposes.

In particular, we first demonstrate cases when users become (micro-)providers of network access themselves. Community wireless networks, where community members share the wireless infrastructure and build wireless mesh networks are one example. We propose a classification of public wireless access schemes and, based on empirical data, discover power-law behavior in their structure. We then focus on the design and implementation of a Wi-Fi sharing protocol on top of resource-constrained, low-cost user devices.

Based on this Wi-Fi sharing scheme, we show how user-provided mobile multimedia services can be built in an autonomous, secure, private and decentralized way. We tackle security threats and address legal concerns by proposing a tunneling-based communication scheme with minimal dependence on centralized infrastructures for signaling. We demonstrate the feasibility of secure, user-centric VoIP services to operate on low-cost, off-the-shelf equipment, using a Quality-of-Experience-driven experimental methodology to estimate upper bounds on VoIP capacity in our architecture.

Finally, we approach issues of optimizing the operation of the wireless infrastructure in a user-centric way, where discovering the wireless topology (a first step towards reconfiguration and interference mitigation) is crowdsourced to roaming users. We design techniques to counter potential attacks by untrusted users, while we design and implement an architecture based on recently standardized technologies, such as IEEE 802.11k. We derive analytic expressions on the topology discovery accuracy of our scheme, showing that a user-centric scheme can offer more than 2× performance improvement over an infrastructure-centric scheme in realistic scenarios, even in the presence of large ratios of attackers.

Επιβλέπων
Καθ. Γεώργιος Πολύζος

Συγγραφέας
Παντελής Φραγκούδης

Χρηστοκεντρικότητα στην ασύρματη δικτύωση

Περίληψη

Το αντικείμενο αυτής της διατριβής είναι ο ρόλος του χρήστη σε διάφορα θέματα ασύρματης δικτύωσης. Η παραδοσιακή θεώρηση αντιμετωπίζει το χρήστη ως καταναλωτή και έχει τον πάροχο της υπηρεσίας στο επίκεντρο. Σύγχρονες εξελίξεις όμως φέρνουν το χρήστη στο προσκήνειο με ενδυναμωμένο χαρακτήρα και δημιουργούν τις συνθήκες για επανεξέταση του ρόλου του. Τέτοιες εξελίξεις περιλαμβάνουν την αυξανόμενη ασύρματη κάλυψη εξαιτίας ασύρματων τοπικών δικτύων που ανήκουν και ελέγχονται από οικιακούς χρήστες, ευέλικτες τεχνολογίες στο οικιακό ασύρματο δίκτυο και τερματικές συσκευές με πολλαπλές δυνατότητες. Ενδεικτική της ενδυνάμωσης των χρηστών είναι η εξάπλωση του λεγόμενου "crowdsourcing," δηλαδή της ανάθεσης ενός παραδοσιακά κεντρικοποιημένου καθήκοντος σε–πιθανόν ανώνυμα ή ψευδώνυμα–πλήθη χρηστών.

Η στροφή προς το σχεδιασμό λύσεων με το χρήστη στο επίκεντρο πηγάζει από την προσπάθεια για καλύτερη αξιοποίηση υποχρησιμοποιούμενων πόρων των χρηστών αλλά και για εκμετάλλευση της κινητικότητάς τους και των υπολογιστικών και άλλων δυνατοτήτων των συσκευών τους, ώστε να επιτυγχάνεται αυτόνομη λειτουργία, μείωση της διαχειριστικής πολυπλοκότητας και παροχή υπηρεσιών χαμηλού κόστους. Από την άλλη μεριά, η ανάγκη λειτουργίας σε μη εξειδικευμένο εξοπλισμό χαμηλού κόστους και υπολογιστικών δυνατοτήτων, χωρίς κεντρικό σχεδιασμό, εγείρει θέματα απόδοσης. Ταυτόχρονα, προκύπτουν θέματα ασφάλειας και αξιοπιστίας λόγω της ανάθεσης καθηκόντων σε εξ υποθέσεως μη έμπιστους χρήστες και έλλειψης κεντρικού ελέγχου. Αυτή η εργασία αποσκοπεί στο να προτείνει σχήματα που εκμεταλλεύονται τα πλεονεκτήματα που προσφέρει η χρηστοκεντρικότητα, ταυτόχρονα αντιμετωπίζοντας προκλήσεις αποδοτικότητας, ασφάλειας και αξιοπιστίας.

Στο πλαίσιο αυτό, έχουμε ορίσει ένα σύνολο από αρχές στις οποίες λύσεις επικεντρωμένες στο χρήστη πρέπει να είναι πιστές: (i) ο χρήστης πρέπει να παρουσιάζεται ενδυναμωμένος, (ii) το κόστος λειτουργίας να μειώνεται, (iii) η λειτουργία να είναι αποκεντρωμένη και (iv) η συμμετοχή των χρηστών ανοικτή, (v) λαμβάνοντας πάντα υπόψη θέματα ασφάλειας και αξιοπιστίας.

Η έρευνά μας στην περιοχή της χρηστοκεντρικής ασύρματης δικτύωσης τοποθετείται σε τρεις άξονες: (i) την παροχή ασύρματης πρόσβασης, (ii) την παροχή υπηρεσιών επικοινωνίας, με έμφαση στις υπηρεσίες φωνής και (iii) τη συλλογή πληροφορίας από τους χρήστες για το ασύρματο περιβάλλον, η οποία μπορεί να αξιοποιηθεί από μηχανισμούς διαχείρισης και βελτιστοποίησης της λειτουργίας του δικτύου. Στην διατριβή αυτή δίνουμε απαντήσεις σε μια σειρά από ερωτήματα σε καθένα από αυτούς τους ερευνητικούς άξονες.

Στην περιοχή της παροχής ασύρματης πρόσβασης, μελετούμε ασύρματες κοινότητες ανά τον κόσμο ως προς τη δομή και τη λειτουργία τους, προτείνουμε μια κατηγοριοποίησή τους και ανακαλύπτουμε εκθετικούς νόμους στη δομή τους. Στη συνέχεια, ασχολούμαστε με θέματα διαμοιρασμού ασύρματης πρόσβασης και ειδικότερα με τις τεχνικές λεπτομέρειες του σχεδιασμού και της υλοποίησης ενός τέτοιου πρωτοκόλλου, με στόχο την εκτέλεσή του σε περιορισμένων δυνατοτήτων οικιακό εξοπλισμό.

Με αυτό σαν βάση, προτείνουμε μια αρχιτεκτονική ασφαλών επικοινωνιών ειδικά για τέτοια περιβάλλοντα πρόσβασης. Προτείνουμε ένα σχήμα ασφάλειας βασισμένο σε τεχνολογίες "tunneling," το οποίο εξασφαλίζει την επικοινωνία πάνω από μη έμπιστα ασύρματα οικιακά δίκτυα, αλλά και λύνει ως ένα βαθμό νομικά ζητήματα που προκύπτουν όταν κανείς μοιράζεται το δίκτυό του με αγνώστους. Επίσης, μελετούμε το αντίκτυπο των χρησιμοποιούμενων μηχανισμών ασφάλειας στην ποιότητα της εμπειρίας των χρηστών υπηρεσιών φωνής, ειδικά όταν μηχανισμοί ασφάλειας υλοποιούνται σε απλό, μη εξειδικευμένο, περιορισμένο υπολογιστικά οικιακό εξοπλισμό.

Τέλος, σχεδιάζουμε και υλοποιούμε μια αρχιτεκτονική για τη συλλογή πληροφοριών για την ασύρματη τοπολογία βασισμένη στην αρχή του crowdsourcing. Ο σχεδιασμός μας βασίζεται εν μέρει σε προτυποποιημένες τεχνολογίες, όπως είναι το πρότυπο IEEE 802.11k, για το οποίο για πρώτη φορά αντιμετωπίζουμε συγκεκριμένες επιθέσεις και προτείνουμε και υλοποιούμε πρακτικές μεθόδους αντιμετώπισης. Συγκεκριμένα, εφαρμόζουμε ένα μηχανισμό φήμης για την αντιμετώπιση επιθέσεων κατά τις οποίες χρήστες υποβάλλουν ψευδείς πληροφορίες και δίνουμε αναλυτικές εκφράσεις για την απόδοση του συστήματός μας ως προς την ακρίβεια με την οποία ανακαλύπτει την ασύρματη τοπολογία. Ποσοτικοποιούμε τα πλεονεκτήματα της προσέγγισής μας σε σχέση με μια προσέγγιση στην οποία οι χρήστες δε συμμετέχουν στη συλλογή πληροφοριών (παρά μόνο τα κεντρικά ελεγχόμενα ασύρματα σημεία πρόσβασης), δείχνοντας ότι η χρηστοκεντρική προσέγγιση οδηγεί σε υπερδιπλάσια απόδοση σε ρεαλιστικές συνθήκες και για πολύ μεγάλα ποσοστά επιτιθέμενων χρηστών.

Acknowledgments

I am greatly indebted to Prof. George Polyzos, my thesis advisor. George has been instrumental in all my research efforts and his influence on me cannot be measured. He is an outstanding scientist and teacher, and an exceptional personality. He is credited with creating a unique working environment at the Mobile Multimedia Lab. It has been my pleasure and honor working under his advice.

I had the privilege of Costas Courcoubetis co-advising my thesis and helping me with his substantial suggestions and critical view. George Xylomenos, with his deep theoretical and technical expertise and thorough understanding of issues spanning across various areas of networking and computer science, also co-advised my work and was always available for discussions and advice.

The help of George Stamoulis and Vasilis Siris was invaluable; their comments and suggestions helped shape important aspects of my work. I have also had the chance to work with Giannis Marias, whom I thank for the collaboration, interesting discussions, and support.

I thank Lazaros Merakos and Iordanis Koutsopoulos, who honored me by joining my evaluation committee.

In the years I have spent at the MMlab, I met or worked with lots of interesting people; in the beginning, Chris Ververidis, Thanasis Papaioannou, Manos Dramitinos, Sergios Soursos, Panagiotis Antoniadis, Costas Calogiros, and George Thanos, and, in the years that came, Nikos Fotiou, Vaggelis Douros, Xenophon Vasilakos, Christos Tsilopoulos, Michalis Kanakakis, Ioanna Papafili, Eleni Agiatzidou, and Alex Kostopoulos.

Ntinos Katsaros, a member of the lab whom I have known since my undergraduate years and admire for his no-nonsense research approach and hard work, has been a good friend and our common interests span far beyond research.

I was working on the P2PWNC project for years, a rewarding experience. There, I had the pleasure of working with a cool group including Vasilis Kemerlis, Dimitris Paraskevaidis, Lefteris Stefanis and Stratos Dimopoulos. In this and other projects, it was always fun hacking, but also hanging out, with Fotis Elianos, Ilias Foudalis and Dimitris Zografos.

I thank Manos Panaousis, Georgia Plakia, and Stamatis Arkoulis for the collaboration.

Nafsika Kokkini, Eleftheria Nyfli, Lina Kanellopoulou, Maria Kanella, Kostas Makedos, and Charis Stais, with whom I usually shared the same working space, were always supportive and a good company, and I thank them for this.

I had the great luck of working closely with Elias Efstathiou for many years. His innovative thinking is hard to match. I consider him one of my mentors and a close friend.

Since my undergraduate studies at AUEB, I was surrounded by some great friends and colleagues: Giorgio Lucarelli, Mike Papadakis and Dimitris Galanis. Together with my other friends, they are to a great extent responsible for the fun time I have had.

Finally, I deeply thank my family for their constant support.

VITA

2012	Ph.D. in Computer Science, Athens University of Economics and Business
2005	M.Sc. in Computer Science, Athens University of Economics and Business
2003	B.Sc. in Informatics, Athens University of Economics and Business

PUBLICATIONS

Journal articles

1. P.A. Frangoudis, G.C. Polyzos, and V.P. Kemerlis, "Wireless Community Networks: An Alternative Approach for Broadband Nomadic Network Access," *IEEE Communications Magazine*, vol. 49, no. 5, pp. 206-213, May 2011.

2. E.C. Efstahiou, P.A. Frangoudis, and G.C. Polyzos, "Controlled Wi-Fi Sharing in Cities: a Decentralized Approach Relying on Indirect Reciprocity," *IEEE Transactions on Mobile Computing*, vol. 9, no. 8, pp. 1147-1160, August 2010.

3. S. Arkoulis, I. Marias, P.A. Frangoudis, J. Oberender, A. Popescu, M. Fiedler, H. de Meer, G.C. Polyzos, "Misbehaviour Scenarios in Cognitive Radio Networks," *Future Internet*, vol. 2, no. 3, pp. 212-237, 2010.

Conference and workshop papers

1. P.A. Frangoudis, D.I. Zografos, and G.C. Polyzos, "Robust client-based Wi-Fi topology discovery," Proc. 8th IEEE Consumer Communications and Networking Conference, Las Vegas, NV, January 2011.

2. E. Dimopoulos, P.A. Frangoudis, and G.C. Polyzos, "Exploiting super peers for large-scale peer-to-peer Wi-Fi roaming," Proc. IEEE Globecom 2010 Workshop on Advances in Communications and Networks (User-Provided Networking session), Miami, FL, December 2010.

3. P.A. Frangoudis and G.C. Polyzos, "Report-based topology discovery schemes for centrally-managed Wi-Fi deployments," Proc. NGI 2010, Paris, France, June 2010.

4. F.A. Elianos, G. Plakia, P.A. Frangoudis, and G.C. Polyzos, “Structure and Evolution of a Large-Scale Wireless Community Network,” Proc. IEEE WoWMoM 2009, Kos, Greece, June 2009.

5. E.A. Panaousis, P.A. Frangoudis, C.N. Ververidis, and G.C. Polyzos, “Optimizing the Channel Load Reporting Process in IEEE 802.11k-enabled WLANs,” Proc. IEEE LANMAN 2008, Cluj-Napoca, Transylvania, Romania, September 2008.

6. P.A. Frangoudis, G.C. Polyzos, “Coupling QoS Provision with Interference Reporting in WLAN sharing Communities,” Proc. IEEE PIMRC 2008 Workshops, Cannes, France, September 2008.

7. K. Katsaros, P.A. Frangoudis, G.C. Polyzos, and G. Karlsson, “Design Challenges of Open Spectrum Access,” Proc. IEEE PIMRC 2008 Workshops, Cannes, France, September 2008.

8. V.G. Douros, P.A. Frangoudis, K. Katsaros, and G.C. Polyzos, “Power Control in WLANs for Optimization of Social Fairness,” Proc. 12^{th} Pan-Hellenic Conference on Informatics (PCI 2008), Samos, Greece, August 2008.

9. P.A. Frangoudis, V.P. Kemerlis, D.C. Paraskevaidis, E.C. Efstathiou, and G.C. Polyzos, “Experimental Evaluation of Community-based WLAN Voice and Data Services,” Proc. 3^{rd} International Mobile Multimedia Communications Conference (MobiMedia 2007), Nafpaktos, Greece, August 2007.

10. P.A. Frangoudis and G.C. Polyzos, “Peer-to-Peer Secure and Private Community Based Multimedia Communications,” Proc. 2^{nd} IEEE International Workshop on Security and Pervasive Multimedia Environments (MultiSec 2006, in conjunction with IEEE ISM), San Diego, CA, December 2006.

11. E.C. Efstathiou, F.A. Elianos, P.A. Frangoudis, V.P. Kemerlis, D.C. Paraskevaidis, E.C. Stefanis, and G.C. Polyzos, “Public Infrastructures for Internet Access in Metropolitan Areas,” Proc. 1^{st} International Conference on Access Networks (AccessNets 2006), Athens, Greece, September 2006.

12. E.C. Efstathiou, P.A. Frangoudis, and G.C. Polyzos, “Stimulating Participation in Wireless Community Networks,” Proc. IEEE INFOCOM, Barcelona, Spain, April 2006.

Demo papers

1. E.C. Efstathiou, F.A. Elianos, P.A. Frangoudis, V.P. Kemerlis, D.C. Paraskevaidis, G.C. Polyzos, and E.C. Stefanis, "Practical Incentive Techniques for Wireless Community Networks," 4th International Conference on Mobile Systems, Applications, and Services (MobiSys 2006) Demo Session, Uppsala, Sweden, June 2006.

2. E.C. Efstathiou, F.A. Elianos, P.A. Frangoudis, V.P. Kemerlis, D.C. Paraskevaidis, G.C. Polyzos, and E.C. Stefanis, "Building Secure Media Applications over Wireless Community Networks," Proc. 13th Annual Workshop of the HP Openview University Association (HP-OVUA) Poster/Demo Session, May 2006.

3. E.C. Efstathiou, F.A. Elianos, P.A. Frangoudis, V.P. Kemerlis, D.C. Paraskevaidis, G.C. Polyzos, and E.C. Stefanis, "The Peer-to-Peer Wireless Confederation Scheme," IEEE INFOCOM 2006 Demo Session, Barcelona, Spain, April 2006.

4. E.C. Efstathiou, F.A. Elianos, P.A. Frangoudis, V.P. Kemerlis, D.C. Paraskevaidis, G.C. Polyzos, and E.C. Stefanis, "The Peer-to-Peer Wireless Confederation Scheme: Protocols, Algorithms, and Services," Proc. 2nd International IEEE/Cre-ateNet Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom 2006) Demo Session, Barcelona, Spain, March 2006.

Posters and abstracts

1. G.C. Polyzos, G.F. Marias, S. Arkoulis, P.A. Frangoudis, M. Fiedler, A. Popescu, H. de Meer, R. Herkenhöner, A. Fischer, and J.O. Oberender, "ASPECTS: Agile spectrum security," Proc. NGI 2011, Kaiserslautern, Germany, June 2011.

2. P.A. Frangoudis, D.I. Zografos, and G.C. Polyzos, "Secure Interference Reporting for Dense Wi-Fi Deployments," Proc. ACM CoNEXT 2009 Student Workshop, Rome, Italy, December 2009.

3. S. Arkoulis, M. Fiedler, P.A. Frangoudis, R. Herkenhoner, G.F. Marias, H. de Meer, and G.C. Polyzos, "Distributed Spectrum Sensing for Spectrum Agility: Incentives and Security Considerations," Proc. 1st Euro-NF Workshop on Future Internet Architectures, Paris, France, November 2008.

4. P.A. Frangoudis, S. Arkoulis, G.F. Marias, and G.C. Polyzos, "Incentives and Security Considerations in Distributed Spectrum Sensing," Proc. 1st Euro-NF Socioeconomics Workshop, Athens, Greece, October 2008 (extended abstract).

5. N.N. Leontiadis, V.P. Kemerlis, P.A. Frangoudis, and G.C. Polyzos, “Secure Network Management Using a Key Distribution Center,” Proc. 2007 Workshop of the HP Software University Association (HP-SUA) Poster Session, Garching/Munich, Germany, July 2007.

6. P.A. Frangoudis, E.C. Efstathiou, and G.C. Polyzos, “Reducing Management Complexity through Pure Exchange Economies: A Prototype System for Next Generation Wireless/Mobile Network Operators,” 12^{th} Annual Workshop of the HP Openview University Association (HP-OVUA), Porto, Portugal, July 2005.

Contents

List of Figures

List of Tables

To my parents, Mina and Andreas, and my brother Kostis.

Chapter 1

Introduction

In this Chapter we lay the background for our work, focusing on recent advances and shifts towards putting the user at the center of the wireless networking landscape. We present our thesis and contributions with respect to user-centric wireless networking, which we approach along three research dimensions: (i) utilization of user-provided infrastructure, where networks are built based on the contributions of individual micro-providers[1], (ii) development of *service* architectures designed with such user-provided wireless networks in mind, and (iii) activating users as providers of information about network topology and conditions, and, in particular, as to the conditions in the radio environment, which can be vital input for a diverse set of applications and services (ranging from frequency planning to wireless positioning services).

1.1 Background

The traditional view of communications has recently been disrupted by the evident user empowerment in all aspects of the communication process. Traditionally, the operator-centric view dominated, where users had a passive role as service consumers. This view seems to change and the factors that have led to this shift are numerous.

Hand-in-hand with the revolution in current Internet usage trends, where we witness a vast increase in the volume and popularity of user-generated content, a new communication paradigm, where users have a central role, has emerged. With the advent of low-cost, ubiquitous, and easy to install and configure wireless equipment, but most importantly, protocols that operate in unlicensed spectrum, users can effectively acquire the dual role of becoming service consumers and *providers* at the same time. This fact has the potential of bringing up new disruptive technologies, where users can enjoy low-cost wireless connectivity via the infrastructure provided by a

[1] The terms *micro-provider* and *micro-operator* will be used interchangeably and denote individuals who take up the role of the access service provider at a small scale, using their private equipment and resources.

heterogeneous crowd of micro-operators. As a matter of fact, communities of users who use low cost wireless equipment for free interconnection, without the need for a provider, have emerged since the early 2000s [40], when the potential of standards for Wireless Local Area Network (WLAN) connectivity began to stand out.

The following key developments and observations are indicative of the shift towards a user-centric view of networking, but also form the basis to propose user-centric approaches to issues that were traditionally tackled in a more strict, centralized, and operator-oriented manner.

Increased user-based wireless coverage WLANs often cover significantly larger area than intended. With their proliferation, Wi-Fi signals pervade modern densely populated urban areas. Most of these networks are managed by individuals, adding to those operated within corporate premises, campus environments, or other public spaces. Coverage is such that is is possible to build a business based just on the presence of Wi-Fi networks and not their communication capabilities. For instance, Skyhook [112], among others, offers a GPS-less Wi-Fi-based positioning service, also applicable to indoor environments where GPS is not available, and with often better performance than GPS in very dense urban environments. Increased Wi-Fi coverage gives rise to the question of whether such user-provided infrastructure could be harnessed to offer a low-cost, ubiquitous wireless access solution that could rival or complement (as a best-effort alternative) 2G/3G/4G cellular services.

Flexible technologies for the wireless home network The technical means to answer the question whether user-based wireless access can offer alternatives to traditional cellular services exist. Off-the-shelf wireless equipment for the home network, available at low cost, is capable of performing far more tasks than simply forwarding user traffic to/from the Internet via a fixed broadband connection. Home wireless routers powered by open-source software have both the necessary flexibility to install custom software, but also potential spare memory and CPU cycles to run more demanding applications.

Versatile technologies at the user end On the other hand, handheld devices have far more capabilities than accomplishing cellular-based phone calls or exchanging SMS text messages. Instead, apart from the apparent increase in processing power, they come equipped with multiple network interfaces (Wi-Fi, 3G, 4G, Bluetooth), high-quality displays, but also with versatile sensing devices, such as motion sensors, cameras and GPS receivers. Also, flexible operating system platforms, such as Google Android and Windows Phone 7 have enhanced application development for mobile environments. The sensing and communication capabilities of modern devices, combined with the inherent user mobility make them powerful platforms to acquire and communicate information about user environment and context, thus giving rise to *crowdsourcing.*

The rise of crowdsourcing Crowdsourcing is a term coined by Jeff Howe as of 2006 [52] to describe a shift towards exploiting user capabilities to “outsource” tasks to (potentially anonymous) crowds; a variety of tasks that would traditionally need significant investment on infrastructure and time can now be delegated to lots of users, who can use their mobile devices with advanced communication, computation and sensing capabilities. Such tasks range from urban sensing [24] to collecting information about the radio environment that can be used for optimization purposes. A significant part of our work focuses on the latter.

User-centrism in Internet usage We make a final observation that is not directly related with the body of our work, but showcases user empowerment, as far as Internet usage is concerned. It is evident that user-generated content (and not only user-distributed, as is the case for content distributed without the mediation of its originator, e.g., software, movies or music exchanged using peer-to-peer technologies like BitTorrent) takes up a significant share in today’s Internet traffic. Such traffic includes web-accessible multimedia content (e.g., videos captured or authored by users, photos, etc.) and interactions via social networking media.

One should note that building infrastructure and services based on user contributions poses a new set of challenges, which are addressed in this work. First, user-provided equipment is typically inexpensive and resource-constrained, standing at the opposite side of powerful and costly infrastructure deployed by operators. Second, users often do not have the technical expertise, cannot spare the resources and time, and lack central control when it comes to configuration decisions, optimizations and planning operations. Therefore, user-provided services, be they connectivity-related, application-oriented, or having to do with content and information provision are assumed to operate in a self-organizing, best-effort style, also due to the unpredictability and variability in user behavior and participation. Designing and implementing user-centric protocols, mechanisms and services, thus, involves addressing significant performance challenges.

Furthermore, crowds of users pooling their resources, executing a distributed task, offering a service or collecting information, cannot always be assumed trustworthy. They are typically not legally bound by contracts and service-level agreements, and are expected to behave strategically, without excluding the potential of purely malicious behavior. In the environment we envisage, user identification is not always to be assumed strong; while this enhances anonymity and privacy and reduces identity management overhead, it calls for careful design of mechanisms and additional protection measures.

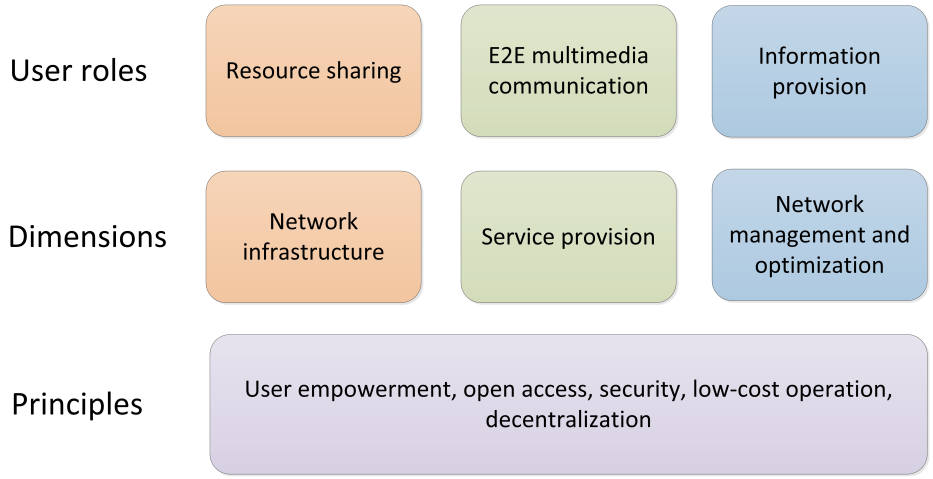
User roles
Resource sharing
E2E multimedia communication
Information provision
Dimensions
Network infrastructure
Service provision
Network management and optimization
Principles
User empowerment, open access, security, low-cost operation, decentralization

Figure 1.1: Aspects of user-centric wireless networking: User empowerment is evident along the dimensions of user-provided wireless access, service provision and assisting in optimizing network operation. Across all dimensions, our work adheres to the principles of decentralization, open access, security and low-cost operation.

1.2 Thesis

We focus on various expressions of user-centrism in wireless networking. In particular, we approach the issues of (i) wireless access, (ii) multimedia service provision, and (iii) information provision for infrastructure reconfiguration from a user-centric perspective. Our thesis is that users can become active as providers, at the access, service, and information levels. We present cases where such behavior has emerged and led to the development of infrastructure sharing wireless communities and then we develop and evaluate user-provided multimedia communication services on top of such communities. At the same time, we put the user at the center of the process of collecting information about the conditions in the radio environment, an important first step before network reconfiguration mechanisms are put in effect to improve performance. Figure 1.1 presents our overall view towards user-centric networking.

1.2.1 Research axes

Our work is positioned along the following three research dimensions.

User-provided wireless infrastructure and network access

Since the emergence of the IEEE 802.11 family of standards, and as the technology matured, a trend towards open wireless access has been evident. Operation in unlicensed bands and the low cost of WLAN equipment, coupled with the enthusiasm and self-organizing spirit of some users have helped towards the direction of building community-based wireless access schemes for open connectivity. We believe that it is important to document this manifestation of network building and get a deep understanding on the internal workings of such communities, both technology-wise, but also from a socioeconomic perspective. It is also significant to attempt to model some of their structural properties in a formal way, based on empirical data, i.e., based on examples of existing large-scale wireless community networks. This will facilitate the evaluation of services and protocols to be deployed on top of them.

User-centric multimedia services

Prior work [29] suggests that user-centric wireless networks may be the answer to the question of providing low-cost open wireless access, complementing cellular Internet services (e.g., 2/3G) for nomadic users. Designing architectures to provide secure end-to-end communication for user-provided wireless networks would add value to a user-centric wireless access scheme and could probably be a necessity in order for the latter to achieve wide adoption. For such architectures, achieving acceptable Quality of Experience (QoE) for multimedia services, while at the same time offering strong security and privacy levels, as well as operating at a low cost, on top of commodity wireless equipment are challenges which this work faces.

User feedback for optimized network management

Our position is that users should have an active role in the process of optimizing the operation of the wireless networks they access. The wireless access schemes we focus on are based on the premise of operation in unlicensed bands. Due to the density of WLAN deployments and unlicensed spectrum scarcity, the problem of interference has recently been aggravated and calls for sophisticated interference mitigation schemes. In order to provide optimized operation, it is important to acquire feedback as to how users perceive wireless coverage and interference. Thus, a user-centric scheme for collecting radio information at client spots is necessary and would offer significant advantages compared to infrastructure-centric schemes, where spectrum measurements are carried out solely by APs (Access Points). We propose an architecture for collecting wireless topology/coverage information, where the task of monitoring is crowdsourced to roaming users. Users, however, should not be considered trustworthy and their feedback should be carefully evaluated as to its validity, since they could engage in fraudulent reporting. For specific types of centrally managed wireless deployments, we believe that such reporting attacks could be countered by simple consensus-based schemes, effectively filtering fraudulent information. Trustworthy measurements by the managed infrastructure and the application of a reputation-based scheme to appropriately weigh user feedback assist in more accurately discovering Wi-Fi topology, tackling reporting attacks.

1.2.2 Principles

Our approach to user-centric wireless networking is based on the common ground of a set of principles which the solutions we propose should adhere to. These principles are stated below.

The user at the center

Throughout this work, our prime principle is that the problems we address should be viewed under a user-centric perspective and the solutions we propose should promote the role of users, exploiting and showcasing user-empowerment.

Open access and participation

As far as network access is concerned, participation and use of the resources of the community is open to all users. A user can join the community by contributing his resources to the pool with the common goal of achieving wider network coverage and enjoying low-cost access to the infrastructure when nomadic. To facilitate organic growth and lift entry barriers to become a service provider, joining the user-centric network should not involve the complexity of setting up contracts, while a loose and decentralized user identification scheme is desirable. Provision of wireless network access is also hindered by the need to acquire a license, a financial and managerial burden impossible for a user to handle. Therefore, operation in unlicensed spectrum is mandated[2]. Beyond wireless access issues, any user-centric solution should facilitate and promote open and voluntary user participation.

Decentralization and distribution of tasks

This principle emerges in any aspect of user-centric networking we have considered in this work. Even though some functions (such as addressing and naming) may be centrally managed, decentralization naturally emerges in user-provided wireless networks at many layers. First, they grow organically by the contributions of individuals. Second, in the case of community wireless mesh networks, they operate using decentralized protocols for routing (more or less as is the case for the Internet itself). Content and services are user-provided, while end-to-end communication can be achieved in a decentralized manner (see Section 4). As we show in Chapter 5, information vital for wireless network management operations can be collected in a decentralized manner, again empowering the role of users.

[2] Even if operating in unlicensed spectrum, though, a separate license would probably be necessary in order to provide commercial telecommunication services.

Security, trust and user rationality

When we design mechanisms and protocols, we do not assume that users are benevolent. Often, one needs to deal with rational behavior which can lead a user to strategic decisions that may violate protocols, or even with pure malice. In any case, such user strategies can lead to attacks which the system designer should tackle. For example, in a user-provided networking setting, trust between service providers and consumers should not be assumed. Communication should be secure and private to avoid attacks by untrusted peers. In the case of collecting information from users, submitted feedback may not be trustworthy. Therefore, effective mechanisms should be in place to evaluate reported information, filter potential fraudulent reports and ensure that potential attackers cannot lead the system to a bad operating point.

Low-cost operation

A key principle in user-centric networking is operation at low cost. This is mandated by the need to amass infrastructure based on private resource contributions of individuals who own and operate inexpensive home (not professional) equipment. Decentralized service architectures can then be laid over this infrastructure and accessed for free and in a peer-to-peer manner. This can enable, for example, free voice and multimedia communication over the user-provided wireless network. This principle is in sharp contrast with the ISP-centric viewpoint; instead of a few large providers, many *micro-operators* could offer a complementary best-effort service at minimal cost.

1.3 Contribution

The main contributions of this dissertation are as follows:

– We study the emergence and evolution of Wireless Community Networks as a prime example of placing the user at the center of wireless access provision. In particular, we provide a classification of such public wireless access schemes in two different dimensions, i.e., with respect to their architectural properties but also the initiatives that drive their emergence and operation. We provide insight as to the reasons that led to their emergence, sustainability and success and, based on historical observations and present-day circumstances, attempt to make predictions for their future.

– Based on data from existing large Wireless Community Networks, we discover that some of their structural properties present power-law behavior; this observation can prove useful for deriving realistic structural models to be used in the performance evaluation of services over such networks.

- Focusing on community networks that can be built on the private contributions of individual WLAN owners by sharing Internet access via their private hotspots, we design, implement and experimentally evaluate user-centric, decentralized, secure and private multimedia services. Using established Quality-of-Experience evaluation methodologies, we show that even with off-the-shelf WLAN equipment, it is possible to set up VoIP calls in a decentralized manner, minimally relying on a centralized rendezvous infrastructure, using strong security mechanisms and offering a level of location privacy to communicating endpoints, protecting (i) call participants from eavesdropping by the provider offering Internet access, and (ii) the provider himself from potential malicious activities on behalf of anonymous visitors.

- We demonstrate the crucial role that users can play in the process of optimizing the operation of the wireless infrastructure by designing, implementing and evaluating a user-centric architecture for collecting wireless coverage measurements, which can be used by the network operator for planning purposes. We propose a crowdsourcing approach, where users are requested to report on the wireless conditions at their vicinity.

- We tackle simple, yet realistic attacks to the reporting process, which are rooted in the fact that users cannot be assumed trustworthy. Our countermeasures are based on consensus-based rules for evaluating user feedback. We propose a reputation-based scheme which is tailored to centrally-managed wireless deployments and show it to work well under various conditions.

- We implement a proposed topology discovery scheme on top of off-the-shelf equipment and making use of standard protocols for authentication, security and reporting. In particular, we utilize the IEEE 802.11i framework for user authentication, authorization and accounting and implement a subset of the IEEE 802.11k protocol for reporting radio resource measurements. We also demonstrate that the attacks we study are feasible by implementing them in modern Linux kernels.

- We analytically calculate the performance of the proposed user-centric scheme and show it to significantly outperform AP-centric schemes, even in the presence of large numbers of attackers. We have selected model parameters based on data publicly available, as well as our own measurement campaign. We present results on the topology discovery accuracy of our scheme for scenarios with varying numbers of attackers and varying AP and client densities.

1.4 Dissertation outline

The remainder of this dissertation is organized as follows: Chapter 2 presents an overview of the state of the art in relevant research fields. Chapter 3 focuses on user-provided wireless networks, where individuals share their infrastructure and resources, becoming micro-providers and building wireless communities. The autonomous and decentralized spirit which characterizes wireless communities is adopted in the design we propose, implement and evaluate for secure multimedia services over user-provided networks, as presented in Chapter 4. Chapter 5 deals with our third research dimension, namely involving users in collecting information about the radio environment in a robust manner, in order to apply it, among others, for optimizing the operation of the wireless infrastructure. Chapter 6 discusses various aspects of our work which can lead to issues for further research, while Chapter 7 concludes with a summary of our work and contributions in the area of user-centric wireless networking.

Chapter 2

State of the art

This thesis approaches three different issues whose common theme is that the user is at the center and provides the wireless infrastructure, services, or information. Another common point is the fact that we deal with security issues that may emerge when dealing with self-interested users, for whom adherence to protocols cannot readily be assumed. By its very nature, the area of user-centric networking is tangent to a large set of heterogeneous research areas, including wireless networking, security and trust, multimedia service provision, performance evaluation of services over wireless links, and more. Figure 2.1 shows that our user-centric approach to wireless networking lays at the intersection of various research fields. This section aims to provide a short overview of the state-of-the-art in each of these relevant fields.

2.1 User-provided wireless networks

User-provided wireless access schemes have recently received research and commercial attention. Based on the private contributions of individuals who operate Wi-Fi equipment, architectures and systems are being proposed with the aim of building resource sharing communities to achieve wide wireless coverage. Chapter 3 provides a thorough survey on such initiatives and their operation.

Originally, research in the area focused on the technical aspects of building wireless community networks, but attention soon shifted towards socioeconomic and incentive aspects. A critical aspect of user-provided wireless networks is to design mechanisms to encourage contribution while limiting attacks by selfish users who aim at *free riding*. Efstathiou [30] proposed a fully decentralized scheme to this end. We provide details on it and build a secure multimedia communications architecture on top of it in Chapter 4.

FON [37] has followed a different architectural approach. It acts as a mediator for the development of a Wi-Fi sharing community, centrally dealing with user authentication and accounting. The role of such mediators as community providers is modeled by Biczók et al. [15, 14]. They analyze their interactions with users and

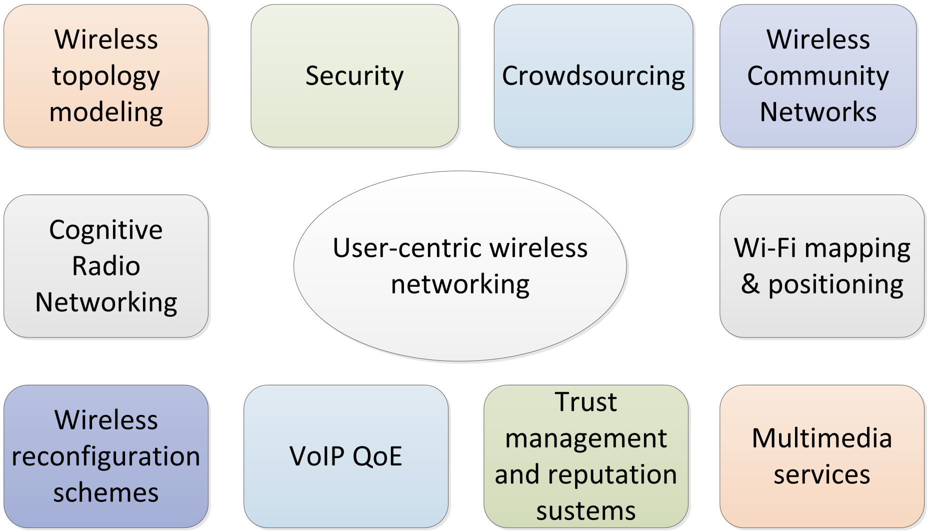
Wireless topology modeling
Security
Crowdsourcing
Wireless Community Networks
Cognitive Radio Networking
User-centric wireless networking
Wi-Fi mapping & positioning
Wireless reconfiguration schemes
VoIP QoE
Trust management and reputation sustems
Multimedia services

Figure 2.1: User-centric wireless networking is a research field relevant with various heterogeneous areas.

ISPs in global-scale wireless community networks and explore the space of available parameters (e.g., roaming cost, ISP's profit share) to determine the benefits of each player when joining the community. The authors present interesting results regarding the role of ISPs: Arguing that ISP endorsement is important for the global scaling of wireless communities, they find that depending on parameters set by the mediator, they will either fully support or abandon (i.e., prohibit Wi-Fi sharing) the community. This conclusion appears to be closely related to the terms of use adopted by ISPs regarding broadband connection sharing over Wi-Fi.

Two significant issues pertinent to wireless communities are studied by Manshaei et al. [80]. First, they study how initial community network coverage and user payoffs and fees affect the evolution of the community. Second, they focus on the competition between licensed wireless access providers and community-based ones, which is an important step towards answering whether wireless communities can be a viable alternative (or complement) to licensed cellular networks.

Ai et al. [4] focus on a critical usability aspect; our approach [30, 31], but also FON [37], require software installation at the AP side, while their scheme only requires software updates at the client side. However, their work requires a central server, which violates the decentralization requirement of our design.

Ben Salem et al. [106] propose a scheme for Wi-Fi roaming in which WISPs have multilateral roaming contracts and register with a central authority that maintains reputation records derived from QoS reports submitted by roamers.

The legal aspects of user-provided wireless access should not be neglected. Sithigh [116] discusses how the adoption of open wireless access is hindered by a diverse set of legal provisions. Legal issues can influence user participation and the adoption of community wireless technologies and should be carefully considered in system design. Often, the terms of service of ISPs explicitly state that sharing one's broadband Internet connection is prohibited, although sharing-friendly ISPs exist [32]. Also, sharing one's Internet access with anonymous visitors may hold the host liable for potential malicious activities by them. Our technical approach to this problem [42, 31] (see Chapter 4) is based on tunneling: Users visiting untrusted Wi-Fi hotspots can tunnel their Internet traffic through VPN gateways. On one hand, they are protected from eavesdropping and, on the other hand, Wi-Fi owners are not liable for illegal activities by visitors, since the attacker can be traced back to the VPN gateway. Sastry et al. [108] and Heer et al. [47, 46] present similar solutions.

2.2 Reputation systems

In environments where no assumption can be made about user "quality," and conformance to protocols and etiquette cannot be assumed, mechanisms need to be in place to evaluate user behavior, build *trust* and encourage cooperation. Each interaction involves hidden information about the quality of prospective participants. A typical approach is to encode this information in a *reputation* value assigned to each entity, which is an indication of the expected quality of the entity's behavior. For example, in electronic marketplaces such as eBay [28], prospective buyers check the reputation of sellers (and vice versa) before they bid for an item. Reputations are updated based on user feedback.

Ensuring honest feedback is a significant challenge. Papaioannou and Stamoulis [99] propose a mechanism that encourages truthful ratings in a peer-to-peer system by applying punishment in the event that ratings between two transacting peers do not match.

Apart from electronic marketplaces, reputation mechanisms appear to be suitable for a wide application domain, including spam filtering [103, 111], packet forwarding in mobile ad hoc networks [19], and resource allocation [109], among others.

Reputations come to play in various aspects of our work, especially when it comes to evaluating user supplied information and user behavior. In Chapter 4 we assume a community-based Wi-Fi sharing scheme where access decisions are based on the subjective estimation of consumer reputations (expressed in terms of service to the community). In Chapter 5 we apply a reputation-based mechanism to weigh user reports in a crowdsourced wireless topology discovery scheme: Reputations are a measure of user trustworthiness, thus reports by reputable users have more weight. Trusted reports by the wireless infrastructure are used as an additional means of validating user reports. We apply an adaptation of a reputation metric proposed by Papaioannou and Stamoulis [98], where user reputation is updated as follows:

$$r' = \beta r + (1 - \beta)\mathbf{1}(success), \tag{2.1}$$

where r' is the updated reputation value, β is the discounting factor for past transactions and $\mathbf{1}(.)$ is the indicator function. In our case, in the place of the indicator function, we use the user's *score* at each reporting round (see Chapter 5).

2.3 Wireless reconfiguration and self-optimization schemes

2.3.1 Topology modeling and discovery

Topology representation

Mechanisms and algorithms that aim at optimizing wireless network operation, but also location-based services, Wi-Fi-based positioning schemes, and handover planning, require information about network topology, that is, information about the location and neighborhood of entities (base stations and users). Part of our work focuses on building such an information database with robustness in a user-centric manner. We designed a topology discovery architecture with mechanisms such as channel assignment in mind, but we opted for a generic topology representation. In particular, we model Wi-Fi topology as a weighted undirected *Coverage Graph* (see Section 5.3). Our coverage graph model is an adaptation of the model introduced and applied to the channel assignment problem by Mishra et al. [85].

Depending on the application, various alternative topology representations are possible, though. To address interference asymmetry between APs and to be able to capture client and AP load, necessary for performing power control, Ahmed and Keshav [3] use an *annotated conflict graph*, which includes additional client vertices, undirected client-AP *association* edges and directed *interference* edges. Another approach [86] is to apply a *conflict set coloring* formulation to the problem of jointly performing channel assignment and load balancing, where, for each client, there is a *range* set (APs in range) and an *interference* set (APs not in range, but with interfering clients associated to them), and the objective is to minimize interference suffered by each client. Jain et al. [66], finally, represent interference by modeling a link between two nodes as a graph vertex and placing an edge between two vertices if the respective links are conflicting.

Stochastic geometry models

To analyze the performance of the Wi-Fi topology discovery scheme we have proposed, we used a stochastic geometry approach. There are numerous such models [7], but we have resorted to homogeneous spatial Poisson Point Processes and an idealized cell coverage model; each AP covers a disk of fixed radius R and each terminal in range can sense the AP. Although this model captures the case for devices with the same fixed transmission power and a free space propagation model, it ignores the effects of fading and shadowing, as well as the case for heterogeneous receivers with different thresholds for sensing the presence of a transmission. Its simplicity, however, made our analysis more straightforward.

2.3.2 WLAN Self-optimization mechanisms

Numerous approaches aim at adding reconfiguration features to Wi-Fi networks. Their common denominator is the need to collect information from the wireless environment. The next step is to apply sophisticated reconfiguration mechanisms by means of frequency selection [85, 86, 70], power control [3, 82], rate adaptation [69], adaptation of the carrier sensing threshold [82, 122], or their combinations. Murty et al. [87] focus on enterprise WLANs where most wireless management decisions are pushed to the infrastructure. Again, they need measurements from clients and APs to perform them. Our work serves in improving the robustness of information collection and providing valid input to the above mechanisms. It should be noted that most of the above schemes [85, 86, 69, 122, 87] require client participation for the collection of input for the respective spectrum sharing mechanisms.

2.4 Crowdsourcing

With the increased sensing, computation and communication capabilities of mobile devices, a diverse set of user-centric applications is possible. Exploiting user mobility and device capabilities, tasks that would require significant investment and personnel costs, if executed in an operator/provider-centric manner, can be delegated to roaming crowds. While crowdsourcing is attractive and is being applied to many different application scenarios, we focus the following discussion on crowdsourced performance and network coverage related measurements, which are more relevant to our work.

When it comes to measurements, cost savings aside, a crowdsourcing approach also provides a user-centric view of the measured conditions. For wireless and mobile networks, it is important to couple performance with user locations. Yao et al. [127] have identified the benefits of utilizing “bandwidth maps” that encode the download capabilities of the mobile broadband network per location to adaptively stream multimedia content to fast-moving users, but also to achieve efficient multi-homing, showing the value of historical observations about mobile network performance at a single location. Constructing such maps naturally lends itself to a crowdsourcing approach. With a similar motivation, Pang et al. [96] present a collaborative service that offers information about AP capabilities, which can be used for improved AP selection. This information is built by crowdsourced reports. They focus on preserving user privacy and try to limit fraudulent reporting.

In our work, we focus on crowdsourcing with the aim of exposing Wi-Fi topology.

Delegating this task to mobile users is a concept that has already been explored in the literature, and is supported by recent standardization activities, such as the IEEE 802.11k [56] standard, which was ratified in 2008. It specified a set of extensions to the IEEE 802.11 standard that provide radio resource measurement functionality. We are particularly interested in requesting and receiving information on the wireless neighborhood of stations, but the standard is much richer. It focuses on defining measurements to be carried out by WLAN devices and provides interfaces to upper layers to request and receive them. What it does not specify, however, is how to utilize this information, which is a higher layer issue. We show how this information can be used to optimize the operation of a managed WLAN deployment by means of channel assignment, but other options are possible. The standard supports requesting more refined measurements which can be used by appropriate sharing and reconfiguration schemes, but also location information from stations, which could be used to offer location-based services. For example, Hermann et al. [49] use IEEE 802.11k reports about user location and neighbor APs to approximate the coverage area of each BSS. Details on carrying out measurements are left to device manufacturers or firmware and driver implementors. For example, in our prior work, we focused on carrying out channel load measurements in an optimized manner, so that accurate values on channel load could be derived while monitoring time is reduced [95]. Our current work is centered around the security and robustness aspects of Wi-Fi topology discovery, a topic rather neglected in the literature.

Our work is also related to the process of distributed spectrum sensing in Cognitive Radio Networks (CRN). In a typical CRN scenario, *secondary* (i.e., unlicensed) users collectively monitor spectrum usage to detect the presence or absence of *primary* (i.e., licensed) ones. Recent standardization efforts within the IEEE 802.22 working group [58] also focus on spectrum sensing. In this context, Chen et al. [20] study two potential attacks, namely Incumbent Emulation, where an adversary's CR transmits signals that emulate the characteristics of a primary user's transmissions, and Spectrum Sensing Data Falsification. In the latter, which is similar in spirit with the attacks we address, adversaries submit fake sensing data to the collecting entity to tamper with the sensing decision. Fatemieh et al. [36] also focus on securing against fake reporting and propose a weighted aggregation process for crowdsourced spectrum reports, a principle that we also adopt in our context.

For purposes of informing users about Wi-Fi coverage, but also for positioning purposes, many web sites host maps of Wi-Fi networks. WiGLE [124] is such an example, where users can scan for Wi-Fi presence using software like NetStumbler and submit their reports to a global database, or even register APs manually. It is not clear how WiGLE tackles fake reporting. Skyhook Wireless [112], on the other hand, taking advantage of the increased Wi-Fi coverage in metropolitan areas, offers a Wi-Fi-based positioning service using a large beacon database developed by "wardriving." A client-driven scheme for updating and expanding this database offers significant coverage benefits. In fact, Tippenhauer et al. [119] demonstrate that Skyhook's *Wi-*

Fi Positioning System (WPS) exploits user reports and show how such a system can be attacked. Notably, they describe a database poisoning attack where users submit fraudulent information about APs in their vicinity and propose database update rules to mitigate this threat.

2.5 Quality-of-Experience for VoIP services

Part of our work focuses on measuring the quality of user-centric VoIP services provided over a community-based wireless access scheme. To this end, we apply a Quality-of-Experience driven evaluation methodology to estimate the VoIP capacity of the proposed architecture in a user-centric manner. In this section we provide a classification of VoIP quality assessment methodologies and delve into the details of the assessment scheme we have chosen. As will shall show, we opted for a scheme which attempts to estimate user-perceived VoIP quality based on measurable quantities (delay and packet loss) instead of resorting to subjective quality ratings by human subjects.

2.5.1 A classification of assessment methodologies

Numerous models for assessing the quality of voice services have been proposed and the ITU-T has been at the forefront of such standardization efforts. Jelassi et al. [67] provide a detailed review of relevant approaches.

Quality assessment models can be classified along two dimensions, i.e., (i) subjectivity and (ii) need for an original reference signal.

Subjective methodologies involve experiments with human subjects who rate the voice (or conversational) quality following the Absolute Category Rating (ACR) method, where each subject rates quality on the 1-5 scale (1 represents the worst quality, while 5 is the score for a perfect call). The average rating is referred to as the Mean Opinion Score (MOS). The ITU-T describes methods and procedures for conducting subjective evaluations of transmission quality in recommendation P.800 [60].

Objective methodologies, on the other hand, attempt to estimate user-perceived voice quality by objective measurements, without involving human subjects. The ITU-T P.862 recommendation describes the well-know PESQ (Perceptual Evaluation of Speech Quality) model [62], where an original input signal is compared to a degraded output signal as a result of its transmission through a communication system. PESQ derives quality ratings using psychoacoustic fundamentals.

ITU-T P.862 needs an original reference signal to operate on. Such methodologies are termed “full reference” or “double-ended.” ITU-T Recommendation P.563 [63], on the other hand, describes a “no reference” methodology aiming at inferring voice quality by measurements at a single point in the mouth-to-ear path.

While fairly accurate, methodologies such as PESQ, which are based on signal comparisons, fail to capture the effects of various operating parameters (e.g., codec settings, jitter buffer implementation) and phenomena in the end-to-end path (network delay, packet loss, delay variation). Parametric models have thus been proposed, aiming at estimating user-perceived voice quality and its dependence on such parameters. In the next sections we describe the E-model, a parametric model for conversational voice quality estimation, standardized by the ITU-T in recommendation G.107 [61] and a methodology to reduce this model to transport layer metrics [23], which can be directly measured.

2.5.2 The E-model

The E-model [61] is a computational model intended to be used as a transmission planning tool. It provides an assessment of the combined effects of various transmission parameters in the mouth-to-ear path on user-perceived conversational voice quality. The E-model takes into account a wide range of telephony-band impairments, in particular the impairment due to low bit-rate coding devices and one-way delay, as well as the "classical" telephony impairments of loss, noise and echo.

The E-model is based on modeling the results from a large number of subjective tests done in the past on a wide range of transmission parameters. The primary output of the E-model calculations is a scalar quality rating value known as the "Rating Factor, R." R ratings can be transformed to estimates of user opinion. Eq. 2.2 [61] shows how R ratings are mapped to MOS values on the ACR scale.

$$MOS = \begin{cases} 1 \quad if \quad R < 0 \\ 4.5 \quad if \quad R > 100 \\ 1 + 0.035 \cdot R + 7 \cdot 10^6 \cdot R \cdot (R - 60) \cdot (100 - R) \\ \qquad\qquad\qquad\qquad if \quad 0 < R < 100 \end{cases} \tag{2.2}$$

For a call of acceptable quality, average MOS should be over 3.60 (R value greater than 70). Table 2.1 shows how the E-model output maps to user satisfaction.

Table 2.1: Relation between R value and user satisfaction [61].

R value (lower limit)	MOS (lower limit)	User satisfaction
90	4.34	Very satisfied
80	4.03	Satisfied
70	3.60	Some users dissatisfied
60	3.10	Many users dissatisfied
50	2.58	Nearly all users dissatisfied

The E-model is based on a mathematical algorithm, with which the individual transmission parameters are transformed into different individual "impairment factors." The basic assumption is that these impairmaint factors are *additive* on the psychological scale:

$$R = R_0 - I_s - I_d - I_e + A. \tag{2.3}$$

R_0 is the basic signal-to-noise ratio at the receiver end, including noise sources such as circuit noise and room noise, and I_s is the combination of impairments that occur simultaneously with the voice signal (e.g., low loudness, non-optimum sidetone, quantization noise). I_d represents delay impairments and I_e is the equipment impairment factor, which includes impairments caused by low bitrate codecs and packet loss. The E-model also includes an "advantage of access" factor A to account for cases where user expectations compensate for impairments: For example, a user may accept some decrease in quality in exchange for mobile connectivity. A discussion on how user context affects the value of A, and, in turn, user-perceived voice quality is available in ITU-T recommendation G.113 [64].

G.107 defines default values for all input parameters to the E-model algorithm and recommends using these values for all parameters which are not varied during planning calculation. Choosing default values for all parameters other than delay and equipment impairment, the rating factor R is reduced to

$$R = 94.2 - I_d - I_e. \tag{2.4}$$

2.5.3 Reducing the E-model to transport layer metrics

Starting from Eq. 2.4, Cole and Rosenbluth [23] have proposed a methodology to reduce the E-model to transport-level metrics. First, the authors assume an all-IP end-to-end path (i.e., there are no circuit-switched components) and use default values for all parameters other than network and codec-induced delay (i.e., talker echo, listener echo, etc.) in the analytic expression for I_d. Then, they use the analytic expression of G.107 to calculate I_d values as a function of mouth-to-ear delay, and fit the resulting curve to a simplified expression.

No analytic expression is available for calculating the equipment impairment factor (I_e), though. Recommendation G.113 [64] lists I_e values as a function of codec type, average packet loss rate, packet loss burstiness and packet size. These values are derived by subjective tests. In a similar spirit as with the delay impairment factors, the authors fit the above empirical data to simple expressions of I_e as a function of packet loss for the G.729a and G.711 codecs.

Combining the expressions of I_e and I_d, the authors derive a fully analytic expression of the R value. For the G.729a codec, when each IP packet carries two 10 ms audio blocks, the algorithmic and packetization delay add up to 25 ms. If we further assume a dejitter buffer which adds 60 ms delay in the playout process and that packet loss is random, we derive the following formula:

$$\begin{aligned} R = {} & 94.2 - 0.024 \cdot (d_{network} + 85) \\ & - 0.11 \cdot (d_{network} - 92.3) \cdot H(d_{network} - 92.3) - 11 \\ & - 40 \cdot \ln[1 + 10 \cdot (e_{network} + (1 - e_{network}) \cdot e_{dejitter})] \end{aligned} \tag{2.5}$$

where:

- $d_{network}$ is the end-to-end network delay in ms
- $e_{network}$ represents the ratio of packets lost in the network path
- $e_{dejitter}$ represents the packet discard ratio at the dejitter buffer due to excessive delay variation
- $H(x) = 1 \quad if \quad x > 0; \quad 0 \quad otherwise$

The quality assessment methodology of Cole and Rosenbluth has been widely used in the literature. Below we cite a few examples. In the context of wireless mesh networks, Kim et al. [73] have used it to explore the interplay between VoIP and TCP flows, while Niculescu et al.[89] have applied it to assess the performance of an IEEE 802.11-based wireless mesh testbed optimized for voice transport, evaluating techniques such as the use of multiple interfaces, path diversity and packet aggregation. The poor performance for TCP traffic as the number of simultaneous VoIP sessions over an IEEE 802.11 WLAN increases is tackled by Verkaik et al. [123], who use this model to evaluate the performance of SoftSpeak, a set of proposed Wi-Fi-compatible software extensions to simultaneously improve VoIP and TCP performance and fairness. Balasubramanian et al. [10] have applied it to assess the performance of *ViFi*, a scheme aiming at improving vehicular Wi-Fi-based connectivity. Markopoulou et al. [81] use a similar E-model-based methodology to assess VoIP over Internet backbones.

We have adopted this methodology for the experimental evaluation of community-based secure VoIP services (Chapter 4), since our purpose is to quantify the effects of various phenomena which are pertinent either to the network topology (i.e., the fact that we focus on a scenario where both call endpoints are attached to Wi-Fi links), network devices in the end-to-end path (resource-constrained residential Wi-Fi routers), or to security services we apply (high-overhead tunneling mechanisms) and correlate them with perceived voice quality; codec settings and other environmental factors were considered fixed or not studied.

2.5.4 VoIP capacity in wireless networks

A significant body of research addresses performance issues of VoIP services over wireless networks, and IEEE 802.11 in particular. Garg and Kappes [43] identified that even for low-bitrate codecs, such as G.729a (8 Kbps), the number of VoIP sessions that a Wi-Fi cell can simultaneously sustain is surprisingly small. This is due to the fact that a VoIP stream is typically composed of small packets and the overhead for their transmission is very high. They propose a simple analytic model to calculate VoIP capacity, which we adapt to our secure peer-to-peer VoIP scheme (see Chapter 4) to estimate an upper bound on the number of VoIP sessions we can sustain.

In the same spirit, Hole and Tobagi [50] propose a similar simple mathematical model of VoIP capacity and use a MOS-driven evaluation methodology. They validate the findings of Garg and Kappes using simulation and expand on various codec settings and scenarios, putting more focus on delay constraints and quality requirements of voice sessions.

With the advent of IEEE 802.11e [54], which added traffic prioritization extensions, efforts have been put in the direction of improving VoIP capacity by appropriately tuning parameters such as the duration of a transmission opportunity (TXOP), minimum contention window value (CW_{min}) and packet buffer size [26, 117].

2.5.5 Security vs. performance

Communications increasingly necessitate security measures, especially when carried out over untrusted wireless networks, as is the case for roaming users visiting foreign Wi-Fi hotspots. Typically, traffic is protected by means of VPN mechanisms, such as IPsec [72] or OpenVPN [92]. However, security mechanisms incur performance overhead, which is particularly important for delay- and loss-sensitive VoIP communication. In our work, we are particularly interested on the effects of security on VoIP performance, especially when both call endpoints are attached to Wi-Fi links and security functionality is implemented on resource-constrained home Wi-Fi equipment; such scenarios have not received significant research attention, to the best of our knowledge.

Voice over IPsec in wireline networks is experimentally studied by Barbieri et al. [11], where a header compression method called *cIPsec* is also proposed and evaluated. IPsec encryption and packetization overhead are studied via analysis and simulation by Xenakis et al. [126], while Miltchev et al. [84] experimentally compare the performance of IPsec and application layer security protocols. A work more closely related to ours is due to Nascimento at al. [88], who present experiments on the effects of IPsec on voice quality when one call endpoint was connected to a Wi-Fi AP and the other connected to a Bluetooth pico-net.

Chapter 3

Wireless Community Networks: A case for user-provided networking

For wireless local area network communications, Wi-Fi emergence was a true revolution. Driven by their low cost and ease of deployment, IEEE 802.11-enabled devices became standard equipment for laptops and handheld devices and appeared as the predominant technology for local wireless connectivity. Operation in unlicensed spectrum facilitated Wi-Fi deployment, since it was straightforward for commercial operators, academic institutions, or even plain radio communications enthusiasts and tech-savvy users to build wireless service architectures on top of it, without the need for acquiring a license.

In modern densely populated urban areas, the coverage that WLANs offer is ever-growing. Performing a scan for wireless network presence reveals so high a number of wireless access points (APs) in the neighborhood that we should be more concerned about interference than coverage.

Wireless Community Networks (WCNs) have been developed as grassroots movements of WLAN enthusiasts, who use inexpensive networking equipment for free interconnection, thus creating all-wireless autonomous networks. Reasons for their emergence can be found in the above discussion. However, their success also depends on local factors (e.g., the degree of penetration of fixed broadband access services in the area). Based on the offered wireless connectivity, such networks aim at providing a variety of services, with free Internet access, community-wide VoIP and file sharing topping the list of the most frequently accessed ones.

The architecture of such networked communities and incentives that keep them operating with robustness are worth studying, and these are the focus of this chapter. Here, we first classify public wireless access schemes (Section 3.1) and report on some of the most significant WCNs worldwide (Section 3.2). Then, we discuss the circumstances under which WCNs emerged and evolved in Sections 3.3 and 3.4, respectively. We present the incentive mechanisms that regulate their operation in Section 3.5 and study their future in Section 3.6. Section 3.7 describes the architecture of community wireless mesh networks. In Section 3.8, based on data from two of the largest wireless community mesh networks worldwide, we discover power-law characteristics in their structure.

3.1 A classification of public wireless access schemes

Here we characterize WCNs based on the initiative behind their emergence and their architecture.

3.1.1 Initiatives

Community initiative

A major focus of this work is on wireless communities which are the result of collective efforts of individual volunteers and function on a not-for-profit basis. Successful communities have emerged, members of which use IEEE 802.11/Wi-Fi technologies to set up a wireless backhaul to connect to one-another, enjoying a variety of broadband services, such as VoIP, online games, FTP or Web access [110, 9, 25, 83, 75]. Sometimes, community members operate public hotspots to offer wireless access to passers-by, attaching them to the community network, or even offering Internet access through community-owned Internet gateways.

In a similar fashion, an individual WLAN owner may open his private hotspot for public access without anticipating monetary compensation. Instead, he is either driven by pure altruism, or expects that his offering will be reciprocated by other community members, when he is roaming near foreign hotspots [30].

Commercial initiative

Following the above trend, commercial players have entered the scene, offering mediation services for the development of wireless communities. FON [37], for instance, has proposed a private hotspot sharing scheme, where WLAN owners can either share their WLANs for a small monetary compensation or in exchange for similar service when they are away from their own WLAN. FON takes care of user registration and authentication and withholds a fraction of the money paid to the hotspot micro-operator for the provided service.

Notably, British Telecom has recently partnered[1] with FON so that hundreds of thousands of BT's subscribers share their home broadband lines over Wi-Fi with other community members.

[1] http://www.bt.com/btfon

Municipal initiative

Municipalities often set up APs in public spaces, offering inexpensive Internet access to citizens. To achieve this, they may get into agreements with private companies, permitting them to deploy their wireless solutions. Authentication with the operator of the network, as well as a fee for the service may be required. This model has been adopted by the municipality of Philadelphia [125], as well as the City of London, which has set up a deal with *The Cloud*, a European Wi-Fi hotspot aggregator.

3.1.2 Design alternatives

Wireless mesh

WCNs sprung from private initiatives often have this structure. Multi-interface nodes set up a wireless mesh. Some also act as gateways to the public Internet and can be reached over the wireless backbone. Section 3.7 offers a more in-depth discussion of this architecture (Figure 3.1a).

Hotspot-based architecture

Hotspot-based WCNs typically target nomadic users who use wireless hotspots to access the Internet (Figure 3.1b). Here, deploying a wireless backhaul is not the norm. Municipality-initiated WCNs usually have this structure. Sometimes, hotspot-based WCNs are built relying on the private contributions of individual WLAN owners, who share their fixed broadband lines over Wi-Fi. We have witnessed both commercially initiated [37] and not-for-profit [30] such attempts.

3.2 Wireless communities around the world

Here we report on some of the most significant WCNs worldwide and categorize them as described in Section 3.1. Our findings are summarized in Table 3.1. The reported network dimensions are based on data found on publicly available node maps and measurement studies [83].

We have chosen representative WCNs and tried to include the most well-known and influential ones. SeattleWireless [110], for instance, has been at the forefront of the WCN movement since the early 2000s.

Also, NYCwireless [91] and the CUWiN Foundation [25] are active in advocating the use of open wireless technologies, fostering the growth of WCNs and developing software for community wireless projects. They are operated by non-profit organizations, whose interest has recently been drawn to developing free wireless access solutions for under-served communities, both in the US and in other countries.

The Athens Wireless Metropolitan Network (AWMN) [9] and guifi.net [45] are two of the largest community mesh networks in the world. Currently, AWMN has

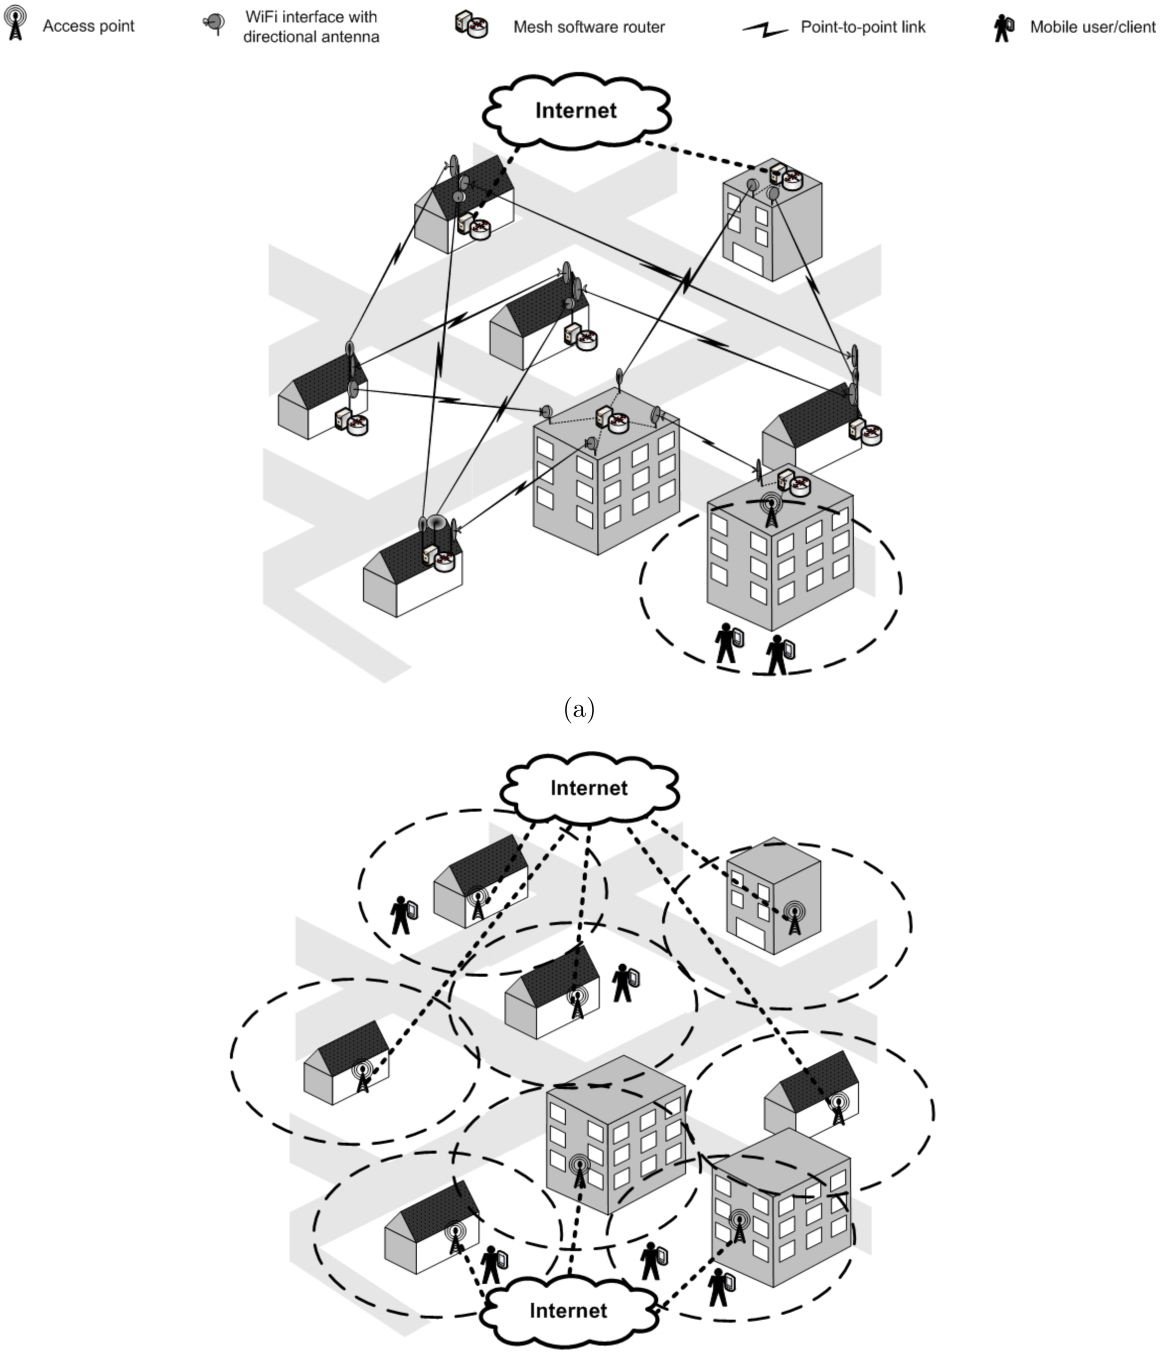

(a)

(b)

Figure 3.1: Architectural alternatives for Wireless Community Networks. A wireless mesh architecture is shown in (a), while the hotspot-based alternative is shown in (b).

more than 9000 registered nodes, with more that 2400 of them being active. guifi.net, based in Barcelona, Spain, counts more than 14000 active nodes as of September 2011.

From data available from the project's website, there are registered nodes in many Spanish cities, but also in other countries across Europe and Africa. Due to physical constraints, the network is composed of isolated zones. Owing to their size, we have selected AWMN and guifi.net as the basis of our work on modeling community wireless mesh networks based on empirical data (Section 3.8).

Wireless Leiden [75] is a similar effort in the Netherlands. Its aim is to provide a free citywide all-wireless network in the city of Leiden and offer free broadband Internet access to nearby villages, where no fast Internet alternatives exist.

Recently, the Freifunk community has gained much popularity. Freifunk mesh networks have sprung in various German cities, as well as in cities in Austria and Switzerland. In Berlin, Freifunk counted 316 concurrent participating nodes on average, according to a 2007 study [83].

The MIT Roofnet [13] mesh network started as a research project focusing on wireless multihop routing and IEEE 802.11 protocol performance, while offering Internet access to nearby residents. It is now less vibrant (approximately 20 active nodes), but its technology is used from other wireless community projects.

Table 3.1: Wireless communities around the world (as of September 2009)

Network	Location	Size	Type[a]	Initiative
SeattleWireless (2000)	Seattle, WA, USA	~80 nodes	M	Community
AWMN (2002)	Athens, Greece	2473 nodes	M	Community
CUWiN (2002)	Urbana, IL, USA	48 nodes	M	Community
Freifunk Berlin (2002)	Berlin, Germany	316 nodes[b]	M	Community
Wireless Leiden (2002)	Leiden, Netherlands	73 nodes	M	Community
guifi.net (2004)	Barcelona, Spain	~14000 nodes	M	Community
NetEquality Roofnet (2007)	Portland, OR, USA	126 nodes	M	Community
NYCwireless (2001)	New York City, USA	145 nodes	H	Community
Wireless Philadelphia (2007)	Philadelphia, PA, USA	15 $miles^2$	H	Municipal
FON (2006)	Worldwide	~700,000 hotspots [c]	H	Commercial

[a] M: mesh, H: hotspot-based

[b] Circa 2007 [83]

[c] Based on information available from `http://www.fon.com`, as of September 2009.

3.3 The birth of wireless communities

The appearance of WCNs dates back to the late 1990s - early 2000s, when IEEE 802.11 was introduced. Radio technology enthusiasts were the first to embrace the new technology and experiment with it for long distance interconnection. Thus, communities of interconnected peers emerged in a period when fixed wireline infrastructure could not always support them. One can consider the first wireless community networks as the evolution of amateur radio.

As wireless technology matured and gained popularity and as more advanced WLAN standards emerged, WCNs started to grow. However, one of the major factors that contributed to their growth was the low penetration of broadband access technologies in some countries.

In addition, some wireless communities [13] have been developed as experimental testbeds for wireless research. In these cases, free wireless Internet connectivity is a side effect, while the main goal of such WCNs is to test novel wireless technologies and evaluate current WLAN standards.

3.4 Growth

One of the reasons AWMN has grown to be one of the largest such networks worldwide was the fact that, in Greece, DSL penetration used to be one of the lowest in the European Union before 2005. Actually, back in 2002, when the first AWMN node was set up, there was a single operator offering DSL to corporate clients, with download speeds of at most 128Kbps and a very high price. Setting up long-distance wireless links using IEEE 802.11 to achieve cheap citywide broadband connectivity seemed to be the only option for tech-savvy users. At the time, the most successful applications were fast file sharing and online gaming, as well as Internet connectivity through AWMN-to-Internet gateways (when some participants shared their fixed broadband Internet lines with the community).

The network evolved and gained publicity through word-of-mouth and dissemination activities of its members, which included demonstrations at technology expos and universities and participation in festivals and other public venues, where free wireless connectivity was offered through the AWMN infrastructure.

A key technological shift was the adoption of the new IEEE 802.11a variant for setting up point-to-point backbone links. Moving from 2.4GHz to the less congested 5GHz band, interference was limited. Nodes with multiple collocated interfaces had more non-overlapping channels to choose from for each of their radio interfaces, while links of higher throughput could be achieved. Backbone nodes could sustain more links, which offered path redundancy and resilience. Citywide VoIP and video conferencing within the community, as well as file sharing have thrived since. Increased capacity, better quality of service and the fact that backbone nodes could now sustain more links caused a population boom in the WCN. A direct consequence was the expansion of its coverage to most of the suburbs of Athens.

3.5 Incentives for sustainable wireless communities

3.5.1 Incentives for participation

Why would one participate in a WCN? Usually it is the enthusiasm stemming from experimentation with new technologies or the joy of creation that building one's own equipment and configuring one's node offers.

Also, the services provided by the community are tempting, especially when the respective services over commercial (fixed) broadband infrastructure are of worse quality, or even non existent.

Those who believe that broadband connectivity should be unleashed and the barriers imposed from the oligopoly of ISPs be removed and that broadband access should be a public good usually are the first to join WCNs. Under this perspective, WCNs can be considered a modern technological "movement". Operation in unlicensed spectrum facilitates their growth and, coupled with the extensive use of open source software, shapes their free and self-organizing character. It is atypical of their participants to think of community-initiated WCNs as the vehicle to make profit or develop a new business.

3.5.2 Conformance and contribution

As far as WCNs like AWMN are concerned, contribution to the community and abiding with the explicit or implicit participation rules is enforced in a tit-for-tat manner. Users that do not conform to such rules can be effectively excluded from the community.

Most WCNs are open for participation, but have some structure and their deployment is not completely anarchic. In order for the system to operate without disruptions, nodes should follow some rules, especially concerning addressing and routing. Also, they should not behave selfishly, since that could potentially cause congestion to the backbone links, as well as saturation of some services. In case a user does not behave, there is high probability that he will eventually be excluded from the network. At a technical level, non-conformant behavior can be easily detected from one's neighbors and reported to the community.

Exclusion can also easily be implemented; neighboring users can simply shut down the links they share with the user that causes anomalies to the WCN's operation. Thus, the latter will be disconnected and isolated from the community. In fact, we have witnessed similar incidents within AWMN.

As described in Section 3.7, there are typically two types of nodes, backbone ones (with multiple point-to-point links) and APs, where clients with a single link attach to. Incentives for the maintenance of the first two links of a backbone node are quite straightforward WCN participation. Additional links add redundancy, network bandwidth, and reputation in the community. There is expressed mutual interest for the two link endpoints to set one up. As soon as one of the two nodes defects, the link becomes inoperative.

The case is different for client-to-AP links. There is no direct and measurable gain for an AP to offer connectivity to clients. However, in the long run, there is a potential gain for the community as a whole, since clients are anticipated to bring more content and are expected to gradually upgrade their nodes to backbone ones. Thus, serving clients is considered a means of attracting attention to the community and gradually recruiting more resources.

3.5.3 Building reputation

Interestingly, senior WCN members and members who own *powerful*[2] nodes are highly esteemed in the community. Because their nodes are important for the operation of the network, their decisions may affect many users.

Building up reputation within a WCN can easily be explained. In such communities, participation anonymity is not easy (or even possible) to achieve, nor is it desirable. Recruiting new members is an incremental process, where older members introduce new ones to the community, offering them connectivity.

The infrastructure-based nature of the network is such that long-term relationships among members are built; nodes are fixed and links are typically permanent. Most of the times, owners of nodes that share a link know each other personally, while the community encourages socialization among its members. Contribution to the community is directly attributed to its initiators, while unacceptable behavior is detected, reported and, potentially, punished, as described in Section 3.5.2.

3.6 A look to the future of WCNs

To gain insight on how WCNs will evolve, one has to reconsider the factors that led to their emergence and booming growth.

First, low-cost broadband services, once one of the highlights of WCNs, have now become a commodity in most modern metropolitan areas. No more is joining a WCN the only way to enjoy high-quality multimedia services or fast file sharing. On the positive side, though, promising new wireless technologies, such as the IEEE 802.11n

[2] Here we refer to backbone nodes with many links, hosting multiple services or operating an AP. They are important for the network, since much of its traffic flows through them and many other nodes rely on them for connectivity.

standard, are about to be integrated. These may revitalize interest in WCNs and cause a new wave of developments in the area.

On the other hand, an issue that is generally still unresolved is how such broadband services will be offered for mobile users at low cost. 3G-based solutions are still considered expensive by many users and their adoption is relatively low. Basing our arguments on the increased wireless coverage in modern metropolitan areas, we believe that open wireless access schemes can become a viable alternative for the provision of nomadic broadband network access. To support this action, traditional WCNs may need to become less exclusive so that they are accessible not only strictly to their participants, but also to roaming non-members.

To this end, opportunistic WLAN access based on the private contributions of individual WLAN owners will likely attract more attention as a vehicle to make inexpensive ubiquitous broadband Internet a reality. WLAN owners can become micro-operators that trade bandwidth for payment in kind or even money. In Section 4.1 we present an architecture towards this vision.

3.7 Architecture of a community wireless mesh

3.7.1 Node types

There are two types of nodes. *Backbone nodes* are those that the backhaul of the network is built upon. They typically have more than two network interfaces and run routing software.

Some backbone nodes operate omni-directional antennas and function also as *APs*. *Client nodes* can then attach to these APs and use the network's infrastructure. *Clients* can be considered the “leaves” of the network.

Typically, WCN participants discouraged new users from operating client nodes, since clients usually do not contribute to the network's operation and coverage. In practice, being a client is usually the first step a user takes when joining such a community.

3.7.2 Links

The backbone of the network is based on directional point-to-point links. Although IEEE 802.11 was designed as a broadcast protocol for local communication, it is widely used for setting up long-distance links, assisted by directional antennas. In the beginning, 802.11b was the protocol of choice, but with the advent of 802.11a, it was replaced in the backbone in order to minimize contention and interference from clients and APs. IEEE 802.11b is still used by APs.

Our recent study [33] of the evolution of AWMN revealed that many backbone nodes maintain up to 6 backbone links. The more links per node, the higher path redundancy is achieved.

3.7.3 Addressing and routing

Mesh-based WCNs have the same functionality, underlying mechanisms and applications as the public Internet. They are based on an IP layer and follow a private addressing scheme. In most of them (e.g., AWMN), a central authority tackles addressing issues. Each newly-registered node is assigned a private IP address range. Every community follows its own routing scheme, often involving BGP, much like the Internet itself. From the BGP viewpoint, each WCN node can be considered a single *autonomous system*. In some cases, experimentation with routing protocols suited for wireless ad hoc networks, such as OLSR, is carried out.

3.7.4 Underlying technologies

Open source software and often hand-crafted hardware are used for interconnection. Linux-powered embedded rooftop equipment is used for controlling the operation of each node, while much experimentation involves the design and construction of low-cost custom-made antennas. Figure 3.2 depicts part of the infrastructure of the AWMN node that we have set up on the rooftop of one of AUEB's buildings.

3.7.5 Services and applications

The most prevalent services are file sharing and community-wide VoIP. Also, on some occasions, members share their fixed broadband connections with the community, so that Internet access is achieved through WCN-to-Internet proxies. For other services see [33].

3.8 Modeling the structure of wireless community networks

In this section we study the topology of community wireless mesh networks. Our intuition was that the way wireless communities are structured is not "random" and our observation about the topology some large WCNs supported this claim: It appears that there are a few "powerful" nodes with many wireless links, while there are lots of users with very few links. A question that naturally emerges is how to derive realistic models for WCN structure and growth. Answering this question would be particularly significant for relevant research. For example, it would facilitate the generation of realistic topologies to be used in the performance evaluation of services and protocols over community wireless mesh networks. We discuss issues pertaining to WCN models in Section 6.1.

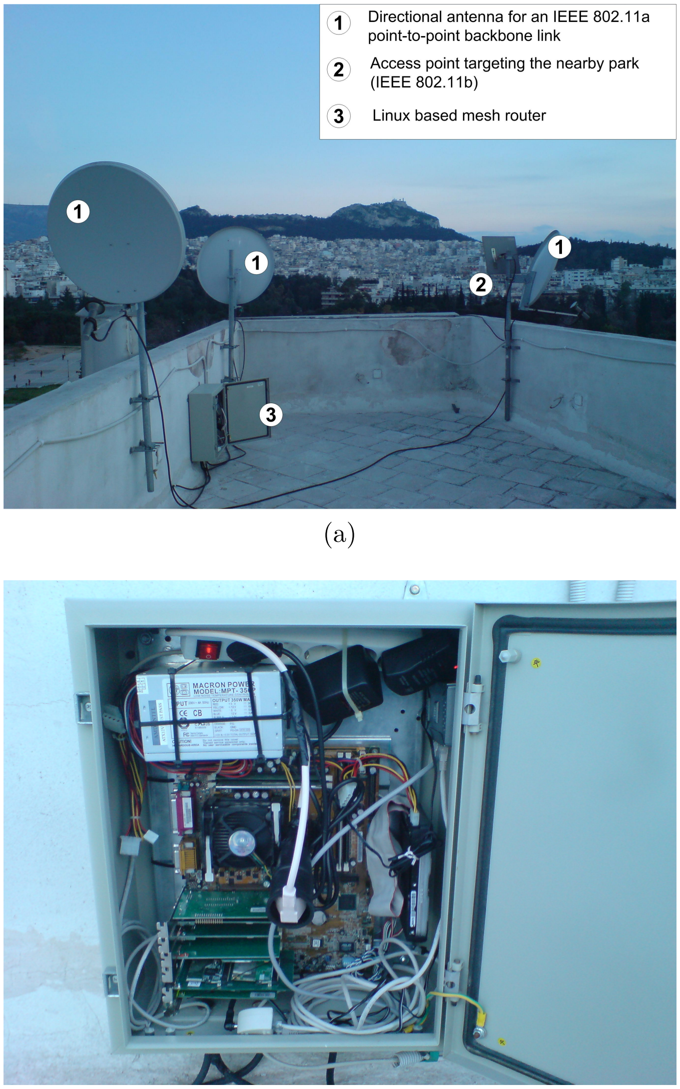
1 Directional antenna for an IEEE 802.11a point-to-point backbone link
2 Access point targeting the nearby park (IEEE 802.11b)
3 Linux based mesh router
1
1
1
2
3
MACRON POWER
MODEL MPT-350

(a)

(b)

Figure 3.2: Pictures from MMlab's AWMN node. Three of our backbone links, as well as our AP are visible in (a). IEEE 802.11 wireless interfaces are attached to hand-made antennas. In (b), the Linux-powered rooftop PC which hosts our wireless interfaces and performs routing is shown.

3.8.1 The emergence of power laws

Based on empirical data, we observed that the distribution of some structural properties of WCNs has a long tail, which leads us to the intuition that power laws could be good fits to describe them concisely. Similar works have been carried out in many other contexts to describe various observable phenomena. In the area of communications and networks, Faloutsos et al. [35] have long shown the emergence of power laws in the AS-level Internet topology. In a similar spirit, we present power laws that can describe the structure of mesh-based WCNs.

A random variable x is power-law distributed if its probability density function is of the form

$$f(x) \propto L(x)x^{-a}, \tag{3.1}$$

where $L(x)$ is a slowly varying function, i.e., a function for which $\lim_{x\to\inf}\frac{L(cx)}{L(x)} = 1$. Here we assume cases where $L(x)$ is the constant function. The exponent a is known as the scaling parameter, where $a > 1$. If we plot $f(x)$ in a log-log scale it will appear linear and a will be the slope of the curve.

If we represent a wireless community mesh network as an undirected graph, where vertices represent nodes and edges represent wireless links between the respective nodes, based on empirical data, we observed that such graphs do not look like random ones; in contrast, few *hub* vertices have a high degree (many links) while there is a long number of nodes with small degrees. These low degree vertices could, for instance, represent backbone nodes with only two interfaces or client nodes attached to an AP. Therefore, the degree frequency distribution is a strong candidate to exhibit power-law behavior. We denote degree frequency, i.e., the number of nodes with degree x as $f(x)$.

This section does not provide an in-depth study of the structural properties of WCNs, which we defer for future work. Instead, it serves as a first step in formally modeling some of their properties, to be used in the generation of realistic WCN graphs. Other features are also of significance in this process, such as node interconnection properties (e.g., hop counts), community neighborhood structure and graph assortativity (i.e., a value indicating the trend of high-degree nodes to connect to other high-degree nodes).

3.8.2 Methodology and results

We base our observations on empirical data from existing mesh WCNs. Unfortunately, even though we discover power laws in these WCNs, it is hard to generalize due to the lack of large communities and relevant data. We focus on two of the largest such networks (to the best of our knowledge), AWMN (Athens, Greece) and guifi (Barcelona, Spain). There are publicly available data about their nodes and links, maintained in web-accessible databases and maps. Node and link information is available in XML and HTML form, which is easy to download and parse. It should be noted that there is an amount of stale information in these datasets, since some nodes and links may not be functional at the time of retrieving data. It is hard to obtain accurate real-time data, especially without access to nodes in the networks (which is the case for guifi) or without the availability of a real-time network monitoring tool. However, this unavoidable lack of accuracy does not invalidate the basic purpose of this work and it is not within the scope of this work to address this issue.

We have preprocessed data to remove registered nodes with zero degree (isolated nodes without any links). Concerning degree frequencies, we have borrowed from the methodology of Faloutsos et al. [35], where few nodes with very high out degree are removed from the dataset. In particular, they plot degrees starting from degree 1 until they reach a degree which has frequency of 1. In both networks we study, though, the removed nodes represent a relatively small fraction of the overall node population: For AWMN, the removed nodes represent 0.36% of the total number of nodes, while for guifi this percentage rises to 0.5% (in both cases, excluding isolated nodes).

It should be noted that, in practice, and by the very nature of the process of fitting a power law to empirical data, one can rarely (if ever) be positive that such data are drawn from a power law distribution. A better fit is possible among the infinite set of distributions. Also, we have chosen to apply linear regression to derive a fit. More robust methods are possible [22].

In Figure 3.3, we present log-log plots of the frequency of degrees for the two WCNs we study. The data points represent empirical data and lines represent power law fits obtained with linear regression. In the case of AWMN, the scale parameter $a = 1.98$ is higher than that of guifi (1.687). It should also be noted that there is a larger percentage of single-degree nodes in the guifi network. The practical interpretation of this phenomenon has to do with the intended use of the network; the AWMN community seems to favor backbone peer-to-peer links, while in the guifi network, there is an increased tendency to operate powerful APs for clients to join.

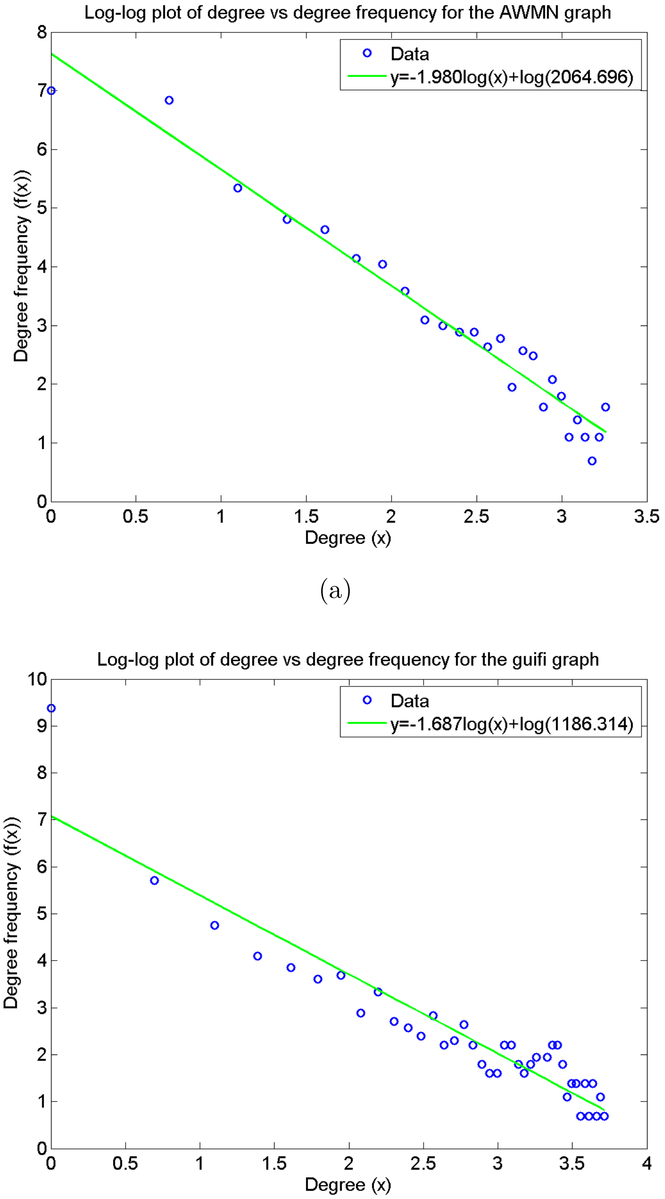

(a)

(b)

Figure 3.3: The emergence of a power law in the distribution of degree frequencies for the graphs representing two WCNs. Power law curves were derived using linear regression. R^2 values are 0.8825 and 0.9616 for guifi and the AWMN, respectively.

Chapter 4

User-centric secure multimedia services

In this Chapter we study hotspot-based wireless community networks, which are built by the private contributions of individual WLAN owners. Our specific focus is on building a service architecture to support peer-to-peer multimedia communications, with an emphasis on security and privacy.

This work is positioned in the context of the *Peer-to-Peer Wireless Network Confederation (P2PWNC)* scheme [30, 31], on which we rely as the wireless access infrastructure. P2PWNC, a peer-to-peer Wi-Fi sharing scheme, was proposed by Efstathiou in his Doctoral dissertation [29].

The question we are called to answer is whether in densely populated metropolitan areas, where wireless coverage is adequate, P2PWNC-based wireless communities can offer a secure, low-cost, user-centric alternative to traditional GSM/3G services. We attempt to answer this question performance-wise, by building and experimenting with secure voice and data services, but also with respect to the basic principles of user-centric networking that we manifested in Chapter 1, i.e., (i) decentralized operation, (ii) low cost, (iii) open access, (iv) security and privacy, and (v) the assumption that users are rational. We argue that our design respects these principles in Section 4.3. Whether or not such a scheme is successful, however, critically depends on the adoption of peer-to-peer Wi-Fi sharing and the coverage such communities enjoy, and is related with economical, societal, and technical factors outside the scope of this work.

To rival traditional cellular services, performance of voice and data communication over user-centric wireless networks is a key issue, in part due to the unpredictable nature of wireless communications, where delay sensitive applications like Internet telephony are known to suffer: Poor signal conditions, but also contention for access to the medium and interference brought by spectrum scarcity, dense and anarchic Wi-Fi deployment, and poorly-configured wireless equipment, account for that. At the same time, important overhead is imposed by the need to secure communication. Our scheme makes use of VPN tunneling, but also involves CPU-intensive cryptographic operations associated with the P2PWNC protocol. These performance penalties are intensified by the limitations imposed by low-cost, off-the-shelf, embedded WLAN devices and mobile terminals.

We adopt a QoE-oriented stance. Most important to us is to estimate the maximum number of simultaneous VoIP calls of acceptable quality–as a user would perceive it–that a typical P2PWNC-enabled WLAN AP can sustain, by measuring how the use of VPNs to secure communications, but also cryptographic operations associated with the P2PWNC protocol affect voice quality.

This Chapter is structured as follows: Section 4.1 provides an overview of P2PWNC and Section 4.2 deals with the design of service architectures on top of it. We show that our design adheres to the principles of user-centrism in Section 4.3. We then focus on the performance evaluation of these services; we use a simple analytic model to estimate upper bounds on VoIP capacity in our tunneling-based architecture in Section 4.4.1, describe our experimental methodology and testbed in Section 4.4.2 and present performance results in Section 4.4.3.

4.1 A peer-to-peer approach to WLAN sharing

In this section we present P2PWNC, an architecture for user-provided wireless networking based on the reciprocity-based exchange of Internet bandwidth over Wi-Fi. We review some fundamental challenges and present its system model, user identification and accounting scheme, internal mechanisms and protocols.

4.1.1 Fundamental challenges

Sharing one's Internet connection with roaming individuals incurs direct and indirect costs, or may be forbidden by the ISP. Users may have to put up with the managerial overhead of configuring their WLAN for shared access. Also, admitting more users to one's private network can result in reduced service level for its owner. Sometimes, WLANs are attached to metered broadband lines. Sharing such a WLAN means that its operator pays for every single byte uploaded or downloaded by visitors. Security concerns also make individuals reluctant to open their WLANs for public access. The challenge for the system designer is to minimize the incurred costs, while revealing the benefits for users to join and contribute their resources.

Joining the system should be possible with minimal installation cost, both financially and effort-wise. User registration should be fast and decentralized, to facilitate the system's organic growth and scaling. From the nomadic user's perspective, wireless access should have low cost compared to its cellular alternatives. Also, roaming privacy is a desirable feature; a user may wish to access foreign APs without disclosing personal information to the service provider or his network point of attachment to his communicating endpoints.

Note that our focus is on low-mobility, *nomadic* users and our emphasis is on access issues rather that true mobility support. Handing-off efficiently between foreign visited WLANs is a significant challenge.

4.1.2 Peer-to-peer hotspot sharing

The core concept in P2PWNC is that wireless Internet bandwidth is exchanged in a reciprocal manner; one shares his Internet connection with anonymous passers-by over Wi-Fi with the anticipation that he will enjoy the same benefit from another peer when mobile. That is, the problem of creating a hotspot-based wireless community is approached based on the peer-to-peer paradigm. Private WLAN owners have an incentive to contribute Internet bandwidth, given that they value much the mobile network access that they will enjoy as good contributors.

To lower the entry barrier to the system, no registration with central authorities is required, nor any strong user identification scheme. Participants are identified by self-issued, uncertified public-secret key pairs. To join P2PWNC, users simply configure their APs for open access and install the necessary software.

Accounting is based on digital proofs of service (*receipts*) that mobile users provide to visited APs. Receipts are stored in repositories distributed across the system; each peer maintains his own repository, which represents his (partial) view of the system's history of service provisions.

Receipt repositories (RR) are the input to the *reciprocity algorithm*, which identifies good contributors and detects *free riders*, i.e., those who consume resources without contributing to the community. Each time a mobile user requests service, the reciprocity algorithm is executed by the visited peer to decide whether the visitor is a good contributor and deserves to be reciprocated.

Since identifiers are free to generate and our design does not exert control on how they are used, nor does it bind them with physical entities and devices, a single identifier may be in use at the same time on multiple devices accessing different P2PWNC hotspots. In other words, our protocol and mechanism design allows an account to be shared by more than one users. This could be the case for a group of users who trust each other, such as family members. See [31] for a discussion on the advantages of grouping users into teams.

P2PWNC is decentralized and no registration is mandated. Other similar schemes, such as FON [37], do not have these properties. Also P2PWNC has some inherent privacy enhancements. The non-persistence of user identifiers, as well as the fact that the disclosure of real-world user identities is not needed, assist in WLAN roaming anonymity.

4.1.3 Mechanisms and protocols

Here we present an overview of the P2PWNC operations, also depicted in Figure 4.1.

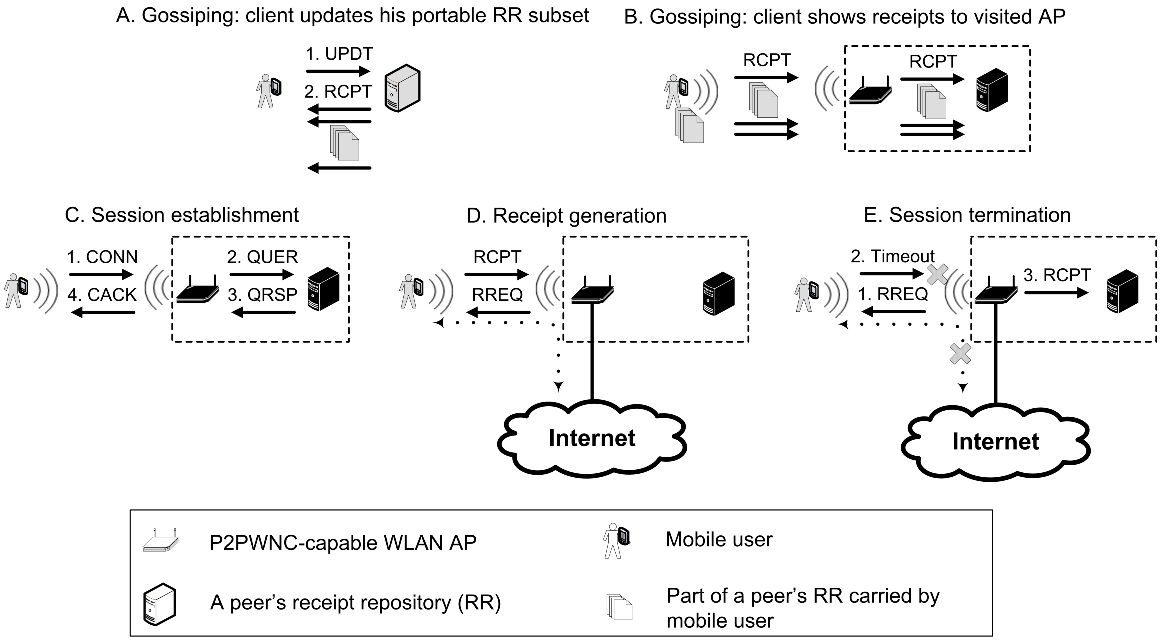

Figure 4.1: P2PWNC operations. First, a user communicates with his home RR and updates the RR subset that he carries with him (A). Then, before requesting service from a visited AP, he presents it with the receipts he carries with him (B). Following, he issues a connection request, upon which the AP consults his own RR, where the reciprocity algorithm is invoked, and decides to accept the user or not (C). In the first case, Internet access is granted and a receipt request-response cycle begins (D). When a client fails to deliver a receipt on time, the session terminates and the AP forwards the last receipt to his RR (E). The messages exchanged during the above processes are summarized in Table 4.2.

Receipts, receipt repositories and the reciprocity algorithm

The P2PWNC accounting unit is a *receipt*, a document digitally signed by a service consumer (mobile user) using his private key. It contains the provider and consumer identities, a timestamp, and information on the amount of traffic forwarded by an AP on behalf of the signer during a P2PWNC session. Receipts represent “debt” among peers, which is assumed to be transitive; if peer A has provided service to peer B and the latter has served peer C, then C indirectly “owes” some service to A. A's contribution can be reciprocated when he visits one of B's or C's APs. With each receipt encoding an “I-owe-you” relationship, a RR can be viewed as a logical directed graph, whose vertices correspond to peers and its edges represent receipts and point from a service consumer to a service provider.

Such a graph is the input to the reciprocity algorithm, which is invoked each time a user requests to be served. Its output is a value expressing a user’s “reputation” in the eyes of a prospective service provider. Given that accounting is decentralized and user identities are self-issued, the reciprocity algorithm should be intelligent enough to detect free riders and prevent from attacks to the accounting mechanism [29].

Session initiation and receipt generation

Receipts are generated on a per-session basis. A prospective service consumer requests access to a P2PWNC-controlled Wi-Fi AP by presenting his identity. Following, the AP invokes the reciprocity algorithm and the RR responds whether the client should be admitted to the visited WLAN. If so, a receipt request-response cycle begins; the AP periodically requests an acknowledgment for the volume of Internet traffic forwarded on behalf of the visitor thus far and the latter responds with a cumulative “fresh” receipt. After a receipt request has timed out (e.g., when the visitor leaves) or upon receiving a malformed receipt (that is, one that the AP is incapable of cryptographically verifying), the session terminates and the AP forwards the last receipt to his RR. This last receipt contains information about the session’s aggregate amount of traffic. Notice that the receipt generation protocol ensures that a visitor’s session has not been *hijacked* by an unauthorized party. In order to sign a receipt and maintain the session, the hijacker would need access to the visitor’s private key.

Receipt dissemination

A visited peer uses only his own view of the system’s history of transactions as input to the reciprocity algorithm. To assist in giving potential service providers a better view of their overall contribution and have better chances of getting access, visitors can also supply parts of their own RRs via a *gossiping protocol*. *Gossiping* takes place at the beginning of a session; the visitor presents the AP with a subset of his own RR carried in his mobile device. These receipts are then merged with the visited peer’s RR and reveal service directly or indirectly owed to the visitor. Note that because the receipts are cryptographically signed, invalid receipts do not help the presenting party. A peer regularly queries his RR for updates to his portable RR subset, since RR contents change frequently.

Implementation issues

P2PWNC entities communicate using a simple ASCII-based protocol running over TCP/IP. For public key encryption, the RSA and ECDSA algorithms are supported for various key sizes. Note that Elliptic Curve Cryptography offers the advantage of smaller key sizes for the same security level. Table 4.1 compares key sizes for the two cryptosystems for comparable security [114].

Table 4.1: Key size comparison between RSA and Elliptic Curve Cryptosystems for the same security level

Security level	**Ratio (RSA/ECC)**
1024/160	6.4 : 1
1536/192	8 : 1
2048/224	9.14 : 1
3072/256	12 : 1

Protocols and the reciprocity algorithm have been designed and implemented with resource-constrained devices in mind, such as off-the-shelf wireless routers and PDAs. P2PWNC typically runs in the firmware of such devices. The protocol messages exchanged among the various entities are summarized in Table 4.2. A detailed description on the protocol, its implementation and its experimental performance evaluation is available [41].

Table 4.2: P2PWNC protocol messages

Message	**Description**	**Direction**
CONN	P2PWNC session initiation request	Client → AP
CACK	P2PWNC session initiation response	AP → Client
RREQ	Receipt request	AP → Client
RCPT	Receipt	Client → AP AP → RR RR → Client
QUER	Query the RR (invoke the reciprocity algorithm)	AP → RR
QRSP	Query response	RR → AP
UPDT	Client (portable) RR update request	Client → RR

4.2 Service architecture

In this section we present our design of user-centric services on top of a community-based wireless network infrastructure. We expect that popular applications would be

VoIP and mobile multimedia and attempt to provide a complementary service to GSM/3G in citywide areas, where Wi-Fi and cellular nowadays have comparable coverage. We envisage an environment where users would be able to enjoy free Wi-Fi roaming around the city and place VoIP calls to other mobile users connected to P2PWNC hotspots. Apart from its low cost, this alternative has inherent privacy enhancements. Given that P2PWNC does not entail central registration to a service provider and that users are free to switch identities at will, our scheme enhances user anonymity.

4.2.1 A tunneling-based scheme

Between an access provider (visited AP) and a consumer (roaming peer), no trust is assumed and no form of cooperation and interaction is expected, other than that dictated by the P2PWNC protocol. Therefore, there is the risk for a visitor that the traffic forwarded by the AP on his behalf is intercepted. On the other hand, the visitor may engage in malicious acts, masking behind the provider's home network. A solution to this issue comes with the use of tunneling. A visited AP only forwards a user's traffic from/to a specific IP address (or a restricted number thereof). This address can be negotiated with the AP after a P2PWNC session has successfully been established and serves as the visitor's trusted gateway.

After negotiation, the user sets up a Virtual Private Network to this gateway and securely relays all his Internet traffic through it. This approach mandates that the user operates a trusted VPN gateway. Considering the overall P2PWNC architecture, and to offer a user-centric feel, we have proposed that this VPN functionality be built within the firmware of the user's home router, which, at the same time, operates the P2PWNC protocol to offer access to other community members. Obviously, for improved performance, and if the user can afford it, he can operate the VPN gateway on separate equipment, and even outside his home network.

With this approach, confidentiality of the user's traffic is ensured, while the service provider cannot be held liable for illegal activities carried out by the visitor, since potential attacks, download or distribution of illegal content, or other malicious behavior will appear to have been performed from the visitors's home network (or the network where his VPN gateway is located).

4.2.2 Rendezvous and call setup

Given that visitors have set up secure tunnels to their home networks, we now show how a voice (or multimedia) call can be set up between two users accessing the Internet via foreign P2PWNC hotspots. In order to spare users the need for extra equipment acting as a VPN gateway, we have built this functionality in the AP's firmware. It should be noted that this AP may serve other P2PWNC visitors at the same time. Note that, although this description is P2PWNC-specific, our approach for

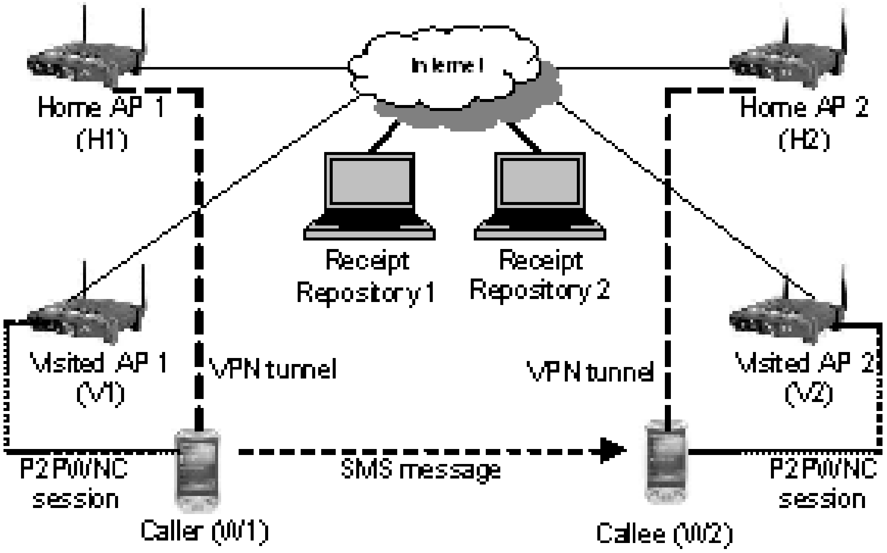

Figure 4.2: A P2PWNC-based secure multimedia call

establishing multimedia communication could be applied to other underlying access schemes. Our basic goal is to demonstrate that a secure multimedia session can be set up in a user-centric manner, with minimal dependence on centralized infrastructures. Figure 4.2 shows the proposed scheme.

A major challenge one has to tackle is for the caller to discover where to initiate the call to. The call endpoints are reachable via their home VPN gateways. Thus, the caller needs to discover the IP address of the callee's home. Public home IP addresses are typically assigned by the users' ISPs and, more often than not, they are dynamically allocated from the ISP's DHCP pool. Here we present a solution for peer discovery based on the exchange of GSM SMS text messages, based on the following assumptions:

- At any time, a user is aware of the IP address of his home VPN gateway, which he can communicate to its peer. This is necessary, not only for user discovery, but also to set up a tunnel in the first place.
- A user who wishes to initiate a call to the other end knows his peer's GSM mobile phone number. This is a reasonable assumption to make, given that either the two users know each other from prior contact and have exchanged such information, or the caller is aware of the callee's phone number via an external channel. In any case, this model follows the traditional GSM paradigm. Also, it is reasonable to assume that cellular phone owners have active subscriptions with GSM operators, and can thus be reachable over GSM.

This is a point where we relax our decentralization assumption: We rely on the centralized GSM infrastructure, but only for reasons of service discovery, and exploiting an external communication channel that is ubiquitous and already in use by peers.

Peer discovery works as follows: Suppose that users W1 and W2 wish to establish a voice call. W1 and W2 are assumed to have established P2PWNC sessions with APs V1 and V2 respectively (V1 and V2 belong to two other P2PWNC peers) and tunnel all their Internet traffic to their home gateways, H1 and H2 respectively. In order to initiate the call, W1 sends an SMS to W2, informing him of his home gateway's (H1) IP address. The SMS can also convey other session parameters, such as the port on which his gateway can receive a media stream, codec parameters, and security-related information. (See Section 6.2.4 for a discussion on achieving end-to-end security.) Then, W2 responds with the voice stream, which is first tunneled to H2, then routed to H1, and, finally, tunneled to W1.

It should be noted that the voice application at the W1 side should be prepared to receive an incoming stream at the agreed upon port. Also, the appropriate state should have been set up at both home gateways so that incoming packets with the peer's home gateway's IP as the source are forwarded to the appropriate tunnel endpoint. (A home VPN gateway may be managing multiple tunnels to roaming stations at the same time.)

Various alternative technologies exist for implementing tunnels. In our measurement-based evaluation we have used OpenVPN [92], a popular and easy to deploy SSL/TLS-based solution. In our earlier experiments, though, we had applied an L2TP/IPsec-based scheme [100], which is more complex (protocol-wise) and with higher per-packet space overhead. VPN tunneling imposes an important data and processing overhead, the effects of which on voice quality are studied in Section 4.4.

As a final note, there are various alternatives to accomplish rendezvous between the two call endpoints, which we discuss in Section 6.2.2. In the IP telephony world, call initiation is typically achieved by means of signaling protocols like SIP [105] or H.323 [65]. However, these involve server components for registration, rendezvous, proxying, etc., and would probably incur additional managerial and performance overhead, especially for non-expert users and home networking equipment.

4.3 Adherence to the principles of user-centrism

We now summarize our arguments that our design is in line with the principles of user-centric wireless networking, as stated in Section 1.2.2.

The user at the center All through our design, users have a central role. They enjoy the role of network access providers, autonomously manage their identities, and operate decentralized security and multimedia services. We also evaluate the performance of these services (see Section 4.4) attempting to measure user-perceived quality.

Open access Anyone can practically join the community by offering a share to his resources. Users also do not have to register with any service; they can simply generate their own, self-certified identity, which they can dispose at any time. Services are built on a voluntary basis. If a two peers need to communicate securely, it is up to them to operate their own VPN gateway at their premises.

Decentralization and distribution of tasks Both at the access and the service provision layer, operation is decentralized. P2PWNC does not entail a centralized user registration and accounting scheme. Instead of relying on a small set of operators to deploy the network infrastructure, we build on the private contributions of (most often residential) WLAN owners. Our service architecture on top of it makes minimal use of centralized infrastructures for peer discovery.

Low-cost operation We have designed and implemented the system to run on off-the-shelf wireless routers and commodity handheld devices. Network access if free, while joining the community requires minimal cost; a user can share the excess capacity of his broadband connection via Wi-Fi and need not spare additional equipment nor operate dedicated servers. Open-source implementation assists in the system's inexpensive operation, while operation in ISM frequency bands facilitates deployment and makes wireless access provision possible without the complexities and cost of acquiring a license.

Security, trust and user rationality We have put much emphasis on the security aspects of our design. The Wi-Fi sharing protocol used is secure against hijacking, while the underlying decentralized accounting scheme protects from excessive free-riding. Free and disposable identifiers, while raising security challenges, enhance roaming anonymity. The tunneling-based approach, where two peers communicate via their home networks, enhances location privacy, since peers do not disclose their point of attachment to each other. (If the public IP address of the visited network is disclosed, the approximate location of a peer could be guessed.) Since we have not assumed that providers and consumers trust each other, this approach also protects users from eavesdropping and service providing peers from illegal activities of anonymous consumers. Some attacks are still possible, which we show how to tackle in Section 6.2.4, in order to achieve end-to-end security.

Our design assumes that users are rational entities; they always act to maximize their utility and never on pure altruism. Efstathiou [29] provides an extensive study on this issue at the access level. At the multimedia service provision layer, this principle is also respected. Service providing APs are not involved in the process of setting up secure communication, since they do not have a direct incentive to do so. (They still need to be aware that their visitors tunnel their traffic home, for security and legal purposes studied above.) End-to-end multimedia communication is performed in a provider-agnostic manner.

4.4 Performance evaluation

We first estimate upper bounds on the number of VoIP calls that can be supported in a community-based Wi-Fi access scheme using a simple throughput model. Then, we perform a set of experiments where we measure user QoE and quantify quality degradation factors that the model cannot capture, such as the processing overhead imposed by VPN mechanisms. In the following sections, the term *VoIP capacity* denotes the maximum number of simultaneous VoIP connections that can be handled in a specific scenario.

4.4.1 Upper bounds on VoIP capacity

Analysis

Prior work [43] reveals that even though the bitrate of a VoIP call may be small, the overhead imposed by the IEEE 802.11 PHY and MAC mechanisms and packet headers is such that an unexpectedly low number of concurrent VoIP sessions can be sustained by a Wi-Fi cell. In this section, we adapt the simple analytical model proposed by Garg and Kappes [43] to our scenario which involves two wireless last hops, to estimate an upper bound on the number of concurrent voice calls of acceptable quality. Here we focus on the effects of the PHY and MAC layers, as well as the packet overhead imposed by VPN mechanisms.

The IEEE 802.11 Distributed Coordination Function (DCF) is a CSMA/CA mechanism for arbitrating channel access. DCF dictates that for a station to transmit a packet, it should sense the medium idle for a specified time duration called DCF Interframe Space (DIFS). If so, the station enters a contention phase, where it senses the medium for a random number of slots. For each idle slot, it decrements a counter and transmission starts when the timer reduces to zero. If the medium is sensed busy at a contention slot, the station “freezes” its counter and needs to sense the medium idle again for a DIFS period before reentering the contention phase. The number of slots the medium should be sensed idle is drawn uniformly at random withing a *Contention Window (CW)*. It is possible that the backoff counters of two stations reach zero at the same timeslot. The stations will transmit and collision will occur. In a collision event, which is identified by the lack of an acknowledgment, stations retry transmission by drawing a new backoff counter value, doubling the size of the CW, after they have waited for a DIFS period. A retry limit is specified, after which a frame is considered *dropped*.

Table 4.3: Fixed overhead for the transmission of an IP datagram over IEEE 802.11. A variable delay is also introduced due to the DCF mechanism.

Overhead	**Size**	**IEEE 802.11b (@11Mbps)**	**IEEE 802.11g (@54Mbps)**
PHY		192 μs	20 μs
Slot time (T_{SLOT})		20 μs	9 μs
MAC header	272 bit	24.7 μs	5 μs
SIFS		10 μs	10 μs
DIFS[a]		50 μs	28 μs
ACK	112 bit	10 μs	2 μs
IP header	160 bit	14.5 μs	3 μs
UDP header	64 bit	5.8 μs	1.2 μs
RTP header	96 bit	8.7 μs	1.8 μs
Security (OpenVPN)[b]	680 bit	61.8 μs	12.6 μs
Security (L2TP/IPsec)[b]	896 bit	81.5 μs	16.6 μs

[a] $T_{DIFS} = T_{SIFS} + 2 \times T_{SLOT}$

[b] The space overhead due to tunneling varies, since padding is added to unencrypted data based on their size. The above values refer to the case for 60-byte unencrypted IP datagrams carrying 20 bytes of audio payload, a UDP and an RTP header.

The above procedure imposes significant MAC-layer overhead, which is more evident when small packets are transmitted, as is typically the case for VoIP services. Also, for each transmission, there is some physical layer overhead: A preamble used for synchronization and the Physical Layer Convergence Protocol (PLCP) header transmitted at a low rate precede each frame. Two preamble types are specified. The short preamble is 72 bits and is transmitted at 1 Mbps, but is not compatible with legacy IEEE 802.11 systems, therefore the long preamble is typically used in IEEE 802.11b. When using the short preamble, the PLCP header (48 bits) is sent at 2Mbps and the total PHY layer overhead adds up to 96 μsec. In contrast, using the long preamble (144 bits), the header (48 bits) is transmitted at 1Mbps and the total PHY layer overhead is 192 μsec. In a Wi-Fi BSS operating in pure IEEE 802.11g mode (when no legacy IEEE 802.11b devices exist) the physical layer overhead adds up to 20 μsec.

For IEEE 802.11b and IEEE 802.11g, fixed PHY, MAC and network (IP) layer overhead is presented in Table 4.3, where frame transmissions, except for the preamble and the PLCP header, are carried out at 11 and 54 Mbps, respectively. Depending on the higher-layer protocols used, overhead due to protocol headers increases.

A simple throughput model for our setup is given by

$$S = \frac{T_P}{T_P + T_{Overhead}} \times R, \tag{4.1}$$

where T_P is the time required to transmit the audio payload at rate R and $T_{Overhead}$ is the time overhead due to PHY, MAC and other higher-layer protocols. Our system handles constant bitrate, bidirectional flows. Let R_a denote the bitrate of a voice call. For G.729a, $R_a = 16Kbps$ ($8Kbps$ for each direction). The total number of simultaneous voice sessions is thus given by

$$n_{max} = \left\lfloor \frac{T_P \times R}{(T_P + T_{Overhead}) \times R_a} \right\rfloor, \tag{4.2}$$

where

$$T_{Overhead} = 2 \times (T_{DIFS} + T_{DCF} + T_{SIFS} + T_{ACK} + T_{Headers}) + T_P. \tag{4.3}$$

$T_{Headers}$ stands for the overhead imposed by MAC, IP, UDP, RTP and potential security related headers (OpenVPN or L2TP/IPsec). Note that $T_{Overhead}$ is the time overhead for *two* wireless transmissions, therefore Eq.(4.3) also includes the "serialization" time (T_P) for retransmitting the packet payload over the second wireless hop.

The only component in the above equation that varies with the number of stations and traffic is the overhead imposed by the DCF mechanism. In IEEE 802.11b, and under various assumptions, Garg and Kappes have approximated it as $T_{DCF} = 8.5 \times T_{SLOT} + T_W \times P_c$, where $T_W = T_{PHY} + T_{Headers} + T_{DIFS} + T_{SIFS} + T_P$ is the time wasted for a packet that has suffered collision and P_c is the collision probability calculated as $P_c = 0.03$. In the following results, we make the following simplifying approximations to derive an upper bound on VoIP capacity:

- The time a station spends waiting for the medium to become idle is composed only of idle slots counted during the backoff phase. Namely, we ignore the case when a station freezes its backoff counter when sensing a transmission and the DIFS time that it has to sense the channel idle before restarting its backoff procedure. Since the transmission of a VoIP packet occurs rarely (once every 20 ms) and occupies little time (31 μs for IEEE 802.11g), it is reasonable to assume that when operating below capacity, and given that VoIP traffic sources are not synchronized, the probability that a packet is generated during an ongoing transmission is low. With this approximation, we get $T_{DCF} \approx \frac{CW_{min}}{2} \times T_{SLOT} = 8 \times T_{SLOT}$ for IEEE 802.11g. It should be noted that this approximation becomes less accurate as the size of packets increases (e.g., due to the space overhead imposed by security mechanisms).

- There are no collisions when operating below capacity.

Table 4.4: Estimated maximum number of simultaneous VoIP calls in our wireless-to-wired-to-wireless scenario

	G.729a	**G.711**
IEEE 802.11b[a]- Unencrypted	7	6
IEEE 802.11b - OpenVPN	6	5
IEEE 802.11b - L2TP/IPsec	6	5
IEEE 802.11g[b]- Unencrypted	30	26
IEEE 802.11g - OpenVPN	27	25
IEEE 802.11g - L2TP/IPsec	27	24

[a] For IEEE 802.11b, we using the approximation of Garg and Kappes [43], where $T_{DCF} = 8.5 \times T_{SLOT} + P_c \times T_W$.

[b] For IEEE 802.11g, we use $T_{DCF} = \frac{CW_{min}}{2} \times T_{SLOT}$.

We numerically evaluate Eq. (4.2) and compare the maximum number of simultaneous voice connections that can be achieved when using unencrypted VoIP streams with the case when they are secured with OpenVPN or L2TP/IPsec. In the latter cases, the total overhead due to protocol headers is shown in Table 4.3. (Details on the structure of a VPN-secured packet are discussed in Section 4.4.2.) We also compare the use of G.729a with G.711, in each case sending 20 ms of audio payload per packet (20 vs. 160 bytes). G.711 trades bandwidth requirements for slightly better audio quality. Our results are presented in Table 4.4.

Silence suppression

A technique that can significantly improve capacity at the expense of implementation complexity is silence suppression by means of Voice Activity Detection (VAD) techniques. In this case, packets are not transmitted if no voice activity is detected, thus reducing bandwidth demands. When the start of a silence period starts, the transmitter signals the receiver with special packets so that the latter generates *comfort noise*. During a talk-spurt, packets are sent at a constant rate.

The performance benefits of using VAD cannot be easily quantified, because they heavily depend on actual voice activity during a session. A simplifying assumption would be that VAD reduces required bandwidth by 50%. However, various studies indicate that a voice activity detector creates silence and activity periods of different durations, thus the *Voice Activity Factor (VAF)* [1], which determines the required bandwidth, may be less than 0.5.

[1] The Voice Activity Factor, denoted as α, is the ratio of time the source is active.

ITU-T Recommendation P.59 [59] describes a method to generate artificial conversational speech based on a 4-state Markov model. It reports the temporal parameters of conversational speech, averaging the results of prior studies which where based on the analysis of speech samples (among which the popular work of Brady [18]). In particular, the reported average duration of a talk-spurt is 1.004 s and the respective duration of a silence period is 1.587, yielding a VAF of $a = 0.38$. Eq. 4.2 then becomes

$$n_{max} = \left\lfloor \frac{T_P \times R}{(T_P + T_{Overhead}) \times a \times R_a} \right\rfloor . \tag{4.4}$$

As shown in Figure 4.3, under the assumptions that (i) $a = 0.38$, (ii) the G.729a codec is used, (iii) each packet carries 20 ms of audio payload, and (iv) IEEE 802.11g at 54 Mbps is used, applying VAD techniques improves VoIP capacity from 30 to 79 simultaneous calls. For the same settings, IEEE 802.11b (11 Mbps) capacity could improve from 7 to 18 calls.

The effects of the payload size

A further option to increase VoIP capacity is to increase the size of the audio payload, for instance by packing multiple samples in a single packet. However, increasing audio payload can arguably make quality worse, due to increased delay, loss and jitter. Most commercial implementations use small payload sizes (10-30 bytes). Figure 4.3 shows the increase in voice capacity as the audio payload size increases in our scenario which involves two wireless hops. Note that while G.729a applies compression such that 10 ms of audio are encoded using 10 bytes, for G.711, 10 ms of audio expand to 80 bytes.

4.4.2 Experimental methodology

Our calculations in the previous section cannot capture the effects of cryptographic operations on VoIP capacity and involve simplifications about MAC layer behavior. To address these issues, but also to offer a user-centric view of how voice quality is perceived and study how transport level characteristics are associated with the achieved QoE, we performed a set of testbed experiments.

Again, we wish to estimate the VoIP capacity of our secure tunneling-based architecture when critical security functionality is build on typical home Wi-Fi devices.

The distinctive characteristics of our approach are that:

- Both call endpoints are assumed to be attached at community-operated Wi-Fi APs.

- Users tunnel their traffic to trusted VPN gateways. In our case, these gateways are collocated with each user's home Wi-Fi router (built into its firmware).

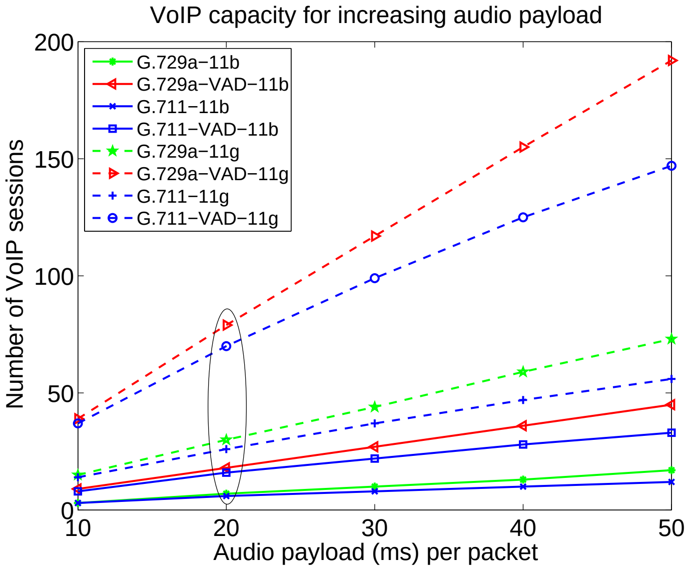

Figure 4.3: Evolution of VoIP capacity as the audio payload per packet increases, for the G.729a and G.711 codecs, also with the application of VAD for silence suppression. We have selected to experiment with 20 ms audio payloads per packet.

– A Wi-Fi sharing protocol may be in place. In our case, we apply the P2PWNC scheme, which incurs performance overhead due to public key cryptographic operations.

Thus, the requirements for a typical P2PWNC-enabled home WLAN AP are the following:

– Route mobile client's traffic using NAT [2].

– Operate the P2PWNC protocol.

– Act as the home VPN gateway for its owner[3], while the latter is visiting other (untrusted) P2PWNC APs. Encryption mechanisms and packet expansion due to additional headers cause processing and data overhead.

Our experimental methodology and testbed setup have to consider the above requirements and study their combined effect on the performance of a P2PWNC-enabled AP.

[2] In a typical WLAN setting, wireless clients are given private IP addresses via DHCP and the AP uses NAT to route traffic from/to the clients. In our architecture, most APs are expected to be connected with a single DSL line and, thus, have only one public IP address.

[3] Note that, at the same time, multiple tunnels with different endpoints may be active, since a single account may be in concurrent use by more than one users.

Testbed desrciption

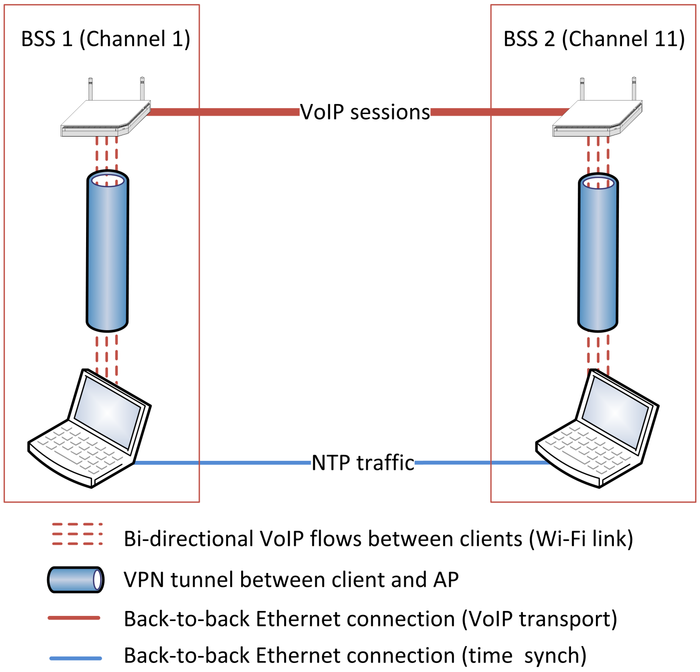
BSS 1 (Channel 1)
BSS 2 (Channel 11)
VoIP sessions
NTP traffic
Bi-directional VoIP flows between clients (Wi-Fi link)
VPN tunnel between client and AP
Back-to-back Ethernet connection (VoIP transport)
Back-to-back Ethernet connection (time synch)

Figure 4.4: Testbed setup for our VoIP quality measurements. Time synchronization between measurement endpoints is carried out over Ethernet.

VoIP call endpoints are laptops running Linux 2.6.38. We have implemented a set of measurement tools to generate constant bitrate bidirectional UDP streams to emulate VoIP traffic and have selected OpenVPN [92] to implement VPN tunnels, due to its configuration simplicity and popularity.

Since we wish to test the capabilities of standard low-cost WLAN equipment, we have used Linksys WRT54GL wireless routers running the Openwrt [93] Kamikaze 8.09 distribution (Linux 2.4.35). We have also included the necessary cryptographic libraries (OpenSSL 0.9.8) and a utility which periodically performs P2PWNC receipt verifications, to measure the effects of P2PWNC receipt operations on VoIP performance. Linksys WRT54GL APs are based on the Broadcom BCM5352 chipset, which includes a 200MHz MIPS processor,16Mb RAM and 4Mb of permanent storage (flash) where the firmware is also stored. A Broadcom IEEE 802.11b/g Wi-Fi adapter (controlled by a proprietary driver) is also included.

Transmission rate was fixed at 54 Mbps to disable rate adaptation and the RTS/CTS mechanism was disabled. Our testbed operated in pure IEEE 802.11g mode (i.e., the physical layer overhead was 20 μs, $T_{SLOT} = 20\mu s$ and $CW_{min} = 15$). Each AP operated in different non-overlapping channels (1 and 11).

To estimate the quality of voice calls we needed accurate measurements of network delay. We achieved this by comparing the timestamps generated at the transmitter and the receiver end for each voice packet. Transmitters and receivers were synchronized using NTP: One of the two laptops was operating an NTP server and the gigabit Ethernet interfaces of the two hosts were connected back-to-back, so that VoIP (wireless) traffic was isolated and the required accuracy level (of a few hundreds of microseconds) was achieved. In our experiments, packets were sent at fixed, 20 ms intervals, so synchronization is considered fairly accurate.

VoIP quality assessment

We emulated voice conversations by setting up bidirectional UDP flows between two laptop PCs. We implemented our own traffic generators, sending 50 packets per second with 20 bytes of audio payload each and 12 bytes for the RTP header. This traffic pattern corresponds to the G.729a codec, which is used by many available VoIP phones. The 20 bytes of packet payload contain 20 ms of voice. Each host was connected to a different IEEE 802.11g WLAN AP and each voice call lasted for at least 120 seconds. We assume that at the receiver end there is a dejitter buffer to ensure smooth playout at the expense of a constant 60 ms delay.

We initiated parallel VoIP calls between the two laptops and collected delay and loss information for each packet at the receiver end for one of the two call directions. Since our testbed is symmetric, the same results were observed for the opposite direction.

Our results reflect the perceived voice quality for a single call in the presence of simultaneous calls. For the VPN experiments, we also set up VPN tunnels between the laptops and the APs each one was connected to.

To estimate VoIP QoE, we used the evaluation methodology of Cole and Rosenbluth [23] to reduce ITU-T's E-model to transport level metrics. This methodology was presented in detail in Section 2.5.3. We can thus derive a score that represents the subjective quality of a voice call based only on network delay, jitter and packet loss information, which are directly measurable in our testbed. For the codec configuration described above, this score (*R-factor*) is given by the following formula (see Eq. (2.5) in Section 2.5.3):

$$\begin{aligned} R = 94.2 - 0.024 \cdot (d_{network} + 85) \\ - 0.11 \cdot (d_{network} - 92.3) \cdot H(d_{network} - 92.3) - 11 \\ - 40 \cdot \ln[1 + 10 \cdot (e_{network} + (1 - e_{network}) \cdot e_{dejitter})] \end{aligned}$$

where:

- $d_{network}$ is the end-to-end network delay in ms
- $e_{network}$ represents the percentage of packets lost in the network path
- $e_{dejitter}$ represents the packet discard ratio at the dejitter buffer of the receiver

– $H(x) = 1 \quad if \quad x > 0; \quad 0 \quad otherwise$

Furthermore, the R-factor can be mapped to a subjective Mean Opinion Score (MOS) as shown in Eq. (2.2). For a call of acceptable quality, average MOS should be over 3.60 (R-factor greater than 70).

Security parameters

The home wireless router operates also as a VPN gateway. In our prior work [31] we experimented with an L2TP/IPsec solution, building the Openswan IPsec implementation into the router's firmware. The L2TP protocol was used for implementing tunnels and IPsec ESP (*Encapsulating Security Payload*) [71] was used to secure them [100]. IPsec operated in transport mode using the AES-CBC algorithm (128bit keys) for data encryption. Preshared keys were used for authentication.

The above solution was complex protocol-wise, but also as far as configuration and maintenance were concerned. As to its space overhead, the original IP packet (IP, UDP and RTP headers and voice payload, adding up to 60 bytes) is encapsulated in a PPP frame (4-byte header), which is carried within an L2TP tunnel, thus an 8-byte L2TP header and an 8-byte UDP header are prepended to it. The resulting packet will be encrypted using the AES-CBC algorithm and encapsulated in an ESP header and trailer. The input data of the encryption algorithm also include the ESP "pad length" and "next header" fields (1 byte each) and their total length is 82 bytes. Before encryption, they are padded to become a multiple of the 16-byte AES-CBC block size, raising their size to 96 bytes. ESP packet contents also include the AES initialization vector (16 bytes), sequence number (4 bytes) and SPI index (4 bytes), as well as an integrity check value of 12 bytes (HMAC-MD5). Finally, the 132-byte packet is prepended with an IP header, adding up to a total of 152 bytes (compared to the 60 bytes of the unencrypted voice packet). Figure 4.5 shows the format of a tunneled VoIP packet using L2TP/IPsec.

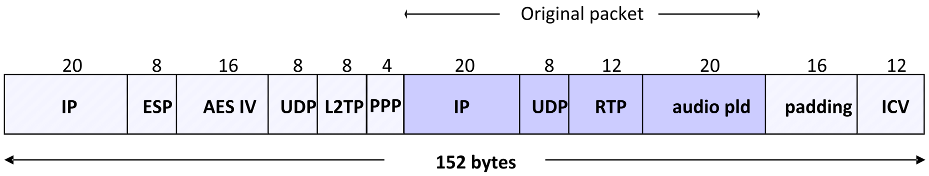

Figure 4.5: Per-packet space overhead for an encrypted/tunneled VoIP packet using L2TP/IPsec.

Note that in a practical P2PWNC-based scenario, NAT traversal would be used for IPsec communication, since users are typically in private LANs, and this would add to the space overhead (additional UDP encapsulation for traversing NAT).

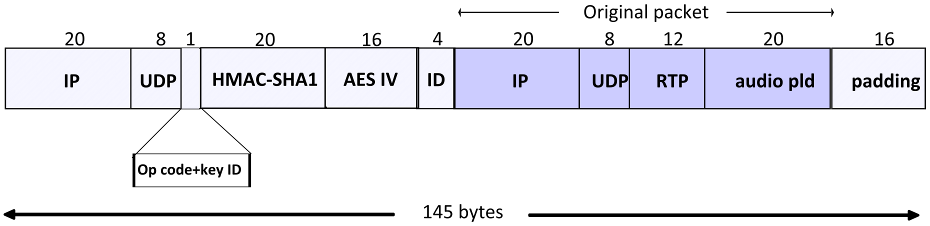

Figure 4.6: Structure of a tunneled packet using OpenVPN.

Instead of L2TP/IPsec, in this set of experiments we opted for an OpenVPN-based solution, maintaining the same security level (128-bit AES-CBC, but also certificate-based SSL/TLS session authentication and key exchange and HMAC-SHA1 packet authentication) and exploiting the popularity, configuration simplicity and wide availability of OpenVPN. In this case, the size of a 60-byte IP datagram expands to 145 bytes (compared to 152 bytes when using L2TP/IPsec). The structure of a packet tunneled over UDP using OpenVPN [113] and the above combination of encryption and authentication schemes is shown in Figure 4.6. Note that the size of the data to encrypt is 64 bytes (audio payload plus the 4-byte ID field). Since this is a multiple of the AES-CBC block size, a 16-byte padding is necessary.

P2PWNC parameters

The main overhead of the P2PWNC protocol is due to CPU-intensive receipt operations. We measured this effect by performing receipt verifications at the APs at regular intervals, instead of executing the full P2PWNC protocol exchange. It should be noted that receipt generation at the client side can also be CPU-expensive (especially when using RSA). However, depending on the number of ongoing P2PWNC sessions, the CPU cost of P2PWNC for the AP can be much larger than that of a single client and the CPU bottleneck which in turn affects VoIP performance is on the AP side. We argue that the effects of receipt signatures (even when these are executed on small form factor processors, such as those of handheld devices) on the performance of a single VoIP session are negligible compared to those due to verifications on the AP. Thus, for reasons of simplicity and clarity, we have decided not to carry out cryptographic operations on the client side in our experiments.

As to the P2PWNC-specific cryptographic parameters, we have used both the RSA (1024 bit keys) and the ECDSA (160 bit keys) algorithms for receipt generation and verification. Equivalent security to 1024 bit RSA keys can be achieved with 160 bit keys respectively [74]. As to the ECDSA-specific parameters, we have used the *secp160r1* verifiably random curve over the F_p finite field [114][115] to generate 160 bit keys.

Table 4.5: Receipt operations CPU times

	Security level (RSA / ECC)	
Key size (bits)	**1024 / 160**	**2048 / 224**
Generation (ms)	300.6 / 20.3	1529.0 / 23.4
Verification (ms)	12.3 / 114.7	37.9 / 135.7

4.4.3 Results

In this section we present the results derived from our measurements. For reference, we have included the pure CPU times needed for a P2PWNC receipt generation and verification on the Linksys WRT54GL platform (Table 4.5). It should be noticed that receipt verifications (public key operations) are performed slower when the ECDSA algorithm is used instead of RSA. The opposite holds for receipt generations, which involve digital signing using the issuer's private key. However, fast digital signing is more important for the typically battery-powered mobile devices[4]. This, combined with the reduced space overhead due to smaller key/signature sizes makes ECDSA more favorable from the viewpoint of mobile devices.

The figures presented in this section include end-to-end network delay, packet loss and dejitter buffer loss statistics for three types of experiments; (i) plain unencrypted VoIP transport, (ii) emulation of the space overhead introduced by OpenVPN tunneling, and (iii) VPN-secured VoIP transport.

For each experiment, we also calculate user-perceived VoIP QoE as a function of the above three quantities and for the codec settings we have assumed. Each data point is the mean of the measured values for 5 iterations of the same experiment, i.e., 5 VoIP calls under the same conditions, presented with 95% confidence intervals. For each iteration, the R-factor was calculated using the mean values for delay, network loss and dejitter buffer loss.

We also discuss results on the effects of P2PWNC operations on VoIP quality and report on our prior results obtained in an IEEE 802.11b testbed.

Experiment 1: IEEE 802.11g performance Our first experiment does not involve any security mechanisms. Instead, it measures transport level characteristics and quantifies user QoE for an unencrypted end-to-end VoIP session. The major quality degradation components are due to IEEE 802.11 MAC and PHY layer mechanisms. This experiment serves as a baseline case for measurements to follow. We find that 30 concurrent VoIP sessions can be sustained, with mean R-score > 75 (Figure 4.10). Beyond that point, network delay reaches more than 100 ms (Figure 4.7), which, combined with the intrinsic encoder delay (25 ms) and the playout delay introduced by the dejitter buffer, but also with noticeable network (Figure 4.8)

[4] P2PWNC receipts are generated by mobile nodes and verified by the service-providing AP.

and dejitter buffer (Figure 4.9) loss makes call quality unacceptable. Results of this experiment are shown as the blue curves in the figures.

The results of this experiment exactly match the upper bound derived using a simple analytic model in Section 4.4.1 (see Table 4.4).

Experiment 2: VPN space overhead In this experiment we quantify QoE degradation due to the space overhead imposed by VPN mechanisms. We do not carry out any cryptographic operations. On the contrary, we transmit packets of fixed size, equal to the size of 60-byte VoIP datagrams tunneled using OpenVPN. As discussed in Section 4.4.2, the size of each such IP datagram is 145 bytes. There is a noticeable drop in VoIP capacity in this case: Space overhead accounts for a 30% decrease in VoIP capacity, which is not fully captured by our analytic model. Instead of the measured maximum number of 21 concurrent high-quality VoIP sesssions (With 21 ongoing calls, R-score is still above the quality threshold, as shown in Figure 4.10.), we had estimated that 27 calls could be sustained. We believe that this is because our analysis underestimated the delay imposed by the IEEE 802.11 DCF. As packet size increases, the probability that a node senses the medium busy while being in its backoff stage increases. This causes its backoff counter to freeze, moving T_{DCF} beyond the optimistic $\frac{CW_{min}}{2} \times T_{SLOT}$ value that we had assumed.

Experiment 3: Combined VPN space and processing overhead Our 3^{rd} experiment focuses on the combined effect of VPN space and processing overhead. Note that we have assumed that each call endpoint operates a separate VPN gateway at his premises, building this functionality in the Wi-Fi router's firmware. Other options which would offer performance enhancements are possible (e.g., dedicating a more powerful device to this purpose or negotiating an end-to-end secure tunnel, which would involve a single encryption and decryption operation per packet), but our solution is more generic and spares the need for separate VPN equipment.

This time, performance is severely impacted. However, a reasonable number of parallel high-quality VoIP sessions is possible to support. In particular, 8 such calls can be simultaneously supported, maintaining very high quality (mean R-score > 80), as Figure 4.10 indicates. However, adding a single secure VoIP call causes all sessions to collapse: One-way network delay increases to more than 220 milliseconds, while 5-10% of transmitted packets are lost. Dejitter buffer loss was minimal in all our experiments. Packet loss obviously does not occur due to congested wireless links, since for the same traffic pattern without performing VPN-related cryptographic operations (see Experiment 2), loss was negligible. Rather, increased processing requirements at the APs incur queuing delays, which cause buffer overflows and thus packet loss.

As a final note, the AP processor is much slower than that of the laptops we used in our experiments. Tests we carried out using the OpenSSL `speed` utility revealed that the cryptographic throughput of encrypting 64-byte blocks using the AES-CBC algorithm (128 bits) on a Linksys WRT54GL was 2 MBps, compared to 90 MBps on

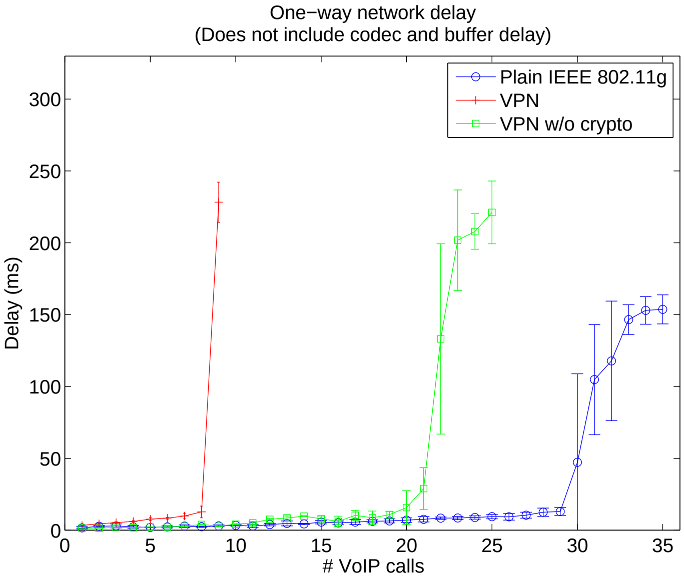

Figure 4.7: One-way network delay for a VoIP session in the presence of increasing numbers of parallel wireless-to-wired-to-wireless calls. Three scenarios are shown: The blue curve (circles) shows the case for a plain unencrypted G.729a call. The green curve (squares) presents a scenario where we emulated VoIP calls secured by OpenVPN, without performing encryption and decryption. The red curve shows the delay for an OpenVPN-secured VoIP call using AES-128.

an Intel i3 quad-core CPU (making use of a single core for the speed test, though) which we used in our experiments. The laptop was capable of encrypting/decrypting approximately 1.4M packets/sec, while 9 concurrent VoIP sessions generated and received 900 packets/sec. Therefore, for 9 concurrent VoIP sessions, the AP seems to be the bottleneck, even though in our testbed the laptop was handling the same amount of traffic to encrypt/decrypt.

Experiment 4: P2PWNC overhead In this set of experiments we measure the effects of CPU-expensive P2PWNC receipt verifications on VoIP performance. In our prior work [31] we observed that frequent receipt requests (in the order of once every few seconds – see also Experiment 5) have an adverse effect on service quality, since they require frequent costly verifications.

Receipt request frequency is a parameter which is controlled by the operator/owner of the community Wi-Fi router. The proper choice of this parameter depends on (i) the expected duration of a P2PWNC session, and (ii) the number of ongoing sessions, as well as (iii) traffic and CPU load on the AP. Frequent receipt requests ensure that little or none of the visitor's traffic will be unacknowledged at the end of the session, but at the same time increase load and in turn affect performance. One would expect

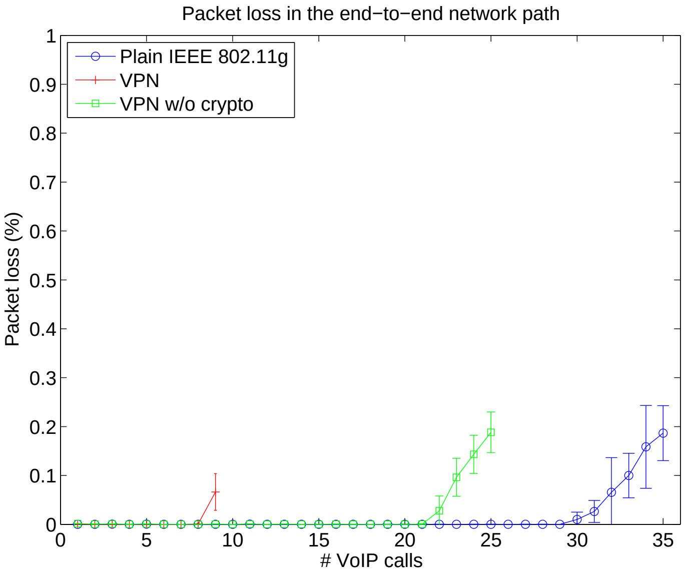

Figure 4.8: Network packet loss for a VoIP session in the presence of increasing numbers of parallel wireless-to-wired-to-wireless calls. Scenarios are the same as with Figure 4.7.

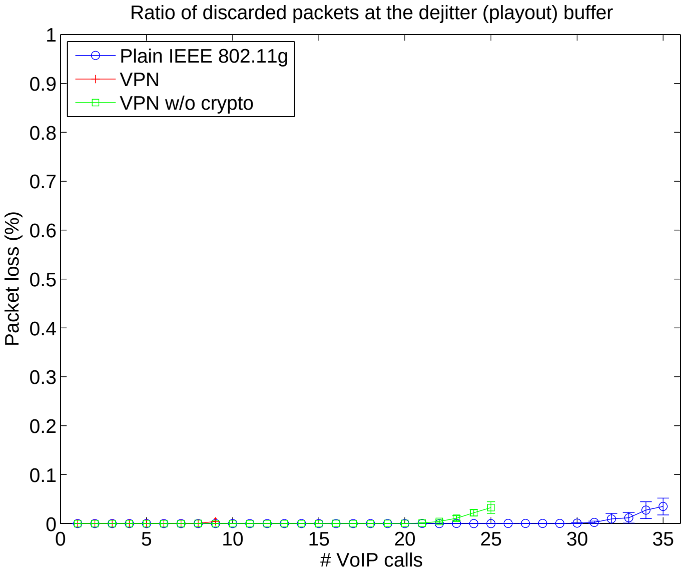

Figure 4.9: Percentage of packets dropped at the receiver dejitter buffer due to excessive jitter. Scenarios are the same as with Figure 4.7.

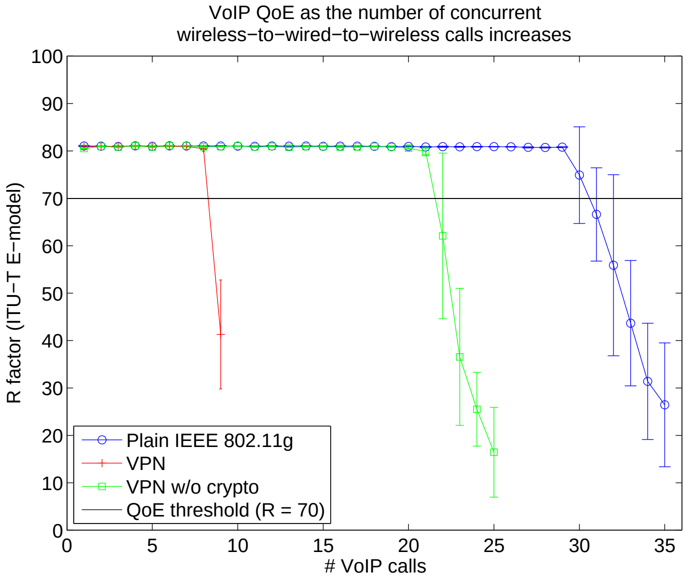

Figure 4.10: User QoE expressed in terms of the E-model's R-factor. Acceptable quality calls score 70 or more (solid straight line).

that a typical P2PWNC session would last at least a few minutes and the number of concurrent visitors would be limited. In such cases, a reasonable choice would be an order of one receipt every few tens of seconds or every few minutes.

Given the above discussion, we measured VoIP quality degradation for a fixed number of concurrent VoIP sessions, while increasing receipt verification frequency. We found that when the RSA algorithm is used, which is known to put little CPU overhead on the wireless router, 20 parallel unencrypted VoIP streams can be sustained with minimal QoE degradation, even when the AP performs (more than) 20 verifications/sec, i.e., requesting on average 1 receipt/sec from each associated client. Even in this extreme case, end-to-end one-way network delay is below 15 ms, and jitter and network packet loss are negligible, yielding an R-factor of approximately 80.

When using the ECDSA digital signature scheme, each receipt verification takes approximately 100 ms on the Linksys platform when no traffic is present. Therefore, no more than 10 receipts/sec can be verified in ideal conditions. Otherwise, the AP may not verify the receipt in time, i.e., before it is time for a new receipt request. When there is traffic to handle, the CPU is shared among routing and cryptographic operations, which further increases verification time, while also affecting network performance. Therefore, to sustain 20 parallel VoIP sessions, receipt request frequency should be reduced.

Finally, when VPN mechanisms are in place (see Experiment 3), the maximum number of parallel VoIP sessions is 8. We measured the QoE of a single call when 8

Table 4.6: Maximum number of VoIP calls of acceptable quality in an IEEE 802.11b setting. L2TP/IPsec are used for tunneling, instead of OpenVPN [31].

Scenario	**Supported calls**
Plain	7
P2PWNC (RSA 1024)	7
P2PWNC (ECDSA 160)	6
VPN - P2PWNC (RSA 1024)	5
VPN - P2PWNC (ECDSA 160)	2

VPN-secured VoIP sessions are ongoing. For high ECDSA receipt request/verification frequencies (e.g, 1 to 8 requests/sec), VoIP quality is unacceptable; network delay increases up to a few hundred milliseconds, while packet loss was more than 50% for frequencies of more than 4 requests/sec. Considering that the space overhead for 8 parallel VPN-secured calls is negligible, we find that this type of packet loss is due to delayed processing by the encryption/decryption engine, resulting in queue overflows. Acceptable quality (R-factor > 70) was observed when the AP carried out one verification every 10 seconds or more.

Note that if our target is to sustain only up to 4 parallel VPN-secured VoIP sessions, even a rate of 1 ECDSA receipt request per second yields acceptable voice quality.

Experiment 5: Comparison with IEEE 802.11b In our prior work [31, 38] we presented results from a similar set of experiments, albeit using IEEE 802.11b, but also L2TP/IPsec for implementing tunnels. Table 4.6 shows the maximum number of simultaneous voice calls of acceptable quality that can be supported in such settings. When no security mechanism is in place and P2PWNC is not in use, in our wireless-to-wired-to-wireless scenario, 7 calls can be sustained. This is in line with our theoretical upper bound of Section 4.4.1. We also experimentally verified the results of Garg and Kappes[43] that 14 calls can be admitted given our codec settings in a single-hop wireless-to-wired scenario, using MOS metrics.

It should be noted that in the experiments involving P2PWNC operations, we assumed that there is a one-to-one correspondence between a voice call and a client and that the AP requests a receipt every 5 seconds, which is a rather aggressive and costly choice, especially as the number of VoIP sessions (clients) grows. Even in such settings, a number of secure calls is still possible.

Summary of results Here we summarize the major results of our experimental evaluation.

- When no security mechanisms are in place, the upper bound on voice capacity is 30 parallel G.729a calls. The experimental result matches the one derived using a simple analytic model.

– Due to packet expansion caused by the application of OpenVPN as a tunneling scheme, we measured that voice capacity drops to 21 calls. Our analysis indicates that the upper bound is 27 calls. This discrepancy is due to the fact that our analysis underestimates delays due to contention for medium access (T_{DCF}), which increases end-to-end delay and packet loss.

– Operating OpenVPN in typical off-the-shelf Wi-Fi equipment is the major quality degradation component: Only 8 high-quality VoIP sessions can be sustained, and this is due to increased delay imposed by the CPU-intensive encryption and decryption operations at the AP. To the best of our knowledge, we were the first to quantify the performance of community-based secure VoIP communication in a setting where both call endpoints connect to community-owned Wi-Fi APs and VPN gateway functionality is implemented in home Wi-Fi equipment.

– Although P2PWNC operations are costly, appropriately tuning receipt request frequency can reduce their impact on VoIP quality: With realistic request frequencies of once per few minutes per session, VoIP quality is practically unaffected, even when VPN mechanisms are in place.

Chapter 5

Crowdsourced Wi-Fi topology discovery

Having studied the users' role in building wireless network infrastructure and secure services on top of it, we now focus on their role as providers of information about the radio environment. Depending on its intended use, such feedback could involve channel load, received signal strength and interference measurements, but could also include location information and higher-level metrics, such as perceived service quality at particular locations.

In this work, we are particularly interested in collecting information about the coexistence of wireless cells, i.e., overlapping wireless coverage, and thus constructing a view of the network topology. We focus on Wi-Fi networks, although our models are generic and could be applied to other cellular technologies with similar characteristics. With controlled Wi-Fi installations being set up in corporate premises, campus areas and public spaces, and residential WLANs being deployed in an autonomous and uncontrolled manner, Wi-Fi signals pervade modern metropolitan areas. Increased Wi-Fi coverage, an aftermath of the low cost, ease of installation and, significantly, operation in unlicensed spectrum, comes with the cost of interference, due to the very scarcity of unlicensed spectrum.

Numerous approaches have emerged that propose sophisticated coexistence mechanisms which focus, among others, on sharing spectrum across frequency (channel assignment) or space (mostly through power control or antenna directionality). To operate efficiently and maximize the offered service quality, these mechanisms rely on accurate network topology information, i.e., knowledge of the wireless neighborhood of Access Points (APs) and client presence.

In this work, we exploit the sensing capabilities of the terminal equipment of mobile users to reveal the wireless topology. Instead of relying solely on measurements carried out by the fixed infrastructure (APs), we argue and show quantitatively that *crowdsourcing* this task to users offers significant performance improvements. The technology to carry out such measurements and convey the results has recently been standardized in the IEEE 802.11k amendment [56]. However, the security aspects of this process have not received enough attention; clients who report on the wireless topology are not necessarily trusted. They may engage in fraudulent reporting which can severely impact the mechanisms that rely on accurate knowledge of the topology to operate. Here, we focus on these aspects in particular; we study specific attacks, comment on their potential impact and propose and evaluate simple measures to counter them.

Our work is relevant in a different context, too; wireless topology is extensively used for positioning services. Skyhook [112] is an example of a WLAN-based positioning service, which relies on a beacon database built by wardriving and user reports. Skyhook can determine a user's approximate location only based on nearby WLANs. For the same purpose, Apple maintains a similar database generated by reports sent by iPhones [8], while Google Latitude [44] and Windows Phone 7 Location Services [78] also rely on Wi-Fi geo-databases to detect a user's location. The above systems are based on the same crowdsourcing approach and, as pointed out in the literature [119], are vulnerable to attacks by untrusted clients.

In this chapter, we make the following contributions:

– We design and implement a user-centric, reputation-based topology discovery scheme tailored to managed Wi-Fi deployments, making use of standards-compliant mechanisms for security, authentication, and reporting.

– Using simple consensus-based mechanisms, we improve the robustness of our system against fraudulent reporting.

– We derive analytic expressions for the accuracy of our scheme and quantitatively justify the need for a crowdsourcing approach, while showing that, for realistic settings, the attacks we study can be effectively mitigated, even in the presence of large numbers of attackers.

This chapter is structured as follows: In Section 5.1 we motivate our work. Then, we present the design of our user-centric topology discovery architecture (Section 5.2), our system model (Section 5.3), reputation scheme, and relevant attacks and our proposed countermeasures (Section 5.4). In Section 5.5, we provide details on our system implementation. Under the attacker model we study, in Section 5.6 we present analytic expressions on the topology discovery accuracy of our scheme. Numerical results follow in Section 5.7. In Section 5.8 we discuss potential AP-centric extensions to further improve performance. We conclude this chapter positioning our scheme with respect to the principles of user centric networking in Section 5.9.

In Chapter 6 we discuss aspects of our approach which future research could build upon, including more sophisticated security attacks and, importantly, our position that the issue of collecting high-quality feedback from users could be jointly tackled with access and service provision.

5.1 Motivation

5.1.1 The need for a user-centric scheme

Many cases of overlapping coverage between neighbor Wi-Fi cells cannot be discovered in a scheme where this process is carried out only by the APs: An AP can directly detect neighboring cells only if the respective APs are in range. However, there may exist cases where there is overlapping coverage but the two APs are not in range of each other. In a pure AP-centric topology discovery scheme, such cases of overlap will be revealed only if there is another AP located there to report this incident. In Section 5.7.2, we quantitatively show that in realistic settings, a pure AP-centric scheme fails to reveal a significant portion of the wireless topology. This motivates us to focus on a user-centric scheme where determining the topology is crowdsourced to users.

User-based reports also help acquire a user-perceived view of wireless conditions; many reports about overlapping coverage between two cells are an indication of a conflict which an optimization mechanism should more urgently resolve. On the other hand, cases where no APs or users are placed in the overlapping region between two APs are not revealed, but are less important since no users are affected.

5.1.2 The need for secure crowdsourcing

Even though crowdsourcing offers measurable advantages, its security aspects should be carefully considered. We assume that users are not trustworthy by default and are expected to engage in fraudulent reporting. Their motives to perform such attacks vary; users may wish to avoid the performance overhead of monitoring, act strategically to manipulate the results of mechanisms relying on topology information, or act out of pure malice, but may also have faulty equipment. A study of such motives is outside the scope of this work.

Effects on efficient channel assignment

Channel assignment (CA) schemes depend on accurate knowledge of Wi-Fi topology. We assume that the set of admissible channels is $\{1, 6, 11\}$, i.e., the set of orthogonal IEEE 802.11b/g channels and there is a centralized CA algorithm which aims to minimize the *sum* of interference across the network[1]. The algorithm operates on a weighted graph whose vertices represent APs and an edge denotes overlapping coverage between two APs (see Section 5.3). Edge weights indicate the number of clients located in the overlapping regions and suffering interference if the APs are assigned the same channel. Clients report the APs in range and we assume, for now, that each report contributes a unit to the weight of the respective edges.

[1] This is a simplified version of the Hsum algorithm by Mishra et al. [85].

In the topology shown in Figure 5.1, the two APs are not in range of each other, but clients are located in the cell overlap area. After reporting, the graph shown is constructed. The weight of edge A-B is 2 and a number of fake vertices appear ($F_i, i = 1..12$), connected to valid ones with unit-weight (fake) edges. Attackers have reported that, for each of the two existing APs, there are 3 interfering APs in each of channels 6 and 11.

Aiming to minimize the network-wide sum of interference, APs A and B would be assigned the same channel (6). To the eyes of the frequency planner, the collective effect of interference suffered by APs F_7, F_8 and F_9, if A were assigned channel 11, is greater than that when channel 6 is selected. The same holds for channel 1 and the set $\{F_{10}, F_{11}, F_{12}\}$. In the same way, channel 6 is also assigned to B. Note that, if the two APs had resorted to local measurements, the A-B edge would not have been discovered, potentially leading to an inappropriate selection of channels.

Simulations we have carried out on the above scenario using *ns3* have shown a 45-50% average downlink UDP throughput reduction per user due to improper CA.

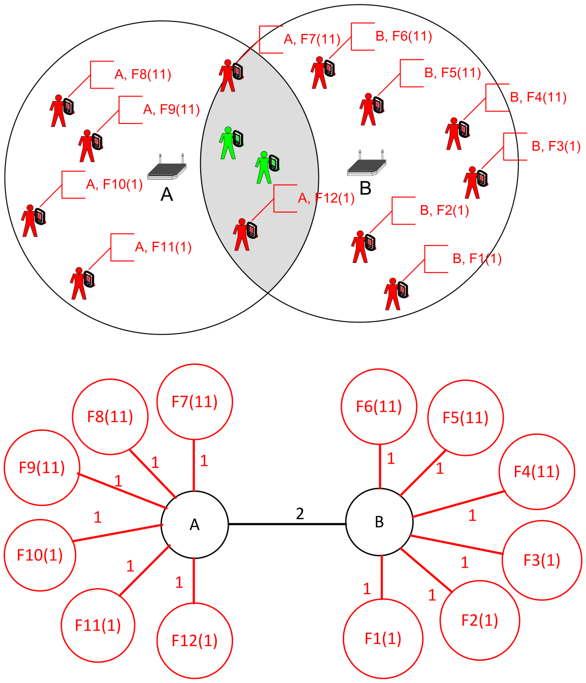

Figure 5.1: Example topology and the reported graph. Clients in green are honest and report that they are in range of both APs A and B. Clients in red are dishonest; each submits a report containing a random fake identifier and the identifier of the AP the client is attached to. The reported channel of each fake AP is in parenthesis. Each fake report contributes a unit-weight edge between the vertex corresponding to the AP each client is associated with and a fake vertex. The weight of edge A-B is 2, since there are two clients reporting it.

Effects on accurate WLAN-based positioning

Positioning systems based on public WLANs generally operate as follows: There is a central beacon database where Wi-Fi identifiers are stored, together with the location where beacons were recorded. This database is built by extensive wardriving, but also by crowdsourcing this task to mobile users. To discover one's location, the user's terminal scans for Wi-Fi presence and based on the list of Wi-Fi identifiers (MAC addresses) present, looks up its location in the beacon database. As Tippenhauer et al. have shown [119], such a location database can easily be manipulated; untrusted users can inject fake data, which can lead to localization errors. This also motivates us to propose methods to improve the security of crowdsourced Wi-Fi topology discovery.

5.2 Architecture

We have designed a topology discovery scheme tailored to deployments whose configuration is centrally controlled and user authentication, authorization and accounting (AAA) are centrally managed. This could be the case for a municipal or University campus WLAN, a WLAN aggregator, such as Boingo [17], managing hotspots belonging to various Wireless Internet Service Providers (WISPs), or even a wireless community network mediator such as FON [37]. Registered users can roam around APs belonging to different federated providers, while the aggregator tackles user registration and AAA.

The purpose of the operator is to collect information from the radio environment around registered APs in order to optimize their operation by tuning parameters such as the transmission power and frequency. To this end, registered APs are periodically requested to collect information about their neighborhood from registered users associated with them and submit them to a central *collector*. Users report details about Wi-Fi cell operation at their spot, such as the ID of each BSS (MAC address) and channel number, information which is carried in IEEE 802.11 beacon frames. Users are certified by the operator and their reports are authenticated. For reasons of robustness, APs also submit their own measurements to the collector. In order to participate in the reporting process, a user needs to be associated with a registered AP and be properly authenticated. We have implemented our scheme (Section 5.5) using standard protocols for authentication and security (IEEE 802.11i), and for collecting topology information (IEEE 802.11k [56]).

Based on received reports, the collector builds a wireless coverage map, on which, for example, CA algorithms can be applied. Since such operations are costly, both in terms of computation but also because they can disrupt the operation of Wi-Fi stations, we assume that they are executed periodically and not frequently. Therefore, a new snapshot of the network topology, i.e., the input to the above schemes, is generated following *reporting rounds*. An order of minutes or even few hours would be a reasonable choice for the interval between rounds. We assume that the collecting entity, APs and reporters are loosely synchronized (in the order of seconds) and devices submit reports, more or less, synchronously and upon request from the associated AP. Note that, although the APs belonging to the operator are under its control and can be carefully tuned, there are many other APs which may be interfering with managed ones. Our reporting scheme aims at also revealing such cases of cell overlap.

To ensure that the collected information is valid and to tackle fake reporting, we apply a reputation scheme. Each user has a reputation record, which is updated on each reporting round based on the information he submitted, which was eventually considered valid. Good reporters are promoted, while the reputation of dishonest ones is discounted and their reports have less weight. Managed APs are trusted by default and their measurements are used for checking the validity of user-provided information.

5.3 Topology model

We model Wi-Fi topology as a weighted undirected *Coverage Graph (CG)*, where vertices represent APs and edges represent coverage overlap between neighbor Wi-Fi cells. As shown in Fig. 5.2, there are two cases of overlap. In the first case (*Type-1* edges), two APs are within range of each other. In the second case (*Type-2* edges), two APs are not within range of each other, but stations (or APs) are located in the overlap area. We set the weight of an edge as a function of the number of reports about it, capturing user-perceived interference. High-weight edges should be more carefully considered while assigning channels or adapting the transmission power of the respective APs, since they affect more users. Our model is very similar to the one proposed by Mishra et al. [85]. Based on reports, the aim of our system is to expose as many CG edges as possible.

We only consider edges which connect two managed APs or a managed AP and a *foreign* one. Edges between foreign APs are irrelevant for the operator, since both APs are outside his control and he is unable to resolve such conflict. E.g., in Figure 5.2, which shows an instance of a CG where each report contributes a unit to an edge's weight, if APs A and B were not managed, the A-B edge would be ignored. Importantly, edges may not always be detected, due to lack of reports (e.g., when no managed AP, or very few clients, are there).

5.4 Trust model, attacks, and countermeasures

While the wireless infrastructure is trusted, this is not the case for users, who can engage in fake reporting, submitting fraudulent information to the collector. Each user is associated with a reputation value in the $[0, 1)$ range, which is a measure of how trustworthy the user is, and is used to weigh his reports. Consider edge e of the CG, which is reported by n users and APs, with r_i denoting the reputation value of

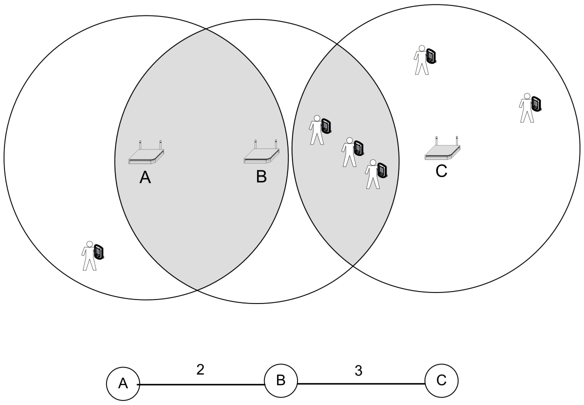

Figure 5.2: The Coverage Graph. A-B is a Type-1 edge and B-C is a Type-2 edge.

the $i-$th entity. We define the weight of this edge to be $w_e = \sum_{i=1}^{n} r_i$. Since managed APs are trusted by default, their reputation equals 1.

To counter potential attacks, we apply consensus-based rules to filter reports. We first make the following assumptions about attacker behavior:

- Each attacker acts independently[2]and submits reports containing a number of random fake AP identifiers.
- The probability that two or more attackers report the same *fake* edge is negligible.
- In a given reporting cycle, a potential attacker may choose not to attack and to report honestly.
- APs never attack and their reports are trustworthy.

A fake report can contribute fake edges to the reported CG, as well as fake vertices for each fake AP. Such edges connect a real vertex (corresponding to the AP the reporter is associated with) to fake ones. Also, fake edges among fake vertices can be added. In the attack scenario we study, the weight of a fake edge is *always* bounded by a unit-weight threshold (since we have assumed no collusion and $r_i < 1$). This leads us to the following observation.

Observation 1. Filtering all edges with weight less that 1 from the reported CG eliminates the possibility that a fake edge appears in it.

Thus, to combat this attack, we simply remove all edges with weight less than $T = 1$ from the reported CG. Also, since the only entities whose reputation always equals 1 are APs and this value is the edge acceptance threshold, we make the following observation:

[2] We will revisit this assumption in Section 5.7.3.

Observation 2. Edges reported by managed APs always appear in the filtered CG.

In each reporting round, the ratio of a user's reported edges that are not filtered is denoted as his *score* and each user's reputation is updated in a weighted manner based on it. We have used an exponential aggregation mechanism for user reputations, which is an adaptation of the metric proposed by Papaioannou and Stamoulis [98] (see Eq. (2.1)). In particular, in the place of the indicator function, we put the user's score, which is our measure of a user's successful interactions at a particular round. If, at round i, a user's reputation is r_i, after a new round and given that at round i his score is s_i, his reputation is updated as follows:

$$r_{i+1} = \beta r_i + (1 - \beta) s_i, \tag{5.1}$$

where β is a discounting factor used to appropriately weigh a user's past history of successful reporting.

Alternative reputation metrics are possible, though. For instance, the Beta aggregation function is often used [97], where the number of service provisions (reported edges, in our case) is encoded in the denominator. The number of service provisions and the number of successes, i.e., reported edges surviving filtering, are exponentially discounted with time. In our case, user reputations could be updated at the end of each round as follows [3]:

$$r_{i+1} = \frac{v'}{n'}, \tag{5.2}$$

where $v' = \beta v + v_i$ and $n' = \beta n + n_i$. v is the total number of edges that have been validated out of the n edges reported by the user across all prior rounds, and v_i is the number of edges validated out of the n_i edges reported by the user at round i.

We chose the simplified version of Eq. (5.1) because the number of edges reported by a user is not only a function of his behavior, but mainly has to do with the radio environment at his location.

All users start with a zero reputation, making reports by APs necessary for system bootstrap. In the first round, the only means for a user to have a score greater that 0 is to have some of its edges also reported by APs; edge discovery thus happens only due to reports by APs and, depending on the deployment density of managed APs, this may be relatively low. However, high initial reputation values for untrusted users make our system more vulnerable to collusion.

As we shall see in Section 4.4.3, as rounds progress and given that the population mix of honest reporters and attackers and the probability that a potential attacker chooses to behave honestly are fixed, the average reputation of the population of honest reporters is expected to increase and converge to a value that depends on the statistics of the topology (AP density, client density, etc.). The reputation of

[3] An alternative would be to perform the reputation update for each edge, i.e., $r_{i+1} = \frac{v'}{n'}$, where $v' = \beta v + \mathbf{1}(\text{edge validated})$ and $n' = \beta n + 1$.

consistent attackers, on the other hand, stays close to 0, since their scores, when they choose to attack, are also 0.

Note, finally, that this consensus-based scheme comes with the cost of filtering edges which are not reported by many clients. However, bearing in mind that we are particularly interested in how users perceive wireless conditions, few reports about an edge are an indication of few affected users and, thus, of less importance.

5.5 System implementation

5.5.1 Overview

A standards-based reporting architecture

In this section we present our implementation of the proposed user-centric reporting architecture. Our system involves implementing a subset of IEEE 802.11k in the wireless networking stack of the Linux kernel, user-space software running on APs requesting reports from authenticated clients and a centralized report collector which aggregates reported information, builds and filters a Coverage Graph and updates user reputations. The collector can execute channel assignment algorithms on the coverage graph and communicate the results to registered APs.

Since only reports from authenticated users are accepted, we have applied an IEEE 802.11i-based authentication, authorization and accounting scheme. The IEEE 802.11i [53] amendment focuses on security aspects of IEEE 802.11 networks. In particular, it deprecated the legacy protection mechanisms (WEP) by proposing stronger encryption and authentication schemes. A radius server, potentially collocated with the collector, is in place to manage user credentials. We have modified it so that it also maintains user reputations, which are updated per reporting round based on each user's score.

We have also implemented a set of reporting attacks on our IEEE 802.11k implementation, which can be easily activated from userspace software on the client side.

Our reference implementation is available for download[4].

Design alternatives

It should be noted that ours is an implementation based on ratified standards on measurements, security and authentication. However, design alternatives exist. A pure application layer solution is possible, where clients report their measurements directly to the collector over a secure channel over the Internet. Advantages of this

[4] `http://mm.aueb.gr/~pfrag/software`

approach include implementation simplicity (no need for kernel-level software development and no need to modify AP firmware) and the fact that clients could report their findings off line, or when connected to an AP that is not managed.

This very flexibility, though, comes with the cost of implementing a user authentication scheme dedicated to report collection; instead, we have chosen to reuse the security components of the IEEE 802.11i standard and thus believe our implementation to integrate better with a centralized WLAN infrastructure. Also, since in our scheme managed APs control report collection, we ensure that only properly authenticated clients currently attached to the APs are allowed to report, which, in turn, ensures the freshness of user feedback (be it honest or not). In the case of an application-layer alternative, an authenticated user can send fake reports even when he is not associated with a managed AP, making the system more vulnerable to attacks.

5.5.2 Background

IEEE 802.11 management plane

The IEEE 802.11 management architecture involves a Station Management Entity (SME), logically separated in two management entities, i.e., the Physical Layer Management Entity (PLME) and the MAC Layer Management Entity (MLME), which provide the service interface to invoke management functions. The SME can be viewed as residing in a separate management plane of the IEEE 802.11 architecture. Typically, it is expected that PLME functions are controlled by the MLME, where the PHY state machine resides.

Typical MLME primitives involve scanning, authentication, association, key management, and others. For details on the above operations, their semantics and parameters to invoke them, the reader is referred to the IEEE 802.11 standard text [55]. Importantly, spectrum management functionality, which includes spectrum measurements, channel switching and transmission power control, among others, are MLME primitives.

Wireless networking in the Linux operating system

Device types Based on how MLME functionality is implemented, in the Linux terminology there are two basic types of wireless devices, i.e., *FullMAC* and *SoftMAC*. The former are those where MLME operations are built in the firmware of the device, leaving limited control to the user. The latter, which have become more widespread, perform frame management in software, allowing for finer control of the hardware, since they offer more options to software developers.

For devices operating in station mode (i.e., *clients*), the current trend [77] is that MLME functionality is handled in kernel space. On the other hand, devices operating in master mode (i.e., *Access Points*) perform such management tasks from user space.

The mac80211 framework The core of the Linux Wi-Fi networking implementation is the mac80211 module. It is a framework that provides consistency between SoftMAC device drivers, offering them a unified interface for interacting with higher layers and implements MLME inside the Linux kernel.

Configuration API The traditional *wireless extensions (wext)* [121], which are based on an `ioctl` interface for configuration options, are still supported in the Linux kernel, although a new netlink-based[5] interface is intended to replace it. The new configuration framework includes a module (cfg80211) that handles device configuration and registration with the networking subsystem and a kernel API (nl80211) for communicating with user-space applications. Like the `iwconfig`, `iwlist` and `iwspy` user-space tools that are based on wext, new utilities have been developed for device configuration (e.g., the `iw` tool), which use the netlink interface (via the libnl user-space library) to invoke cfg80211 functions.

For SoftMAC devices, cfg80211 interfaces with mac80211, which implements cfg802-11 callbacks. For modern FullMAC devices, this functionality is implemented in the device driver. Legacy drivers bypass the mac80211/cfg80211 subsystem and communicate with userspace applications via the wireless extensions.

Figure 5.3 presents the architecture of the wireless networking stack in modern Linux kernels. It caters for drivers compatible with the mac80211/cfg80211 framework, but also for legacy drivers using the wireless extensions to communicate with user space.

User-space MLME

As of the current trend in the Linux operating system, for devices operating in master mode (APs), MLME functionality is performed in user space. The `hostapd` [51] daemon handles client authentication, key management, and other aspects of the wireless infrastructure. This daemon can operate with mac80211-based drivers, but also supports a few legacy ones. In the first case, `hostapd` communicates with mac80211/cfg802-11 via the netlink interface (nl80211). `hostapd` executes the MLME state machine, offers a control interface for other applications to use it to configure the device, implements the radius protocol [2] and can operate as a radius server or radius authenticator for IEEE 802.1X-based client authentication, utilizing the Extensible Authentication Protocol (EAP) [1] framework.

[5] Netlink is a socket interface which is typically used for communication between the kernel and user-space processes. The netlink protocol is specified in RFC3549 [107].

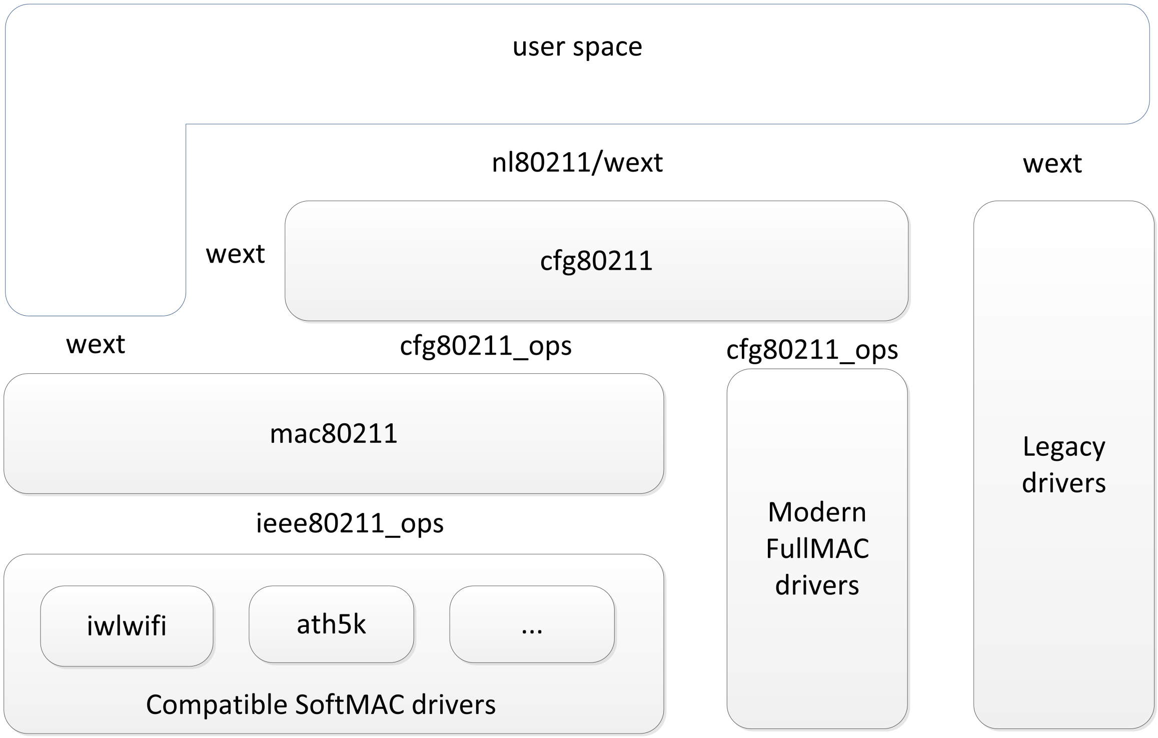

Figure 5.3: Architecture of the Linux wireless stack and the interfaces between its building blocks

5.5.3 Authentication, Authorization, and Accounting

Our system is suitable for managed WLAN deployments, where user authentication, authorization and accounting is centrally managed, as is also the case for network planning and configuration. We apply radius-based IEEE 802.11i authentication, where managed APs (NAS devices, in radius terminology) and users are registered with a central database which manages their credentials. Users are identified by username-password pairs. (Certificate-based authentication is also an option.) We have chosen EAP-PEAP as the authentication protocol. A user associates with an AP and his credentials are checked with the radius database. If they are valid, a session starts. The wireless link is protected by WPA2 (AES encryption) with frequent pairwise rekeying. We have extended the radius database so that it stores reputation records for registered users, which are updated per reporting round, based on the amount of truthful information submitted by each user, as explained in Section 5.5.5.

5.5.4 Reporting protocol

In our work, we make use of the IEEE 802.11k standard to request spectrum measurements and convey the results. Topology information can be encoded into *beacon report* action frames. Each beacon report includes information about the beacons received by the clients (i.e., the BSSs operating in the client's vicinity). An AP sends a specific *beacon request* action frame to an associated client, when requested by the application layer. For instance, in our case, a daemon running on the AP instructs `hostapd` to collect measurements from clients, responding, in turn, to the centralized collector controlling APs which executes a reporting round. Figure 5.4a shows the format of a measurement request, which includes the Information Element (IE) [6] that corresponds to a beacon request.

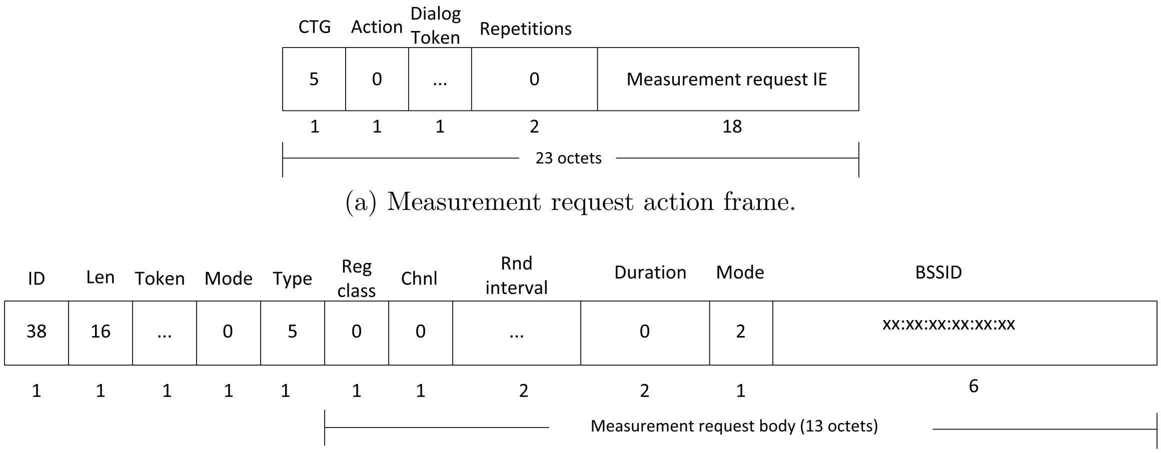

CTG	Action	Dialog Token	Repetitions	
5	0	...	0	Measurement request IE
1	1	1	2	18

(a) Measurement request action frame.

ID	Len	Token	Mode	Type	Reg class	Chnl	Rnd interval	Duration	Mode	BSSID
38	16	...	0	5	0	0	...	0	2	xx:xx:xx:xx:xx:xx
1	1	1	1	1	1	1	2	2	1	6

(b) Beacon request Information Element.

Figure 5.4: Beacon request frame.

Table 5.1 presents the values each field of a beacon request frame can take. Note that the first octet of an action frame denotes the type of the action frame, which, in our cases is "Radio measurement" (value 5). Also, note that action frames carry a (potentially variable) number of IEs.

Since we have opted for a user-space implementation of MLME functionality at the AP, beacon requests are generated by `hostapd`. A beacon request is sent to each associated client. Upon reception, a client replies with a beacon report. Based on the "measurement mode" field of the beacon request, the client is expected to perform an active or passive scan, or submit its beacon table. In the latter case, he submits information about all beacons received in any supported channel.

The beacon report carries one IE for each captured beacon, which includes information about the operating channel, the received signal strength, etc. Figure 5.5a shows the format of a typical beacon report action frame. Each of the IEs is 31 octets

[6] In the IEEE 802.11 terminology, IEs are blocks of information with a specific format that are contained in management frames. The format of each IE is specified in the standard and depends on the type of management frame they are included in. For instance, a beacon frame which announces BSS information includes IEs for various operating parameters.

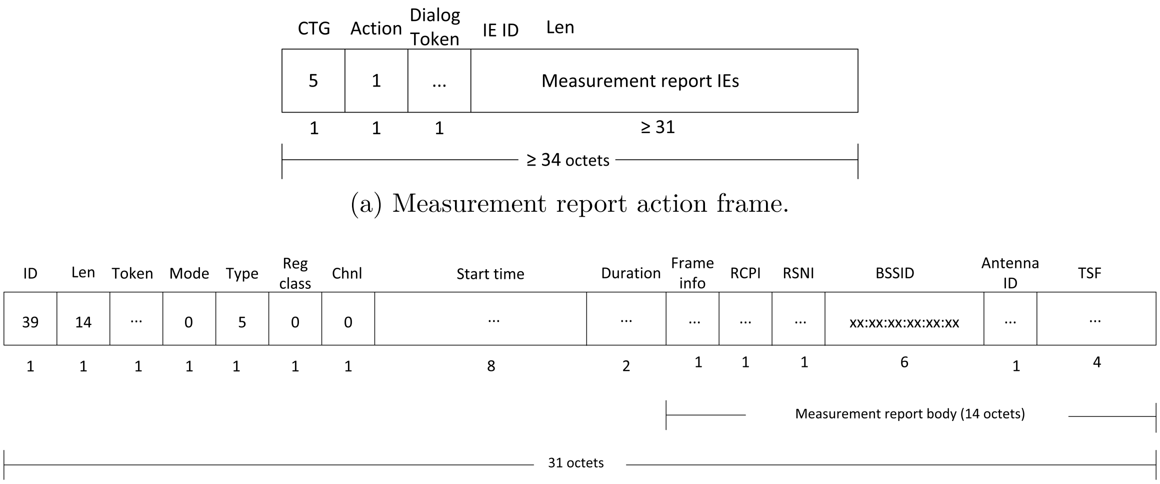

(a) Measurement report action frame.

(b) Beacon report Information Element.

Figure 5.5: Beacon report frame.

long and its format is shown in Figure 5.5b. Table 5.2 presents details about each field in a beacon report frame. Note that a list of optional fields can be present in the beacon request and report IEs.

Table 5.1: Beacon request fields

Field	**Description**	**Value**
CTG	Action frame category	Spectrum measurement (5)
Action	Type of action	Request (0)
Dialog token	Value generated at request time and identifies the request-report conversation	Random octet
Repetitions	Number of times the measurement should be repeated	0 for no repetitions
IE ID	IE identifier	Beacon request (38)
Len	Size of the beacon request IE not including the first 2 octets	16
Token	Random octet uniquely identifying each IE within the request frame	Random, ≥ 0
Mode	Bitfield controlling request parameters	Accept only reports performed on request by the AP (01100000)
Type	Measurement type	Beacon request (5)
Regulatory class	Channel set for which request applies	4 (2.4GHz, Europe)
Channel	Channel number	0 for all available
Randomization interval	Upper bound of the random delay prior to making the measurement	Ignored
Measurement mode	Type of beacon request	2 (Beacon table)
BSSID	BSSID for which a beacon report is requested	FF:FF:FF:FF:FF:FF (all BSSs)

Table 5.2: Beacon report fields

Field	**Description**	**Value**
CTG	Action frame category	Spectrum measurement (5)
Action	Type of action	Report (1)
Dialog token	The dialog token value included in the respective beacon request	Random octet
IE ID	IE identifier	Beacon report (39)
Len	Size of the beacon report body not including the first 17 octets	14
Token	Random octet uniquely identifying each IE within the report frame	Random, ≥ 0
Mode	Bitfield indicating reasons for a failed/rejected measurement request	0 (ok), 01000000 (incapable), 00100000 (refused)
Type	Measurement type	Beacon report (5)
Regulatory class	Channel set for which request applies	4 (2.4GHz, Europe)
Channel	Channel number	0 for all available
Start time	Measurement start time	Actual value of the STA's TSF counter when measurement started
Duration	Duration over which the Beacon Report was measured	Expressed in TUs
Frame information	Information about the PHY type over which frame was captured and type of reported frame	7 bits for PHY and 1 bit for frame type (0: beacon/probe rsp)
RCPI	Received channel power for reported frame	Measured in dBm
RSNI	Received signal to noise indication for reported frame	Measured in dB
BSSID	The BSSID of the reported frame	AP MAC address
Antenna ID	ID of the antenna(s) used for this measurement	0: unknown, 1: single antenna STA
Parent TSF	STA's TSF timer at the start of reception of the first octet of reported frame timestamp	Lower 4 octets of TSF timer

5.5.5 Topology information collection

Collector

Topology information is collected in reporting rounds. The interval between rounds is configurable by the network operator, but an order of minutes or even few hours is a reasonable choice, given that the purpose of this process is to come up with a new channel assignment and one would not expect this to occur very often.

We have implemented a centralized report collector, which periodically requests topology information from registered APs. Based on received reports, it builds the CG, filters it to remove potentially fake information and executes a channel assignment algorithm for the APs under its control.

The report collector could be collocated with the radius server, but this is not a requirement of our system. The basic requirement is that the collector has access to the radius database, so that it can retrieve the list of managed APs which participate in the collection process, as well as registered user identities and their reputation values.

The reports submitted by a user are evaluated based on his reputation, as described in Section 5.4. For example, if, at round k, user i reports a CG edge e and his current reputation value is r_i, then the weight of e at the reported CG is incremented by r_i. At the end of the round, the graph is filtered (see Section 5.4), user scores are calculated and their reputations are updated accordingly in the radius database.

AP software

A daemon running on the AP (`11kd`) is responsible for collecting reports from authenticated users associated with it, responding to requests from the collector. Upon receiving a request to collect measurements, `11kd` communicates with `hostapd`, which we have appropriately extended to

- handle communication with `11kd` using IPC methods, and
- transmit IEEE 802.11k beacon requests to stations and receive beacon reports.

In our reference implementation, communication between `11kd` and `hostapd` is carried out over a message queue; `11kd` informs `hostapd` that a new reporting round is in progress, while the latter sends each client a beacon request. Each beacon report received is placed in the message queue. When all clients have responded (or a timeout has occurred), `11kd` compiles a report batch and submits it to the collector.

```
RBAT
Content-Length: 90
Num-bss: 2
00:0b:6b:4e:e0:4a/00:0f:66:c7:ab:52/11
00:0b:6b:4e:e0:4a/00:0b:6b:4e:63:46/4
```

Figure 5.6: A report batch (RBAT) message

Collector-AP communication

Communication between `11kd` (AP) and `collectord` (collector) is carried out using a simple text-based protocol on top of either TCP or UDP. It should be noted, though, that other options would be possible. For example, an SNMP-based solution could be applied, both for the collection of topology information, but also for configuring the operating channels or other parameters of the managed APs. Below we summarize the set of messages that can be exchanged between the collector and an AP.

Report request (RREQ) This message is sent from the collector to each registered APs. It has no other parameters.

Report batch (RBAT) This message is compiled by the AP when IEEE 802.11k beacon reports have been received by stations. Its body includes <*station MAC - BSSID - channel*> tuples (but can be extended to encode other BSS information, such as the received signal strength). An example RBAT message is shown in Figure 5.6; it includes the report of a single station, which has two APs in range, one operating in channel 11 and one in channel 4.

Channel assignment (CASS) The collector can instruct an AP to switch to a new channel, according to the outcome of the execution of a channel assignment algorithm. Note that the details of such algorithms are outside the scope of our scheme.

Our architecture and protocol operations are shown in Figure 5.7.

5.5.6 Implementing attacks

To demonstrate that attacking the IEEE 802.11k protocol is feasible, we have implemented three different attacker strategies that clients can adopt when requested to submit beacon reports:

1. Ignore IEEE 802.11k beacon requests.

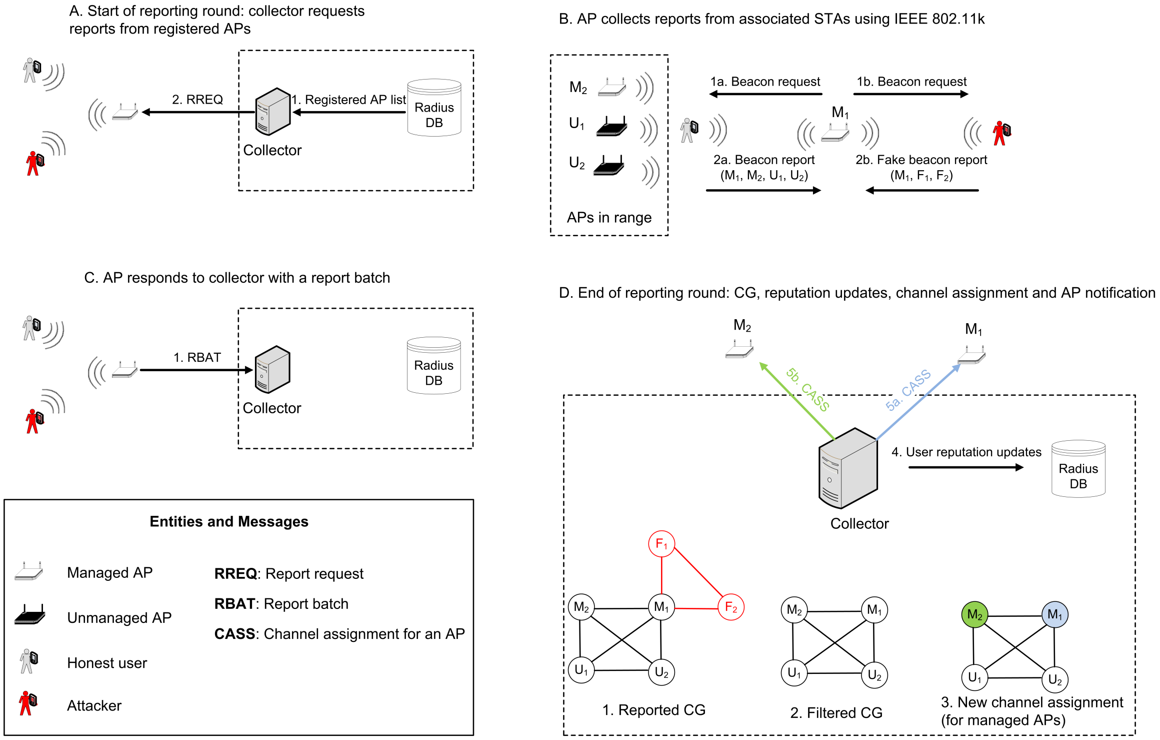
A. Start of reporting round: collector requests reports from registered APs
2. RREQ
1. Registered AP list
Radius DB
Collector
B. AP collects reports from associated STAs using IEEE 802.11k
M_2
U_1
U_2
APs in range
1a. Beacon request
1b. Beacon request
M_1
2a. Beacon report (M_1, M_2, U_1, U_2)
2b. Fake beacon report (M_1, F_1, F_2)
C. AP responds to collector with a report batch
1. RBAT
Collector
Radius DB
D. End of reporting round: CG, reputation updates, channel assignment and AP notification
M_2
M_1
5b. CASS
5a. CASS
4. User reputation updates
Radius DB
Collector
F_1
F_2
M_2
M_1
U_1
U_2
1. Reported CG
2. Filtered CG
3. New channel assignment (for managed APs)
Entities and Messages
Managed AP
Unmanaged AP
Honest user
Attacker
RREQ: Report request
RBAT: Report batch
CASS: Channel assignment for an AP

Figure 5.7: An IEEE 802.11k-based Wi-Fi coverage reporting architecture. This figures presents the steps carried out in a reporting cycle and the relevant protocol messages and operations.

2. Report a random fake set of BSSIDs.

3. Report a predefined set of BSSIDs.

The above strategies are implemented purely in software, by modifying the mac80211 kernel module. A malicious user can activate them from user space via the *proc* file system. He can simply write the appropriate data to a specified file within `/proc` to set the attack mode and relevant data. For example, for the second attack, the user should define the number of random BSS identifiers to generate, while for the third attack, the BSS identifiers and channel numbers should be set. Then, upon receiving a beacon request, the mac80211 module compiles a beacon report with the fake data. The third attack mode can be used by a number of colluding users who have agreed to report the same fake set of BSSIDs to make their (fake) feedback mode credible.

The fact that the number on Linux-powered user equipment, such as Android mobile phones, keeps increasing, makes attack potential higher. Report authentication and encryption mechanisms on their own help mitigate specific attacks, such as multiple fake reports by the same entity, but cannot stop a user from faking the submitted information.

5.6 Performance analysis

5.6.1 Preliminaries

Assumptions

We assume that stations (clients) and APs are distributed according to homogeneous spatial Poisson Point Processes in the 2D space, with intensities λ_c and λ_{AP}, respectively. We also assume an idealized model where AP coverage is a disk of (fixed) radius R and each client and AP in range can decode AP beacons. Each AP is centrally managed with uniform probability p_m and each user associates with a random *managed* AP in range. If none of the APs in range are managed, the user does not participate in the reporting process. The system operates in rounds. AP locations are always fixed, while at each round clients move to new locations, independently from their previous ones.

There are two types of clients; users that are always honest (with probability p_t) and potential attackers (with probability $1 - p_t$) and, for each of the scenarios we study, the ratios of the two types of users over the total population are fixed. Attackers follow the model we described in Section 5.4. At any round, a potential attacker may choose to attack with probability p_a.

Distance distributions

The probability that a CG edge is detected is a function of the size of the area of overlap between the two respective APs. The latter, in turn, depends on the distance between the two APs. Since cell radius is assumed fixed, the area of overlap $A(x)$ between two neighbor APs whose distance is $x \in [0, 2R]$ can be calculated as the intersection of two circles with equal radii and its size is double the size of the shaded area in Figure 5.8. E_s is the area of a circular sector of angle 2ϕ and it is given by

$$E_s = \pi R^2 \frac{2\phi}{2\pi} = R^2\phi = R^2 cos^{-1}(\frac{x}{2R}). \tag{5.3}$$

E_T is the area of the triangle AK_1B (two times $E_{AK_1\Gamma}$, i.e., the area of triangle $AK_1\Gamma$), which is given by

$$E_T = 2E_{AK_1\Gamma} = \frac{x}{2}y = \frac{x}{2}\sqrt{R^2 - \frac{x^2}{4}} \tag{5.4}$$

Therefore,

$$\begin{aligned} A(x) &= 2(E_s - E_T) \\ &= 2(R^2 cos^{-1}(\frac{x}{2R}) - \frac{x}{2}\sqrt{R^2 - \frac{x^2}{4}}) \\ &= 2R^2 cos^{-1}(\frac{x}{2R}) - \frac{x}{2}\sqrt{4R^2 - x^2}. \end{aligned} \tag{5.5}$$

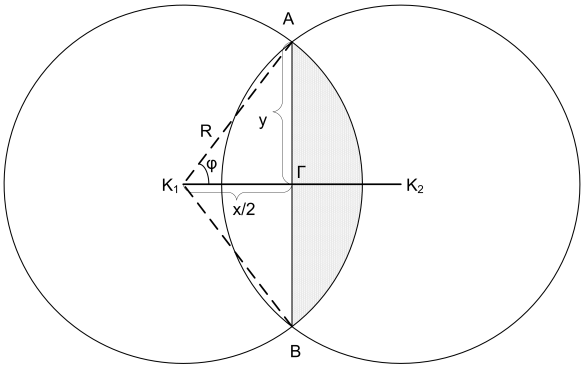

Figure 5.8: Circle-circle intersection

We have then calculated the CDF of the distance between an AP and a random neighbor AP (i.e., a random AP in a $2R-$radius disk centered at the former AP) as:

$$F_X(x) = P(X \leq x) = \frac{\pi x^2}{4\pi R^2} = \frac{x^2}{4R^2},\ 0 \leq x \leq 2R, \tag{5.6}$$

thus the respective PDF is given by

$$f(x) = \frac{x}{2R^2},\ 0 \leq x \leq 2R. \tag{5.7}$$

Finally, to be able to calculate user scores (see Section 5.6.3), we need an expression of the distribution of the distances between reported APs, i.e., APs located in a disk of radius R centered at the client. The CDF of the distance between two randomly picked such APs is given by ([16]):

$$L_X(x) = 1 + \frac{2}{\pi}\left(\frac{x^2}{R^2} - 1\right) cos^{-1}\left(\frac{x}{2R}\right) - \frac{x}{\pi R}\left(1 + \frac{x^2}{2R^2}\right)\sqrt{1 - \frac{x^2}{4R^2}}, 0 \leq x \leq 2R, \tag{5.8}$$

from which we derive its PDF $l(x)$:

$$l(x) = \frac{4x}{\pi R^2} cos^{-1}\left(\frac{x}{2R}\right) - \frac{x^2\sqrt{4R^2 - x^2}}{\pi R^4}, 0 \leq x \leq 2R. \tag{5.9}$$

Evaluation metric

Our performance evaluation metric is the ratio of discovered CG edges to the existing ones:

$$\mathcal{E} = \frac{N_d^{(1)} + N_d^{(2)}}{N_e^{(1)} + N_e^{(2)}}, \tag{5.10}$$

where $N_e^{(1)}$ and $N_e^{(2)}$ denote existing Type-1 and Type-2 edges, while $N_d^{(1)}$ and $N_d^{(2)}$ are discovered Type-1 and Type-2 edges. Based on the observations of Section 5.4, in the attack scenarios that we study, it is not possible for fake edges to appear in the CG. Therefore, the performance of our scheme is only limited by *false negatives*, i.e., by real edges that did not meet the filtering threshold or by cases where a CG edge did not get reported.

Number of CG edges

Instances of cell overlap are potential CG edges: If the two overlapping APs are in range of each other, an edge exists in the CG if at least one of the APs is managed. If the distance between the two APs is in the $(R, 2R]$ range, the respective edge exists only if one of the two APs is managed and clients or APs are located in the overlapping region.

From the above discussion, and given the distribution of distances between neighbor APs (i.e., $f(x)$), we can derive formulae for calculating the estimated number of CG edges. For Type-1 edges, we have

$$N_e^{(1)} = N_{pe} F_X(R), \tag{5.11}$$

while for Type-2 edges

$$N_e^{(2)} = N_{pe} \int_R^{2R} f(x)(1 - e^{-(\lambda_c + \lambda_{AP})A(x)})dx, \tag{5.12}$$

where N_{pe} is the number of pairs of APs with overlapping coverage where at least one of the two APs is managed (potential edges).

5.6.2 The performance of a pure AP-centric scheme

Intuitively, involving clients in the topology discovery process should contribute in discovering more cases of cell overlap. Here, we explore the performance limitations of a pure AP-centric scheme and demonstrate quantitatively the need for a user-centric one. We show that for realistic wireless deployments, and as the density of clients increases, relying solely on APs is not adequate.

Type-1 edges are always detected, since, by definition, at least one of the two APs involved is managed and will report the edge, therefore $P_d^{(1)}(x) = 1, \forall x \in [0, R]$, where $P_d^{(1)}(x)$ is the probability that an $x-$distance Type-1 CG edge is discovered. Things are different for Type-2 edges. Since the two APs are not in range of each other, a managed AP should be located in the overlap area. Otherwise, the edge is missed since clients located there do not participate in the reporting process. Therefore, an edge is missed when

- there are clients or unmanaged APs in the overlapping region, and
- there is no managed AP there to report it.

The probability that an $x-$distance Type-2 edge is discovered in the pure AP-centric scheme is given by

$$P_d^{(2)}(x) = 1 - (1 - e^{-(\lambda_u+\lambda_c)A(x)})e^{-\lambda_m A(x)}, \tag{5.13}$$

where λ_c, λ_u and λ_m are the intensities of the distributions of clients, unmanaged APs and managed APs respectively. Note that the above distributions are independent; the Poisson process according to which APs are distributed is split with probability p_m, with $\lambda_u = (1 - p_m)\lambda_{AP}$ and $\lambda_m = p_m\lambda_{AP}$.

The estimated number of discovered Type-2 edges is

$$N_d^{(2)} = N_{pe}\int_R^{2R} f(x)[1 - (1 - e^{-(\lambda_u+\lambda_c)A(x)})e^{-\lambda_m A(x)}]dx, \tag{5.14}$$

while all Type-1 edges are discovered, i.e., $N_d^{(1)} = N_e^{(1)}$. The efficiency of an AP-centric scheme then follows from Eq. (5.10).

5.6.3 The performance of a user-centric scheme

Number of discovered CG edges

In the user-centric scheme, the conditions for a potential edge to be discovered are as follows:

- At least one managed AP is located in the overlapping region, or
- A sufficient number of clients exist such that the sum of their reports meets or exceeds the (unit) threshold. For an $x-$distance edge at round i, we derive a formula for this probability (denoted as $P_d^{[i]}(x)$) in Section 5.6.3.

As with a pure AP-centric scheme, all Type-1 edges are reported by at least one of the two APs involved and thus discovered. The expected number of detected Type-2 edges is given by

$$N_d^{(2)} = N_{pe}\int_R^{2R} f(x)P_d^{[i]}(x)dx. \tag{5.15}$$

In the remainder of this section, we step through the calculation of $P_d^{[i]}(x)$.

Isolation probability

There are cases where a user cannot contribute reports because no managed APs (mAP) are in range, so he cannot be authenticated with the collector. A special case is when a user has a single managed AP in range and no unmanaged (uAP) ones, and thus cannot report any cases of cell overlap. We term these users *isolated*. If, at round i, a user is isolated, his reputation does not get updated. The probability that a user is *isolated* is given by:

$$\begin{aligned} p_i &= Pr\{0 \text{ mAP}\} + Pr\{1 \text{ mAP}\}Pr\{0 \text{ uAP}\} \\ &= e^{-\lambda_m \pi R^2} + \lambda_m \pi R^2 e^{-\lambda_m \pi R^2} \cdot e^{-\lambda_u \pi R^2} \\ &= e^{-\lambda_m \pi R^2} + \lambda_m \pi R^2 e^{-\lambda_{AP} \pi R^2}. \end{aligned} \tag{5.16}$$

Efficiency ($\mathcal{E}$) at round i

The number of reports necessary for an edge to be accounted for varies and depends on the mix of honest users, potential attackers who do not attack at a particular round, and managed APs located in the overlap area. To begin with, let $\overline{r_a^{[i]}} > 0$ and $\overline{r_t^{[i]}} > 0$ denote the mean reputation of honest users and attackers respectively at round i. By design, each report weights as much as the respective user's reputation. At round 0, $\overline{r_t^{[0]}} = \overline{r_a^{[0]}} = 0$ and efficiency equals that of an AP-centric scheme.

Again, Type-1 edges are always detected, so we focus on Type-2 ones. Consider two APs whose distance is $x \in (R, 2R]$. To calculate the discovery probability for an $x-$distance edge, we calculate the probability that the edge is missed or filtered by counting all the possible outcomes such that the edge does not meet the filtering threshold. We introduce the following notation as to the random variables denoting the number of entities located in the overlapping region:

- A: number of managed APs in the overlap area; Poisson distributed with mean $\lambda_m = p_m \lambda_{AP}$.
- X: number of honest users; Poisson distributed with mean $\lambda_t = p_t \lambda_c$.
- Y: number of potential attackers who choose not to attack at this particular round; Poisson-distributed with mean $\lambda_a = (1 - p_t)(1 - p_a)\lambda_c$.

The distributions of the above classes of users, as well as of managed (and unmanaged) APs are independent Poisson Point Processes. Assuming that $X = j$ and $Y = k$, honest users contribute $j\overline{r_t^{[i]}}$ to the weight of the edge. If there are no managed APs in the overlap area and $j\overline{r_t^{[i]}} < T$, where $T = 1$ is the filtering threshold, the edge will be filtered if non-attacking attackers cannot contribute the remaining $T - j\overline{r_t^{[i]}}$. If $j\overline{r_t^{[i]}} \geq T$, on the other hand, the edge is discovered irrespective of the distribution of managed APs or potential attackers. In other words, if no managed APs are in the overlap area and $j < \frac{T}{\overline{r_t^{[i]}}}$, at least $k \geq \frac{T - j\overline{r_t^{[i]}}}{\overline{r_a^{[i]}}}$ non-attacking attackers should be present. Therefore, the discovery probability for an x-distance edge at round i follows (by independence):

$$\begin{aligned} P_d^{[i]}(x) &= 1 - \sum_{j=0}^{\left\lfloor \frac{T}{\overline{r_t^{[i]}}} \right\rfloor} \sum_{k=0}^{\left\lfloor \frac{T - j\overline{r_t^{[i]}}}{\overline{r_a^{[i]}}} \right\rfloor} Pr\{A=0\}Pr\{X=j\}Pr\{Y=k\} \\ &= 1 - e^{-(\lambda_m+\lambda_t+\lambda_a)A(x)} \sum_{j=0}^{\left\lfloor \frac{T}{\overline{r_t^{[i]}}} \right\rfloor} \sum_{k=0}^{\left\lfloor \frac{T - j\overline{r_t^{[i]}}}{\overline{r_a^{[i]}}} \right\rfloor} \frac{(\lambda_t A(x))^j (\lambda_a A(x))^k}{j!k!}. \end{aligned} \tag{5.17}$$

From (5.10), (5.15), (5.17), and setting $N_d^{(1)} = N_e^{(1)}$ (all Type-1 edges are discovered), we get the value of $\mathcal{E}$ at round i.

Score calculation

A user's reputation is updated based on his score, i.e., the ratio of the edges reported by the user that meet the weight threshold. For the attack scenario studied, our filtering mechanism guarantees that all fake reported edges are removed from the CG, therefore an attacker's score is always 0. Reputations of isolated users are not updated, since they do not report any edges.

We focus on the calculation of an honest user's score (a potential attacker's score is calculated in a similar fashion). First, apart from the random variables A, X and Y defined in Section 5.6.3, we define the following events:

- D: An $x-$distance edge is discovered.
- B: An $x-$distance edge is reported by (at least) an honest user, i.e., $X > 0$.
- C: The sum of the weights of all client reports exceeds the filtering threshold.

The probability that an $x-$distance edge contributes to a user's score (i.e., is detected, given that it is reported by the user) is

$$P_s^{[i]}(x) = Pr\{D|B\} = \frac{Pr\{D \cap B\}}{Pr\{B\}}. \tag{5.18}$$

$Pr\{B\}$ easily follows from the Poisson distribution. $Pr\{C\}$ is calculated using a similar approach as with (5.17):

$$Pr\{C\} = 1 - \sum_{j=0}^{\left\lfloor \frac{T}{\overline{r_t^{[i]}}} \right\rfloor} \sum_{k=0}^{\left\lfloor \frac{T - j\overline{r_t^{[i]}}}{\overline{r_a^{[i]}}} \right\rfloor} Pr\{X=j\}Pr\{Y=k\} \tag{5.19}$$

$Pr\{D \cap B\}$ represents the cases when either the edge is reported by a managed AP and at least a truthful user, or no managed AP is in range, client reports are adequate to meet the filtering threshold, and there is at least a truthful user among the reporters, and is given by

$$Pr\{D \cap B\} = Pr\{A > 0\}Pr\{B\} + Pr\{A = 0\}Pr\{B \cap C\}, \tag{5.20}$$

where

$$Pr\{B \cap C\} = 1 - Pr\{X = 0\} - \sum_{j=1}^{\left\lfloor \frac{T}{\overline{r_t^{[i]}}} \right\rfloor} \sum_{k=0}^{\left\lfloor \frac{T - j\overline{r_t^{[i]}}}{\overline{r_a^{[i]}}} \right\rfloor} Pr\{X = j\}Pr\{Y = k\}. \tag{5.21}$$

Finally, making use of the distribution of distances between two reported APs (see Eq. (5.8) and (5.9)), the probability that an x-distance edge reported by a truthful user is discovered, and the fact that all Type-1 edges reported by a user are discovered, we can estimate the mean score at round i:

$$\overline{s_t^{[i]}} = L_X(R) + \int_R^{2R} l(x)P_s^{[i]}(x)dx. \tag{5.22}$$

The same analysis can be applied to calculate the mean score of a potential attacker at round i ($\overline{s_a^{[i]}}$).

Reputation updates

Honest reporters From the reputation update rule (5.1), an honest user's reputation at round $i+1$ is a function of his score at round i. With probability p_i, a user is isolated and his reputation is not updated. Therefore, the average reputation of an honest reporter is given by

$$\overline{r_t^{[i+1]}} = p_i\overline{r_t^{[i]}} + (1 - p_i)(\beta\overline{r_t^{[i]}} + (1 - \beta)\overline{s_t^{[i]}}). \tag{5.23}$$

Attackers Users who attack at round i get a zero score and their reputation is discounted. On the other hand, there may be some potential attackers who choose to cooperate, who contribute to the average reputation of potential attackers. The average reputation of potential attackers, also considering the probability that an attacker is isolated, is given by:

$$\overline{r_a^{[i+1]}} = p_i\overline{r_a^{[i]}} + (1 - p_i)\{p_a\beta\overline{r_a^{[i]}} + (1 - p_a)\left[\beta\overline{r_a^{[i]}} + (1 - \beta)\overline{s_a^{[i]}}\right]\}. \tag{5.24}$$

The term $(1-p_a)\left[\beta\overline{r_a^{[i]}} + (1 - \beta)\overline{s_a^{[i]}}\right]$ represents the contribution of potential attackers who cooperate to the average reputation of the total population of potential attackers.

It should be noted that Eq. (5.24) reduces to Eq. (5.23) if we set $p_a = 0$ (i.e., when potential attackers never attack).

5.7 Numerical results

5.7.1 Scenarios and parameter selection

We explore two different cases; first, that of an ISP whose subscribers set up residential Wi-Fi APs connected to their broadband lines and, second, that of a campus Wi-Fi network operator who manages a public WLAN at a university's premises. In the first case, we assume an idealistic scenario where the operator has control over home Wi-Fi routers and, with the appropriate firmware installation, can request for measurements and collect results to centrally plan their operation. This approach represents a best case for a pure AP-centric scheme, since, in practice, due to the heterogeneity of WLAN APs, the fact that APs could be switched off by their owner or their firmware be replaced, the number of centrally-managed APs should be much smaller. The second case is more straightforward; in a campus or enterprise environment, the Wi-Fi installation is typically centrally configured and managed, as is user AAA.

From the work of Jones and Liu [68], it is realistic to assume that the density of APs in metropolitan areas is in the order of a few hundred to a few thousands per km^2. Our own measurements in a densely populated area in the center of Athens, Greece, verify this result.

On the other hand, the value of p_m varies significantly and depends on the size of the (Wireless) ISP or the organization managing the Wi-Fi deployment. The question that naturally arises is what is the ratio of APs that can be managed by the ISP (the p_m value), or, in other words, (given our assumption that the ISP has all registered home Wi-Fi APs of its subscribers under its control) what is the ISP's share in the residential broadband market. From data publicly available [48], as of December 2009, the incumbent operator accounted for 55.3% of the number of DSL lines in Greece, while the rest of the market was shared among alternative providers.

We also measured the ratio of the APs centrally managed by AUEB on one of the Campus' buildings and the average value was typically[7] between 5% and 10%. We measured the number of APs in range indoors and on each measurement, their average number was 6. We also measured the average radius of a few Wi-Fi cells we set up in various parts of the building, which was around $30m$ on average, yielding an estimated $\lambda_{AP} = 2123$ APs/km^2.

Finally, we have set $\beta = 0.2$. Since users start with zero reputation, there is a tradeoff between the importance of a user's reporting history and the speed of building reputation; a large value for β makes the system more robust against attackers who aim at quickly increasing their reputations by reporting honestly, before they switch to an attacking strategy. On the other hand, this makes it slow for honest users to

[7] The number of APs in range varies with time and day and includes other APs operating in offices and labs, and residential or corporate APs set up in nearby buildings.

build their reputation, which in turn causes the system to suffer from low performance for many rounds, before it finally converges to a high topology discovery accuracy. The choice of β is left to the system operator.

5.7.2 Limitations of a pure AP-centric scheme

Based on values we selected for p_m and λ_{AP}, we have numerically evaluated Eq. (5.10) for the AP-centric scheme as client density increases. Figure 5.9 shows the percentage of discovered edges in a topology of 2123 APs/km^2 for the two scenarios we study. Performance drops as the number of clients increases because, as the density of clients grows, more Type-2 edges appear which cannot be discovered if managed APs are not located in the overlapping regions. For increasing client densities, in both cases, we note that a significant share of the Wi-Fi topology is not discovered (around 50% when p_m is high and around 70% for low p_m values). We argue that the second case is expected to appear more often in practice, and show that a client-centric scheme, even in the presence of large numbers of attackers, can do much better.

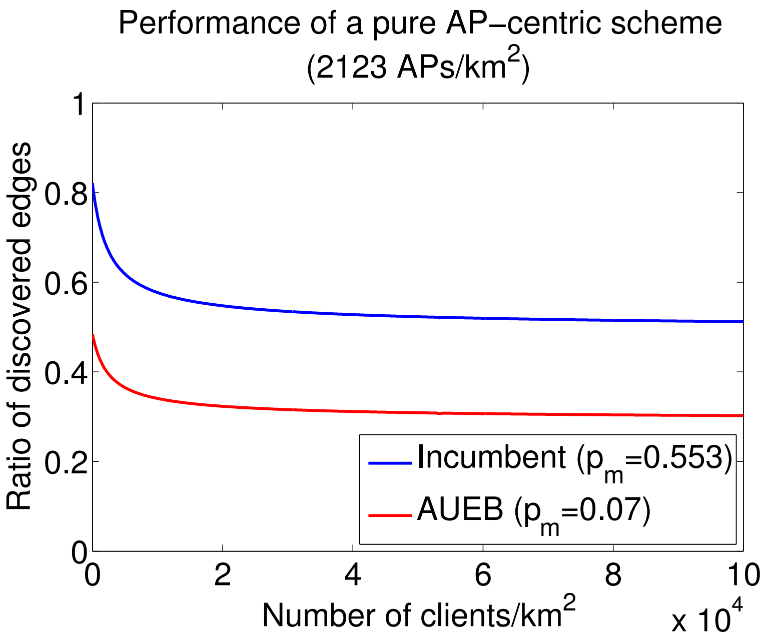

Figure 5.9: Ratio of discovered CG edges in an environment approximating a campus building at the AUEB vs. the case for residential Wi-Fi APs managed by Greece's incumbent ISP.

5.7.3 Performance of a user-centric scheme

The advantages of crowdsourcing

Here we compare the performance of our user-centric scheme to a pure AP-centric one in the two scenarios studied in Section 5.7.2. We fix client density to $\lambda_c = 10000$ clients/km^2, which is realistic for densely populated urban areas. For comparison, and to demonstrate the benefits of a client-driven scheme, we also present the performance of the pure AP-centric scheme. We have selected to experiment with large numbers of attackers; 50% of the total client population are potential attackers and each attacks at a particular round with probability $p_a = 0.9$.

Numerical results from our analytic model shown in Figure 5.11 indicate significantly improved performance for the user-centric scheme, which is more apparent for smaller managed deployments. The mean reputation of honest users converges to a value close to 1, while potential attackers have a mean reputation of less than 0.1. Since our model involves numerical evaluation of integrals, and there is potential loss of accuracy due to rounding in Eq. (5.17) and (5.21), we also programmed a custom simulator in C to validate our analytic results.

After a number of rounds, the performance of a user-centric scheme converges. In such a state, Figure 5.10 shows the relative advantage of a user-centric scheme over an AP-centric one, for the whole range of ratios of honest users and for fixed client and AP densities. Even for large attacker ratios, for realistic managed AP ratios, the user-centric scheme can achieve significant performance improvement. For reasons of clarity, in Figure 5.10, we have set the probability that a potential attacker submits a fake report to $p_a = 1$, i.e., we consider users that either always attack or are always honest. The relative performance improvement is more evident as the ratio of honest users increases.

Dependence on AP density

Here we quantify the effects of the AP density on the performance of the user-centric scheme, for fixed client density. We have also fixed the ratio of honest users to $p_t = 0.5$, the probability to attack to $p_a = 0.9$ and cell radius to $30m$. We present results for the campus Wi-Fi network scenario ($p_m = 0.07$).

Interestingly, Figure 5.12 indicates that the performance of the user-centric scheme, as rounds progress, converges to the same value, irrespective of the AP density. When AP density is high, the user-centric scheme simply converges faster towards this value. On the other hand, as also shown in Section 5.7.3, performance improvement compared to the pure AP-centric scheme is more notable for smaller managed AP densities.

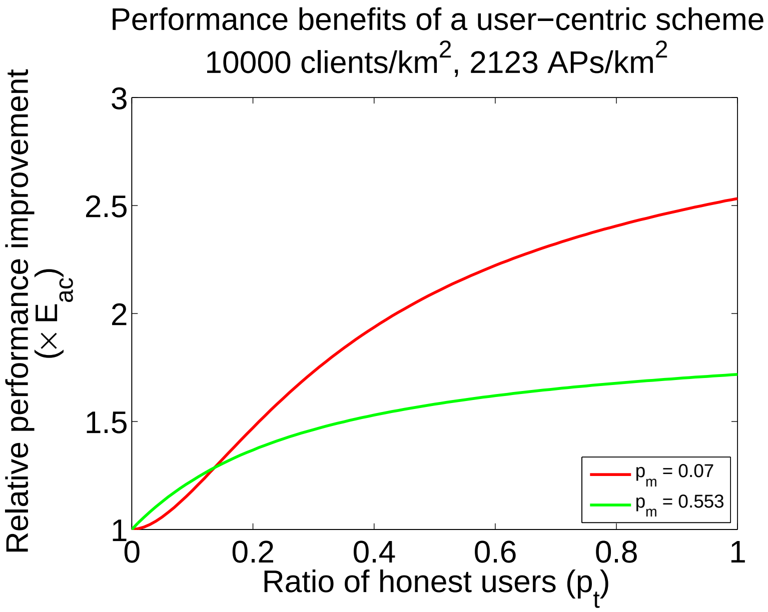

Figure 5.10: Relative performance of a user-centric scheme vs. an AP-centric one, for fixed client and AP densities and for varying ratios of honest users. Each curve shows a different management scenario. At each round, all potential attackers submit fake reports (i.e., $p_a = 1$).

Dependence on client density

In Figure 5.13 we demonstrate the performance of our user-centric scheme for the campus Wi-Fi network scenario ($p_m = 0.07$) for 100 reporting rounds, for the same attack parameters as in Section 5.7.3 and increasing client density. We show that as the density of clients grows, the performance benefits of a user-centric scheme are more apparent (UC curves). Flat lines represent the performance of the AP-centric scheme for each case. The slight decrease in the performance of the AP-centric scheme when client density increases is due to more Type-2 edges missed. In the settings we evaluated, the user-centric scheme achieves a 1.6× to 2.7× performance increase compared to a pure AP-centric one (e.g., when there are 30000 clients/km^2, 87% vs. 32% CG edges are discovered).

For the same managed AP ratio and AP density, Figure 5.14 shows the performance of the user-centric scheme for increasing client densities relative to the density of APs, and after a large number of rounds. Each curve corresponds to a different ratio of honest users; potential attackers always submit fake reports ($p_a = 1$).

For the same scenario, i.e., when $p_m = 0.07$, in Figure 5.15 we quantify the improvement achieved over the AP-centric scheme for varying ratios of honest users. Potential attackers always attack (i.e., $p_a = 1$).

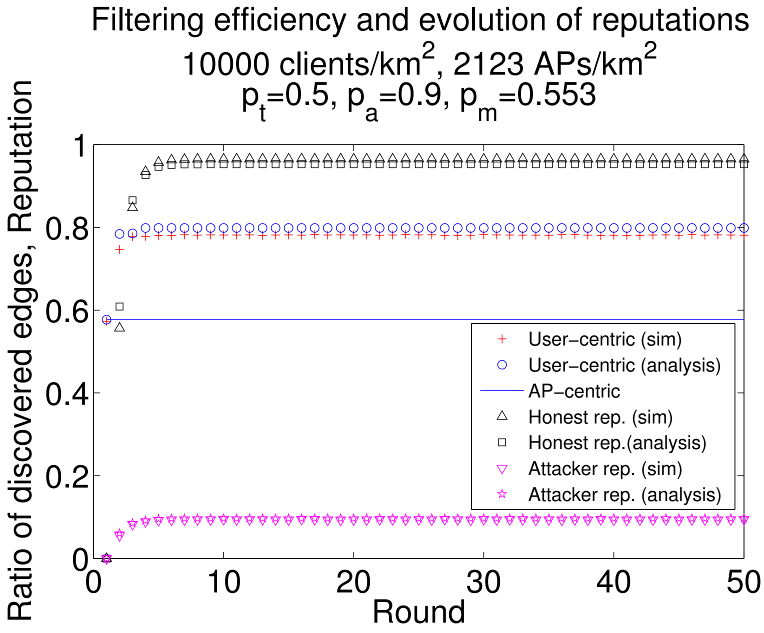

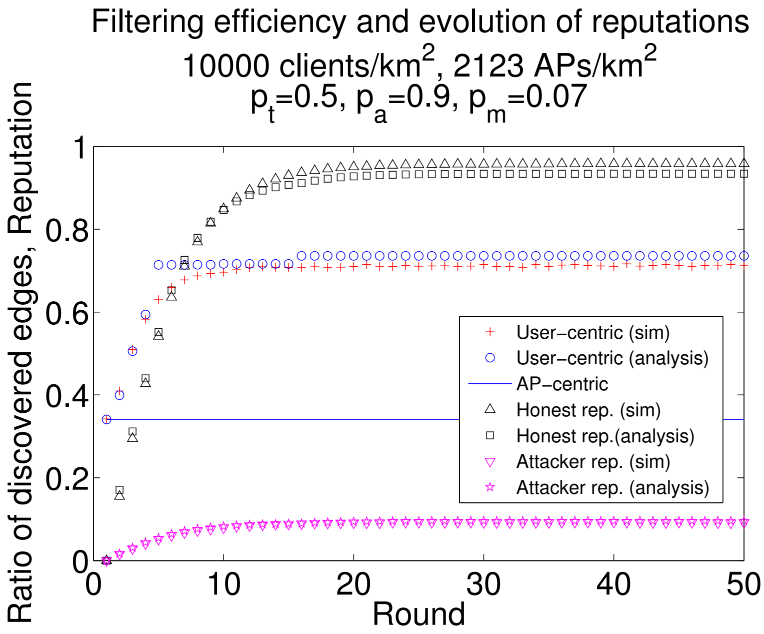

Figure 5.11: Numerical results for scenarios with different managed AP ratio. The first figure shows the case for Greece's incumbent ISP, where 55.3% of the total AP population is centrally managed. The second represents a scenario where 7% of them are managed. The ratio of discovered CG edges (the $\mathcal{E}$ metric), and the evolution of reputations of honest users and potential attackers are shown.

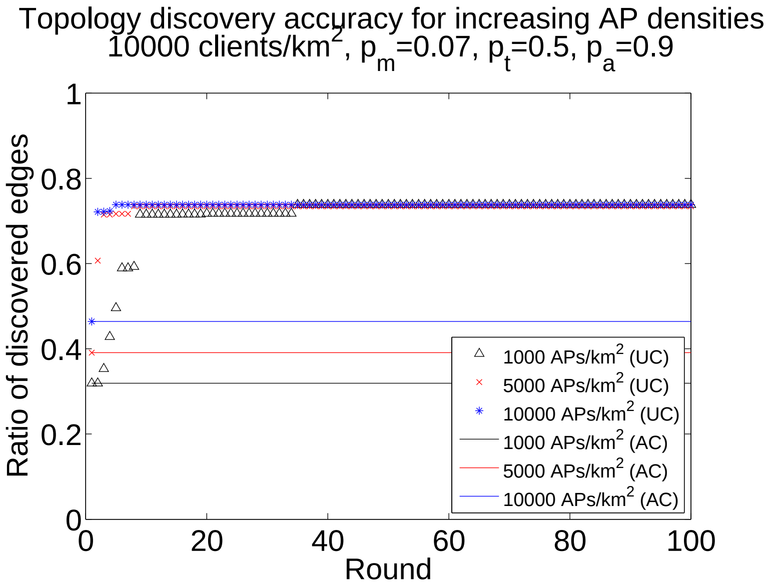

Figure 5.12: Ratio of discovered CG edges for increasing AP densities.

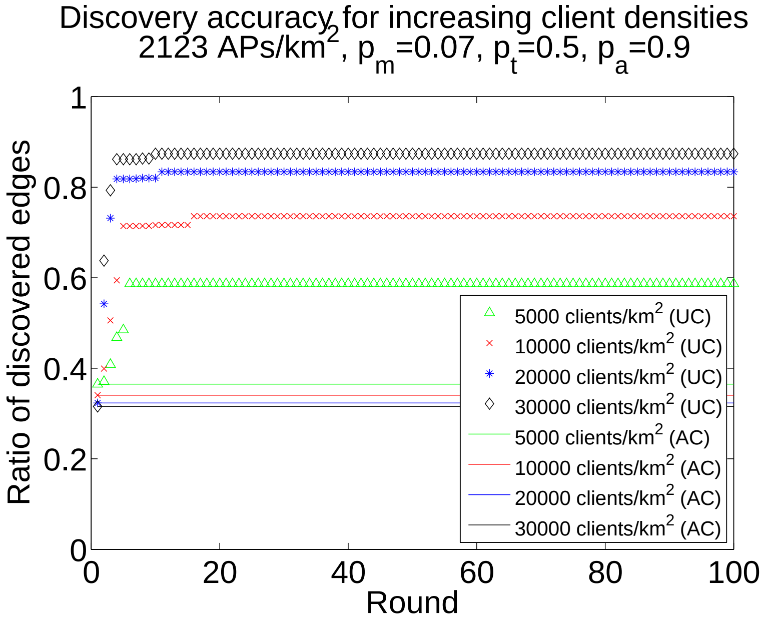

Figure 5.13: Ratio of discovered CG edges for increasing client densities.

Performance for varying ratios of honest users

The performance of our scheme critically depends on the ratio of honest users. For the case where no attackers are present, simplified analytic expressions can be

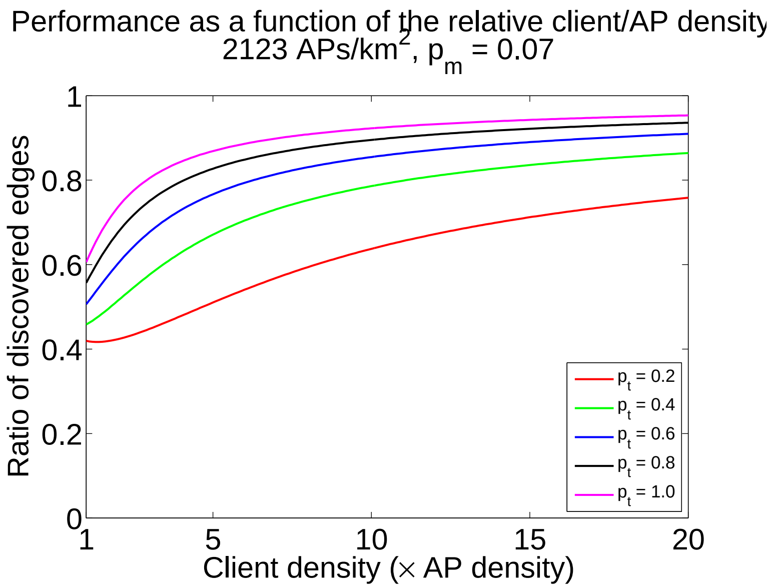

Figure 5.14: Dependence of topology discovery accuracy on the relative density of clients and APs. We assume that $p_m = 0.07$ and $p_a = 1$. Each $x-$axis value is a client density, expressed as a multiple of the density of APs (from 1 to 20 times).

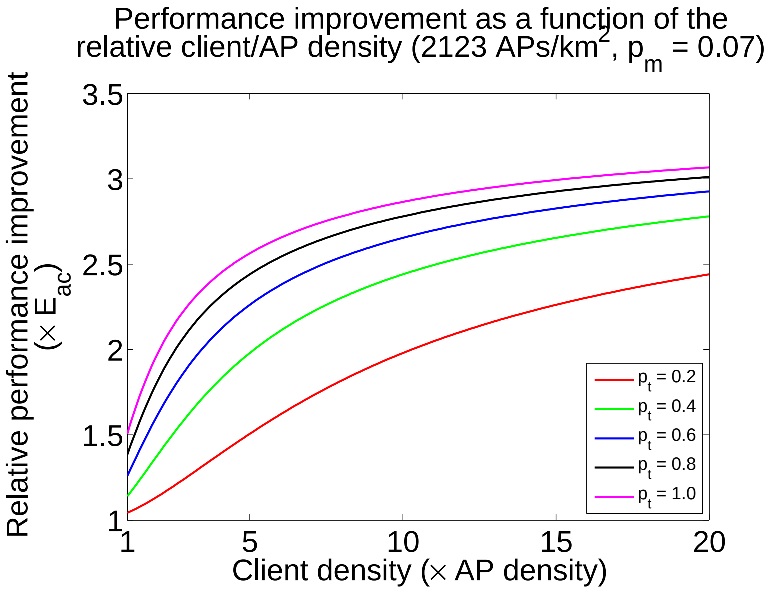

Figure 5.15: Performance improvement of the user-centric scheme vs. the AP-centric one, as a function of the relative density of clients and APs. We assume that $p_m = 0.07$ and $p_a = 1$. Each $x-$axis value is a client density, expressed as a multiple of the density of APs (from 1 to 20 times).

derived. As Figure 5.16 shows, in an idealistic setting where all users are honest, after a number of rounds necessary to build high reputations, most of the CG edges can be discovered.

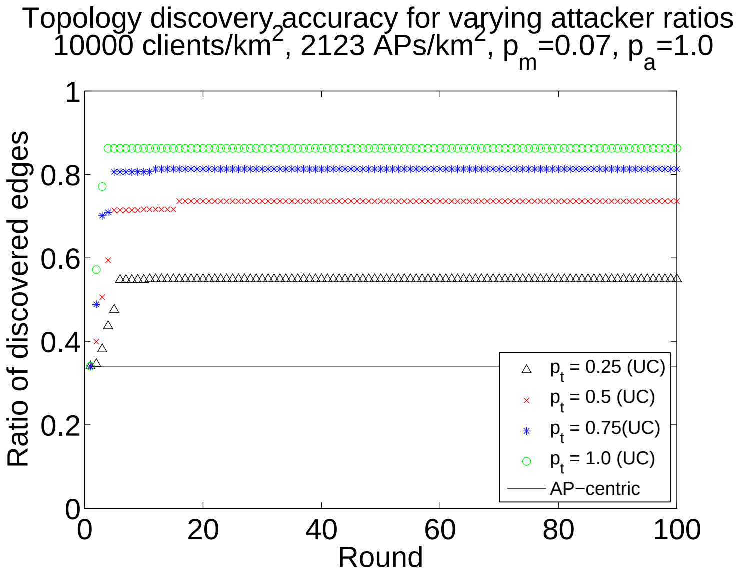

Figure 5.16: Ratio of discovered CG edges for various ratios of honest users.

When the performance of the system has reached a steady state, that is, after a number of rounds has passed such that topology discovery accuracy has stabilized, we present system performance as a function of the ratio of honest users in Figure 5.17. We fix client and AP densities and show one curve for each network management scenario (ranging from a case where 20% of the total AP population is managed to an extreme setting, where all APs are under the operator's control). When $p_t = 0$, the performance of the user-centric scheme is reduced to that of a pure AP-centric one, since no useful information is provided by the users (all users attack).

The role of a user's reporting history

Here we quantify the role of a user's reputation history in the system's performance. The value of the discounting factor in the reputation update rule of Eq. 5.1 is left to the system operator. In principle, a small value for β gives less importance to a user's reporting history. Thus, his score at the particular round has more weight than his history and affects the evolution of his reputation more severely. For instance, a user who is always truthful, builds reputation more quickly. On the other hand, as noted in Section 5.7.1, a small value for β allows potential attackers to quickly build their reputation with few successive honest reporting rounds, before they exploit their increased reputation to make their fake reports more credible.

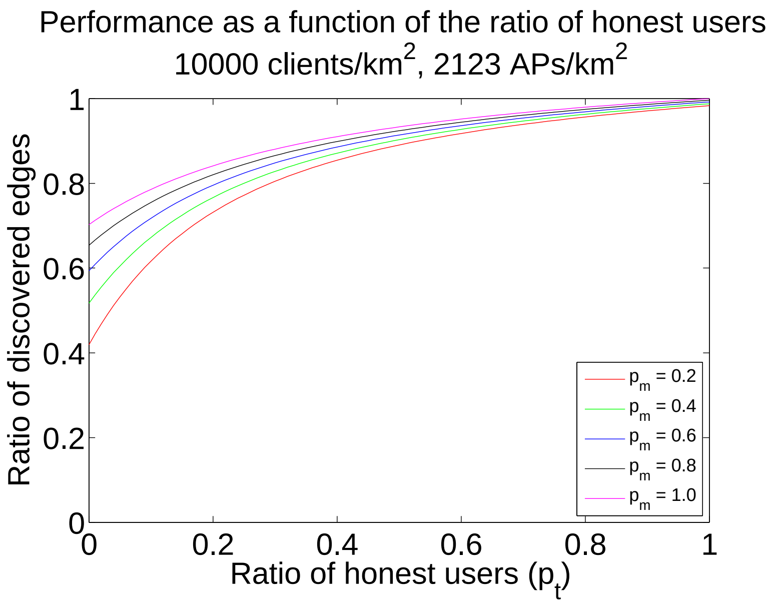

Figure 5.17: Performance of the user-centric scheme as a function of the ratio of honest users. We assume that potential attackers always attack ($p_a = 1$). Each curve shows the percentage of discovered edges for a different managed AP ratio (p_m value).

As shown in Figure 5.18, starting from the lower bound, i.e., the performance of a pure AP-centric scheme, eventually, the system's topology discovery accuracy converges to the same value for different values for β, as is the case for user reputations. We observe that for the first few rounds, there is significant performance advantage when giving less weight to a user's past reputation ($\beta = 0.2$) and valuing more his current score.

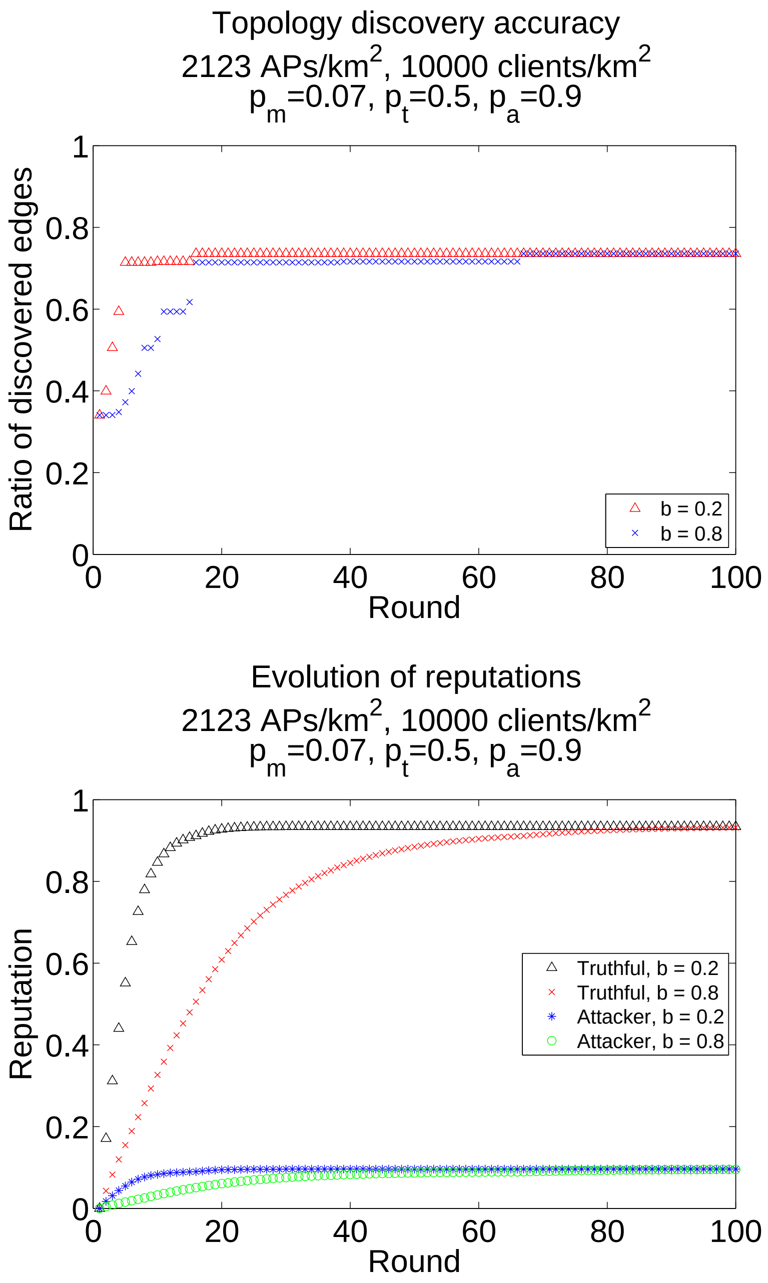

Figure 5.18: The role of reporting history. When the value of β is small, reputations of honest reporters grow quickly, as is the case for the system's discovery accuracy, since we value a user's score more, compared to his reporting history. However, this makes the system more vulnerable to specific types of attacks.

A note on collusion resistance

Here we revisit our basic assumption that attackers act independently and report random fake sets of AP identifiers, focusing on the case of *colluding* attackers, who, aware of the threshold-based filtering mechanism, coordinate to submit the same fake edges to increase their weight and pollute the information stored in the collector's graph. Attackers may report honestly at some rounds, to build some reputation and launch an attack at a subsequent round. The capability of a group of attackers to successfully add fake edges to the graph is limited by their mean reputation. If p_a is high, potential attackers rarely report honestly, which keeps their reputation low. For example, for the campus Wi-Fi scenario we studied, when $p_a = 0.9$, mean attacker reputation is kept as low as 0.0959, which makes it necessary that a colluding group has at least 11 members reporting the same fake edges so that the latter can survive filtering. Since users are authenticated via IEEE 802.11i mechanisms, *Sybil* attacks [27], i.e., attacks in which a user appears with multiple identities, are limited; in our example, a potential attacker would need to have access to 11 different accounts to perform an attack on his own.

However, we have shown that after a few rounds of honest reporting, a user can build a high reputation. He can then switch to an attacking strategy to exploit it by colluding with other attackers reporting the same fake edges. We briefly discuss such sophisticated collusion attacks, which we consider an important topic for further study, in Section 6.3.1.

Non-zero initial reputations

For the scenario we study, initial reputation values higher than zero would help honest users build reputations faster, while after a few rounds, the average reputation of attackers would eventually converge to a low value. High initial reputation values offer the system better performance in the first few rounds, but eventually performance converges to the same value. The system is also more vulnerable to collusion if users do not start with zero reputation. It is up to the system operator to select this value (potentially based on a priori knowledge or expectation about user behavior). Figure 5.19 shows the evolution of system performance for different initial reputation values (0.0 vs. 0.5). The fact that user reputations converge to the same values even when different initial reputations are applied is shown in Figure 5.20.

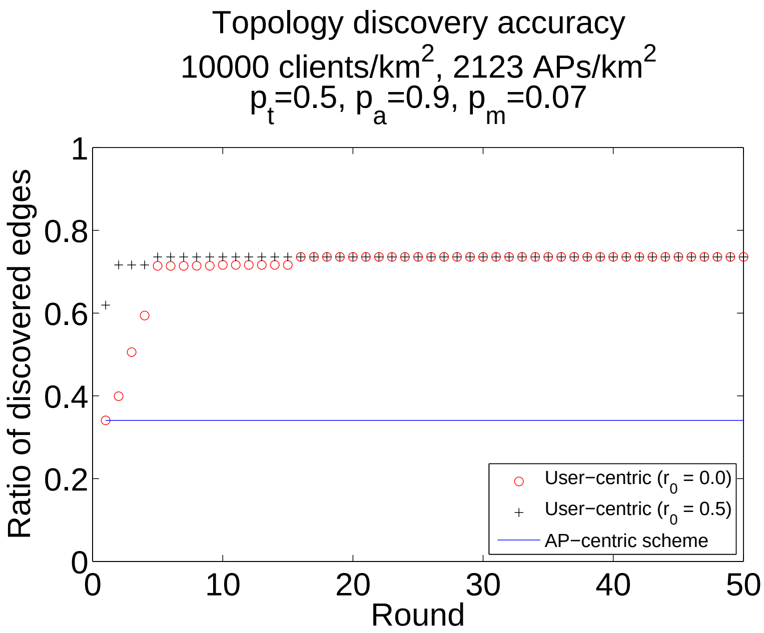

Figure 5.19: Evolution of topology discovery accuracy as rounds progress for different initial reputation values.

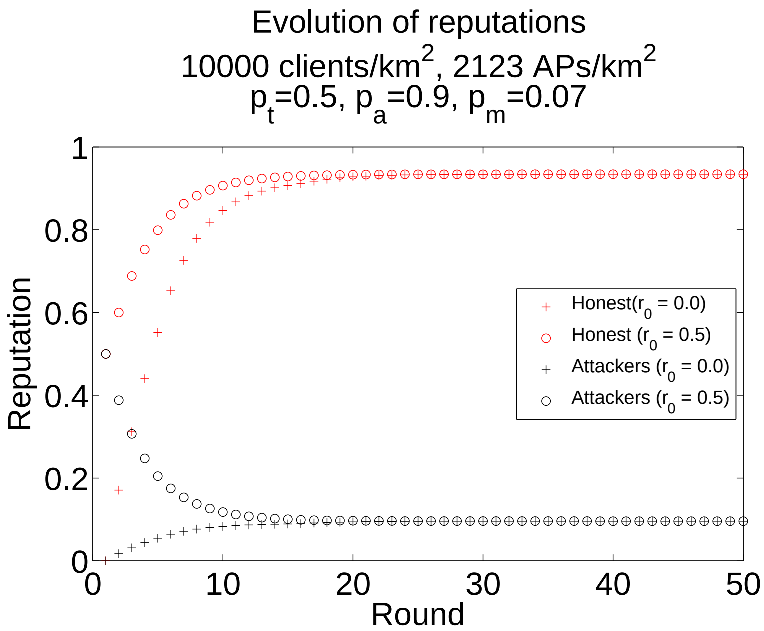

Figure 5.20: Average user reputations converge even for different initial values.

5.8 AP-centric extensions to improve performance

Topology discovery accuracy can be improved at the expense of implementation and hardware complexity or application performance, if we enhance the monitoring functionality of managed APs. There are cases of cell overlap which can be positively discovered even if no managed AP is located there, nor the necessary number of users.

The basic premise of this extension is that APs can monitor transmissions by all clients in range. For an AP with a single wireless interface, this could bring significant performance penalties for user applications, though: The interface should be switched to monitor mode and iterate over all available channels in search of client transmissions to sniff. During this time, packets cannot be sent by the APs, causing potential data loss, disconnections or increased delay for stations. Modern drivers allow for virtual wireless interfaces among which the wireless card is shared. Even if a separate virtual interface is dedicated to channel monitoring, performance penalties persist. A solution would be to use additional dedicated Wi-Fi hardware for monitoring. This, however, entails significant infrastructure cost.

Assuming that the above monitoring functionality is in place, by snooping client traffic APs may be capable of indirectly discovering neighbor APs, by capturing client transmissions towards these APs. This can be accomplished by inspecting frame headers to determine the source and destination MAC addresses, as well as the “To DS” and “From DS” fields, which show whether a frame is sent to or from the AP.

There are still cases where this approach misses cases of cell overlap. First, all clients in the region of overlap could be associated with the monitoring AP. In this case, there is no traffic from/towards the neighbor that the AP would be able to capture. This situation is depicted in Figure 5.21. AP B cannot sense any transmissions towards AP A.

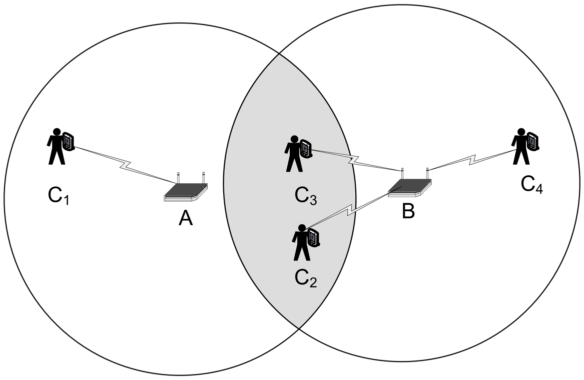

Figure 5.21: Case 1: All clients in the overlapping region are associated with the same managed AP.

Second, links may be asymmetric. For example, a monitoring AP would fail to capture the transmission of a client towards a neighbor AP if the transmission range of the client is sufficiently small. Such an example is depicted in Figure 5.22. Client C_1, probably due to lower transmission power or more severe signal propagation conditions, has a smaller transmission range. While being able to communicate with AP A, it cannot be sensed by AP B. It should be noted that, in practice, link asymmetry is common.

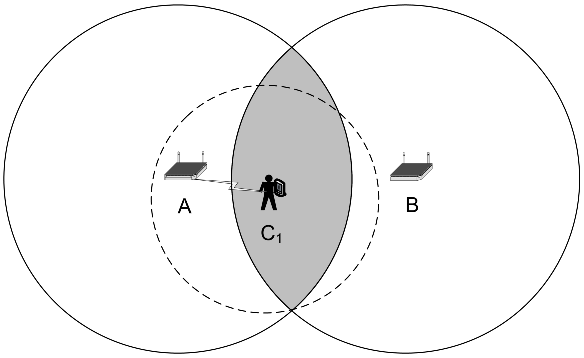

Figure 5.22: Case 2: Asymmetric links.

5.9 Adherence to the principles of user-centrism

In a similar spirit with the discussion of Section 4.3, we conclude this chapter by demonstrating that our crowdsourcing approach to wireless topology discovery respects the principles of user-centrism.

The user at the center Our scheme is user-centric by design. Our information evaluation mechanisms are based on consensus among reporters. We crucially depend on user feedback, which we show to be vital to improve topology discovery accuracy, making user-empowerment evident by showing that an infrastructure-centric approach is significantly outperformed by a user-centric one. The particularities of involving users in the process of collecting information, instead on relying on trustworthy, centrally-controlled infrastructure are carefully addressed.

Open access The architecture we propose in this chapter is not related to network access issues. However, open participation is a desirable feature. The more the participating users, the better the quality of the collected information. On the other hand, users have the freedom to provide feedback or not. Considering their motives to abstain from participation, it is up to the system designer to offer incentives to users to cooperate and contribute truthful information. The reputation-based system that we have proposed works towards this end, and can be utilized to build incentive mechanisms on top of it.

Decentralization and distribution of tasks Although our architecture relies on a centralized infrastructure for user registration, authentication, and reputation management, the crowdsourcing concept which is at the heart of this work is an expression of decentralization. Considering optimizing the operation of a WLAN deployment, we provide a decentralized solution to information collection to support centralized network management tasks. It should be noted that our approach could also be adapted to work in decentralized settings, where the collected information could be made available to a larger set of interested parties who, based on it, could take autonomous decisions; this is something we consider future work.

Low-cost operation We reduce infrastructure cost by delegating the task of collecting information about the radio environment to crowds of roaming users: Instead of using dedicated monitoring infrastructure, the network operator exploits the spectrum sensing capabilities of user devices and takes advantage of the inherent user mobility to more efficiently unveil network topology. We have shown that our architecture can be built on commodity wireless equipment and open-source software. Importantly, we use standards-based technologies and reuse proven security solutions.

Security, trust and user rationality Our work focuses on security aspects of information collection, building on the trust relationship between reporting entities and the system operator. Security threats stem from the fact that users are not trustworthy. By proposing a reputation-based report evaluation scheme and exploiting a set of trusted measurement points, we improve the robustness of decentralized Wi-Fi topology discovery in the presence of reporting attacks. We have not assumed benevolent users. Rather, the latter are rational or even malicious. In the first case, they may misbehave for performance-related or other purposes. (See Section 6.4 for a discussion on the disincentives of users to contribute truthful measurements.) In the second case, they may attempt to disrupt system operation without a particular expected benefit. Our design aims at thwarting attacks by both types of users.

Privacy implications of our approach are yet to tackle. By reporting wireless coverage, approximate user location may be disclosed to the collector. Future research will focus on location privacy issues.

Chapter 6

Design alternatives and open issues

In this section we discuss some aspects of our work which lead to issues that are to be addressed in the future. The points discussed here span across the three research dimensions that we have focused on.

6.1 Developing models for community wireless mesh networks

In Section 3.8 we showed that the frequency of the degrees of nodes in mesh-based WCNs is power-law distributed, based on empirical data. Knowledge that some properties of WCNs could be described accurately with power laws can help build realistic topology generators. Therefore, future work can focus on answering the question on which model is more appropriate to accurately create a network with the desired properties. Two well-known network growth models are due to Barabasi and Albert (BA model) [6], and Fabrikant, Koutsoupias and Papadimitriou (FKP model) [34].

Taking into account the physical constraints that are imposed when setting up links, i.e., the fact that higher quality is achieved for shorter distance links, and there are potential natural obstacles (e.g., hills) that can intervene in the line of sight, our intuition is that a preferential attachment model (BA) would produce less accurate results than the FKP one. In FKP, nodes join the network in discrete time and try to connect to other nodes according to the following criterion:

$$\min_{j<i} \left(\alpha \cdot d_{ij} + ecc\left(j\right)\right) \quad , \tag{6.1}$$

where d_{ij} is the Euclidean distance between nodes i and j, where i is the new node entering the graph. The second term, $ecc(j)$, is the node *eccentricity*, i.e., the hop distance of j from the center of the network (assumed for simplicity to be the first node to join) and α is a weight factor, capturing the relative importance of the two objectives. It is proven [6, 12] that, for all but extremely small and extremely large values of this parameter, the degree sequence of the resulting graph is power law distributed.

An interesting issue to address would be a different interpretation of the eccentricity objective. The center of the network, for instance, instead of being the first node to join, could denote the most powerful node in terms of links and services. For social reasons, these two nodes often coincide; senior nodes are quite often more active in the community.

6.2 Peer-to-peer multimedia services

6.2.1 Evaluation of video communication

With the evolution of video capturing technologies for mobile devices and the proliferation of devices with larger displays and of better quality, it is important to extend our performance evaluation of the secure service architecture we proposed towards video communications. In a similar spirit with the E-model, QoE measurement methodologies and tools for video transmission are available [102]. The main challenge arises with the increased bandwidth demands that video communication entails. However, we argue that by selecting low-overhead security solutions, designing optimizations aiming to shorten the end-to-end communication path[1] and with the advent of more powerful hardware in handheld devices, which will increase the speed of cryptographic operations, the QoE of secure video communication can be enhanced.

6.2.2 Rendezvous and call setup

In this section we study alternative mechanisms to initiate multimedia communication in the context of the architecture we presented in Chapter 4. We assume that the caller and the callee are away from their home networks, visiting foreign (and untrusted) Wi-Fi APs and having set up secure tunnels to their home VPN gateways, through which their traffic is routed towards/from the Internet. We also assume that the IP address of their home VPN gateway is public, be it static or dynamically allocated by the ISP, as is usually the case for ADSL lines. Finally, there are two basic premises in our design, namely that the use of centralized infrastructures should be minimized and that peers perform actions when they have mutual incentives to do so.

[1] In our architecture presented in Chapter 4, both endpoints tunnel their traffic to trusted home gateways, which then decapsulate and forward packets to each other. A selection protocol where one of the two endpoints is picked as the VPN gateway for both users would offer significant performance improvement, given that trust issues and constraints are not an obstacle to this.

Rendezvous between two peers who wish to communicate can be carried out using various mechanisms and protocols and each has advantages and disadvantages in its own right. In any case, users need to discover each other's service parameters, namely the IP address and port where they listen for incoming multimedia traffic and potentially negotiate security parameters to achieve an end-to-end encrypted channel. In our system, we have opted for call initiation based on the exchange of GSM SMS messages, on which we comment first. Then, we discuss the role of DNS in peer discovery and present two solutions based on the Session Initiation Protocol (SIP) [105].

Use of a GSM SMS

In Chapter 4 we proposed a rendezvous scheme that was based on the exchange of a simple GSM SMS to setup a multimedia call. The SMS includes the IP address and port on which the user is reachable, i.e., the IP address of his home VPN gateway. Upon reception, the software agent residing in the callee's mobile device will reply with a voice stream which will be routed to the caller's VPN gateway via the callee's home network, and eventually data will reach the caller, after forwarding state has been set up at the two home gateways. We adopted this approach for its simplicity and for being in line with the assumption that the caller knows only the mobile phone number of the callee, which closely follows the paradigm of current GSM network use. We still resort to a centralized infrastructure (GSM network), but this is only for call setup (after all, typically, mobile device users already own GSM subscriptions; our approach does not mandate using a separate system, such as DNS or SIP registrars for rendezvous).

In Section 6.2.4 we comment on a potential attack by the GSM operator and a means to counter it and achieving end-to-end security without resorting to a PKI.

Dynamic DNS

In a DNS-based alternative, instead of initiating a call using the callee's mobile phone number, if we assume that a user's home IP address can be resolved from a DNS name, a caller can directly stream his data to his peer's home. The procedure then follows the same way. It should be noticed that since VPN gateways would normally reside in users' home networks, typically accessible via a dynamic IP address, user equipment should dynamically update the name-IP address binding. Many dynamic DNS services exist [90, 79] and most home network equipment have the capability to perform such updates built into their firmware.

Relying on DNS hides one vulnerability similar to when using GSM SMS text messages, since it is still implied that (i) users trust the Dynamic DNS name service, and (ii) name servers that peers use to resolve each other's host names are trusted. These conditions have to do with how calls are set up: The caller tunnels a request to resolve the callee's name to his home name server through his VPN gateway. If this name server is under the control of an adversary, it is possible that the callee's name resolves to a host controlled by the adversary, who could then easily launch a man-in-the-middle attack. An end to end security scheme where users authenticate each other is necessary to combat such threats.

Using the Session Initiation Protocol (SIP)

A multimedia call could also be established using the *Session Initiation Protocol (SIP)* [105]. Peer discovery, however, is again an important issue. Users are identified using SIP URIs, which mandate the existence of SIP registrars where users are listed and, typically, SIP proxies in each user's domain. In the decentralized scenario we study, and given the assumption that cooperation between two peers happens only when there are mutual incentives to do so, and interaction cannot be assumed for purposes other than requesting Internet access when visiting a foreign AP or setting up a call, this would mean that each peer should maintain its own SIP server. Typically, the caller communicates with a SIP proxy in its domain, which uses DNS to resolve a SIP proxy in the callee's domain. DNS procedures can also be used to discover service parameters (such as the transport protocol to use or the port the other server listens to). Locating SIP servers is specified in RFC 3263 [104].

Applying a SIP-based solution has some advantages: First, it is a standards compliant approach, with SIP being heavily used and tested. Also, many SIP phones or implementations of SIP user agent software are readily available. On the downside, the need for every peer to operate SIP servers increases management effort and home gateway complexity, while it is still necessary to rely on DNS.

Peer-to-Peer SIP

The problem of locating call endpoints in a peer-to-peer manner is being studied within the P2PSIP IETF Working Group [94]. For reasons of scalability, resource limitations that do not allow for the installation of SIP servers, but also for reasons of trust and reliability, e.g., when an organization or individual does not want to rely on external centralized service infrastructure, a distributed alternative to SIP is being standardized. The core concept is to rely on a P2P network (a Distributed Hash Table (DHT) such as Chord [118]) built on the equipment of the users of the system to distribute the core functionality of SIP, such as user registration, proxying, and locating users. The REsource LOcation And Discovery (RELOAD) protocol has been specified by the WG for signaling in a P2PSIP network. Cirani et al. [21], on the other hand, have proposed a discovery architecture similar in nature, but without utilizing the RELOAD protocol. Again, instead of using SIP servers to locate call endpoints, a DHT is used as the IP address/port lookup structure.

Interestingly, the security aspects of P2PSIP architectures have recently received attention. For a thorough review of relevant issues, the reader is referred to the work of Tuceda et al. [120].

Applying the above approaches on the user-centric wireless networking scenario we study implies that either a network of special nodes forms the DHT, or, some of the home gateways join it, which in turn implies cooperation between community members. While in some wireless communities this could be possible, in our case, it violates the assumption that peers do not engage in any cooperative activities if it is not to their interest; clearly, in order to build a reliable DHT on top of residential Wi-Fi equipment belonging to P2PWNC peers, cooperation is necessary on the peers' behalf. A peer would have to maintain a set of records and respond to lookup requests for multimedia calls he is not involved with, for which a mechanism to stimulate cooperation at the DHT level is necessary. An interesting direction could be to extend the P2PNNC accounting and incentive mechanisms to incorporate DHT storage and lookup operations.

6.2.3 Dealing with highly-mobile users

We have thus far considered low-mobility, nomadic users. Achieving efficient handoffs for highly-mobile users in the peer-to-peer Wi-Fi sharing scheme of Chapter 4, while maintaining ongoing multimedia sessions, requires addressing problems hard to tackle. First, when a user moves from one micro-provider to the other, and not taking into account the time required to complete the association with the new Wi-Fi AP at the MAC level and the time required to request and be assigned with a new local IP address by the visited peer's DHCP server, he needs to re-negotiate Internet access. Cooperation between the two visited peers to perform the handoff faster is ruled out by design. (A visited peer does not have a clear incentive to assist the user in getting Internet access from a neighbor peer, as the user moves.) Instead, the user will have to prove his contribution to the new prospective provider, first by providing him with a subset of his receipt repository (*gossiping* step), and then by the latter executing the reciprocity algorithm. If access is granted, the tunnel towards the user's home gateway should be re-established. A user's ongoing VoIP and multimedia sessions are expected to suffer significant packet loss during the hand-off process. Quantifying this overhead and improving handoff performance in such settings are significant issues for future research, where potential solutions could involve mechanisms for mapping P2PWNC APs, predicting a peer's path and proactively initiating P2PWNC sessions over the Internet with prospective providers nearby.

6.2.4 End-to-end security and avoiding man-in-the-middle attacks

Providing end-to-end security in a peer-to-peer manner in the multimedia communications architecture that we have proposed is yet to be tackled. One would notice that although the two endpoints of the multimedia call set up secure VPN connections with their trusted home gateways, the path between the two gateways is still unsecured.

If we assume that the IP address of the home gateway of the caller has been communicated to the callee via a GSM SMS (see Section 4.2), it is straightforward for the GSM operator to sniff on the multimedia call, performing a simple man-in-the-middle (MITM) attack: It can modify the contents of the SMS pointing to a gateway of their own. The callee responds with the multimedia stream that comes unsecured from the VPN gateway of the callee to the gateway that belongs to the GSM operator. Then, the operator forwards the traffic to the caller and none of the call endpoints is aware that their traffic is being intercepted.

To combat such an attack and achieve end-to-end secure and private peer-to-peer multimedia communication, the unprotected path of the call has to be secured. A simple means is that of conveying the caller's public key to the callee during the call setup phase (GSM SMS exchange), and using it to exchange a shared key for traffic encryption. However, the callee has to verify that the public key of the caller has not been changed by an adversary (e.g., the GSM operator) performing a MITM attack. Thus, after communication has been set up and assuming a voice or video call, the two parties can use some form of vocal acknowledgment to verify that the public key exchanged is the appropriate. This way, an end-to-end secure VoIP or video call can be set up without resorting to trusted certification authorities.

Methods and systems that implement this type of in-band key exchange have been proposed in the literature. The ZRTP [129] protocol uses public key cryptography without resorting to a centralized PKI and, thus, without the need for certificates. During call setup, the two parties perform Diffie-Hellman key exchange in the media path and set up a shared secret. Then, a *Short Authentication String (SAS)* based on the shared secret is derived, which appears on the displays of the participants' devices. To ensure that a man-in-the-middle attack is not being performed, both users should confirm the same SAS value.

All ZRTP messages are multiplexed with the media stream, making it independent of the signaling protocol used for call setup (e.g., SIP). The fact that it is purely peer-to-peer makes is suitable for application to our architecture. Vocal verification of a key exchanged between the two parties engaging in a phone conversation is a concept presented as early as 1996 [128].

One should note the performance implications of applying such end-to-end security schemes to our service design. While the two peers maintain VPN tunnels with their home networks to protect themselves from untrusted visited APs and suffering the overhead of tunneling, they also have to encrypt their data end-to-end before tunneling them home. Therefore, there is redundant use of security mechanisms. Along the mobile user-home gateway path data are encrypted twice: Application data are first encrypted with the shared key established between users (call endpoints) when the media session was initiated. Then, data are encrypted again and tunneled home, where they are decapsulated, sent to the peer's home gateway, encrypted and tunneled to the peer, where they are decapsulated. The resulting data are eventually delivered to the application, where they are decrypted using the shared key. It is important that future research efforts emphasize on removing redundancy across security layers, and optimizing the end-to-end media delivery path.

6.3 Attacks on topology discovery schemes

In Chapter 5 we studied simple attacks on the crowdsourced topology discovery process. Here, we discuss some more sophisticated attacks, their importance, and potential methods to counter them. Note that the list of potential attacks can become lengthy; new countermeasures put in effect may bring up new strategies to evade them. The following discussion includes some relatively straightforward cases of misbehavior that we have not considered in our analysis.

6.3.1 Strategy changes

Reputation-based schemes face the problem of strategy changes; a well-behaving user can build reputation and exploit it to perform attacks. In our case, a user may report honestly for a number of rounds and then switch to an attacking strategy. As described in Chapter 5, in many cases (e.g., when the deployment of managed AP is dense and a user's reporting history does not have much weight compared to a user's current reporting score), it does not take long before the reputation of an honest user converges to a high value.

Note that in the scenarios we studied, attackers can only exploit this by colluding. (Otherwise, their reports are filtered by default.) They can then perform targeted attacks by reporting the same fake edges. In turn, these very edges help them achieve a high score and thus keep increasing their reputations, while affecting the collector's view of the topology. This poses an important threat to our system and calls for sophisticated countermeasures. What magnifies this problem is the fact that if the reputation of attackers is high, only few of them are enough to place a fake edge in the filtered graph.

One potential countermeasure would be to take into account the identity of the AP contained in a fake report, as well as the number of occurences of this identifier across all user reports: The fake AP identifier which would be connected to real ones via fake edges would appear only in the reports of colluders. Things are different in case colluders attempt to place a fake edge between two existing managed APs (see also Section 6.3.2). If location information about managed APs is available, the fake edge could be filtered, if the two APs are very far from each other. Otherwise, i.e., if the two APs are located such that coverage overlap is possible, the reports of nearby users (e.g., users associated with either of APs in question) or managed APs could be utilized: absence of a report about the fake edge by them is an indication of collusion.

Given the severity of colluding attacks, the reputations of detected colluders should be discounted in a stricter manner.

6.3.2 Fake reports including existing AP identifiers

Thus far, we have considered attackers who report random fake AP identifiers, even in the case when attackers collude. Here we revisit this assumption. In the case of attackers who act independently, i.e., they submit fake edges between existing APs, the attack is equivalent to the one addressed in Chapter 5. Such edges will be filtered by definition, since their weight will always be below the threshold.

In the case of colluders, this attack becomes more significant. If there is a sufficient number of colluders such that a fake edge eventually survives filtering, an edge can join two existing vertices, and this can have more severe impact (affecting more network nodes) to any mechanisms relying on topology information.

Again, such attacks can often be tackled by exploiting location information. For example, an organization deploying a user-centric topology discovery scheme may have knowledge of the actual locations of its APs and an edge between two such APs which are located far from each other may be directly filtered from the CG, due to physical limitations in signal propagation.

6.3.3 Partial attacks

Here we study the case where attackers keep a mixed strategy: Their reports include some fake AP identifiers, as well as the identifiers of some true Wi-Fi cells in range. The observations of Section 5.4 ensure that all random fake edges will be removed (or, in the case of collusion, for a fake reported edge to survive filtering, a certain number of reports by colluders is necessary, based on attacker reputations). Existing edges will appear in the CG if their weight is above the threshold. Therefore, from an efficiency point of view (measured by the number of discovered edges), this attack is of less significance than the case when no true edges are reported. The downside is that true edges help an attacker sustain a higher reputation, which can be exploited in the case of collusion.

6.3.4 Refusing to report

Counting the number of discovered edges, refusing to report is equivalent to reporting fake edges only. However, in our design, this does not lead to reputation decrease, since when a user delivers no reports he is considered *isolated* and his reputation remains unchanged.

Refusal to report can be performed either by sending an IEEE 802.11k beacon report with the “refused” or “incapable” bits set in the “mode” field (see Table 5.2), by simply not replying to a beacon request or by sending a normal beacon report including only the AP he is associated with. The fact that we do not update the reputations of users who have a single managed AP in range may offer incentives not to report. For example, a user who is a good reporter consistently, after a number of rounds, will develop a good reputation. He can then refrain from submitting reports without suffering a reputation reduction and, in the eyes of the collecting system, he will still look like an honest (but isolated) reporter.

Technical means to counter this behavior are available. In the first two cases, the AP can notify the collector about the client's failure to report. In the third case, where the client is considered isolated, as rounds progress, the collector can identify cases where a user consistently abstains from the reporting process. Also, reports by neighbor nodes can be utilized to get a clearer view of the conditions in the user's vicinity. A user who implements this strategy can be suspected if the majority of users associated with the same AP as him submit (similar) reports while the user reports no edges. Then, the collector can apply appropriate measures, e.g., by appropriately discounting their reputation. Reconsidering our rule that reputations of isolated users are not updated, an approach where user reputations are discounted with time (i.e., rounds) is also to be studied.

6.3.5 Compromized APs

In our work, reports by APs are crucial for system bootstrap and for increasing topology discovery accuracy. APs are trusted by default, an assumption that we have not yet discussed. In a centralized setting, where an organization deploys a Wi-Fi network and has all the APs under its full control, it is realistic to assume that such an assumption is valid (unless an AP is compromised, which would potentially require physical access and a firmware upgrade—such cases our outside the scope of this discussion).

However, when dealing with potentially competing service providers, their trustworthiness is debatable; incentives could be such that an AP participating in the topology discovery process would submit fraudulent information, if it was to its owner's interest, just as a normal user could do. In such cases, alternative measures should be put in effect. For example, the APs could also be associated with reputation values, or a set of trusted roaming users could assume the role of cross-checking AP-based measurements. Punishment could be applied as a means of enforcing truthful behavior, which could be implemented in various ways (e.g., excluding a misbehaving operator from the confederation).

6.4 Incentives for trustworthy reporting

Although we considered attacks towards the crowdsourcing mechanism for collecting Wi-Fi topology information, we have yet to address the motives of users to perform attacks in the first place, as well as provide incentives for truthful reporting.

We identify two basic families of incentives for users to submit fraudulent feedback, i.e., *performance-oriented* incentives and *economics-oriented* ones. Here, we consider a more generic reporting architecture, without focusing only on centrally-managed Wi-Fi deployments. This could be the case, for instance, for a Wi-Fi sharing community of users or a coalition of commercial WISPs who agree upon using a central service to maintain Wi-Fi topology and make the coverage map available to registered users, who can then run distributed reconfiguration schemes outside the control of a central entity.

Performance-oriented incentives are associated with the overhead of spectrum monitoring. While scanning for Wi-Fi presence is a standard and relatively low-cost operation when performed infrequently, and since clients can passively record Wi-Fi beacons, more advanced spectrum monitoring schemes, which a more sophisticated scheme would necessitate, could reduce user QoE. It should be noted that research in the area of Distributed Spectrum Sensing in Cognitive Radio Networking environments has identified such performance implications, as well as the tradeoff between accuracy of measurements and user-perceived performance [5].

Even in the case of requesting a client to perform an active scan, from a set of experiments that we have carried out [39], we have shown that frequent requests can disrupt voice conversations that the user has in progress. In particular, using the same codec settings (G.729a) and evaluation methodology (E-model) as in Chapter 4, we emulated a VoIP call over a Wi-Fi link, at the same time requesting the client to perform an active scan. We experimented with varying scanning frequencies and discovered that even when no other traffic is present, if a report is requested frequently, QoE can drop below an acceptable level. While, in practice, reporting frequency is expected to be relatively low, instantaneous quality degradation is possible if, for each report, the client is requested to perform repeated measurements for improved confidence[2]. Figure 6.1 shows the results of our experiments. For scanning frequencies of a few seconds, the *R-score* for the voice conversation drops below the acceptance threshold. The reason is that an active scan takes 250 ms on average, during which packet transmission and reception are impossible.

From the above discussion, it appears that a user who is not willing to suffer potential performance effects may wish not to join the reporting process. Sending random fake reports is one potential strategy that such a user could follow.

Next, we address the other type of incentives. Consider the following scenario: There is a confederation of WISPs and each user is affiliated with one of them. This affiliation may be in the form of a paid subscription, but the confederation allows roaming across providers. A central entity is responsible for maintaining wireless coverage information, built by reports from roaming users. Since this information

[2] The IEEE 802.11k standard allows for repeated measurements; see Section 5.5.4, and, in particular, the *repetitions* field of a beacon request frame.

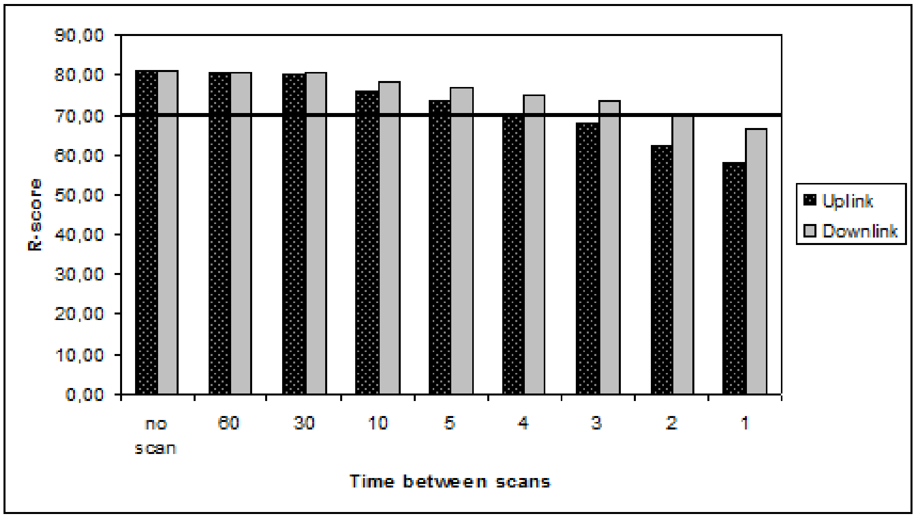

Figure 6.1: The effects of active scanning on VoIP performance

can be used for frequency assignment or transmit power control, a group of users affiliated with provider A can collude and submit the same fake reports, making an AP belonging to provider B appear to cause much interference. Then, an overlaid power control scheme could dictate this AP to transmit at lower power, thus limiting its coverage area and causing it to fail to serve some users, who would then have to associate with another AP in the area. Another example would be to cause B to assign a channel to its AP such that performance would be suboptimal for users associated with it. In both cases, driven by competition, provider A attempts to either cause dissatisfaction to B's clients indirectly, or manipulate the spectrum sharing schemes that are based on topology information to its advantage, always with the aim of maximizing its own utility at the expense of its competitors.

The reputation-based scheme we proposed in Chapter 5 helps limit some such attacks; consistent attackers have low reputation values and a large colluding group may be necessary in order to trick the threshold-based filtering mechanism. However, even applying such countermeasures, more sophisticated forms of collusion are possible. Therefore, an important future research step is to provide incentives to users not to attack in the first place. Such an approach could involve linking a user's reputation with the service quality offered. Below we propose a simple method to achieve it.

The basic requirements are that APs are aware of user reputations and that they can perform bandwidth sharing. The former is implied in the architecture that we have described in Chapter 5, but is also fairly straightforward to assume for other types of network structures, even in decentralized settings where each AP may have a different view of the reputations of visitors, as is the case for the P2PWNC Wi-Fi sharing scheme. The latter assumption, i.e., bandwidth sharing capability, can be achieved by appropriately limiting egress traffic in the AP's interface towards the

Internet and limiting the egress traffic in the wireless interface. Many APs today are Linux-powered. The `tc` tool [76] can be used to implement such traffic shaping. It should be noticed that uplink traffic at the wireless link cannot be controlled; malicious users can jam the channel by sending packets at a high rate, even disrespecting the IEEE 802.11 Distributed Coordination Function (DFC). Such behavior is outside the scope of this work. Countering jamming attacks is an issue widely addressed in the literature [101].

Bandwidth rewards can then be offered to users with high reputations. The AP operator can partition its bandwidth in a way that a fixed portion of it is shared evenly among users and the rest is distributed to them as a bonus for their truthful feedback in a weighted manner, according to their reputation. Each user associated with an AP, receives the following amount of bandwidth as a bonus:

$$B_i^{[k]} = \frac{r_i^{[k]}}{\epsilon + \sum_{j=1}^{n} r_j^{[k]}} B_b, \tag{6.2}$$

where $r_i^{[k]}$ is the reputation of the i-th user associated with the AP at round k, B_b is the total bonus bandwidth (in bits per second) and $0 < \epsilon << 1$ is used to avoid a zero denominator in case all users have zero reputation.

Depending on the particular network setting, the network operator, the AP owner, or any other entity responsible for network planning and operation, can decide on the ratio of the bandwidth that will be allocated as a bonus, which depends on how much he values truthful topology information. It should also be noted that $B_b^{[i]}$ values are adjusted when necessary, that is, when a visitor joins or leaves the network and at the end of each reporting round, when reputations are recomputed.

Other forms of reward or punishment are, of course, possible in order to motivate users towards truthful behavior. Any such approach should be carefully studied from various perspectives, such as their implementation simplicity, susceptibility to attacks and incentive-compatibility, an aspect which can be evaluated using tools from game theory.

6.5 Alternative uses of topology information

The topology model we have applied is suitable for use with specific types of channel assignment schemes, namely those that operate on a weighted undirected graph with vertices denoting Wi-Fi APs and edges denoting conflict. We have shown how to collect such information using IEEE 802.11k frames. Different uses are, however, possible, which may require model changes. In this section we discuss the potential for different applications of topology information, as well as the modifications to the model necessary to accommodate these applications, and new security threats that may emerge.

6.5.1 Transmit power control

Schemes that aim to optimize Wi-Fi performance by means of transmit power control could operate on the Coverage Graph. However, our model should be augmented so that it can encode information to more accurately quantify the amount of interference suffered by users located in regions of overlapping cell coverage. This information can be conveyed by clients using the RCPI and RSNI fields of beacon reports. Appropriate edge weight functions should be designed to simultaneously capture the reputations of users reporting an edge and the cumulative amount of interference caused to them. It is important to note that such an approach requires carefully tackling another potential attack, which we have not addressed: Along with reporting fake BSS identifiers, users can also falsify received power measurements. This calls for sophisticated schemes to quantify interference in the presence of such attackers, to which end research that is active in other areas (e.g., distributed spectrum sensing in Cognitive Radio Networks with fraudulent reporters) could appear relevant.

6.5.2 Network planning

In our work thus far we have not considered the time dimension. Namely, on each reporting round, we derive a new snapshot of the network topology, which can be used for reconfiguring operating parameters of the network (e.g., frequency), but will only be valid until the next reporting round. Generating statistics about network coverage and user presence over long periods of time, for different hours of the day, or days of the week is, however, important for network operators, since it can reveal trends in user behavior, as well as locations which are over- or under-provisioned. It is rather straightforward to extend our system so that the time dimension is taken into account, and it is then a matter for network operators to exploit the available information for network planning purposes.

6.5.3 Handover planning

Wi-Fi topology knowledge could prove important for reducing hand-off latency. When mobile, a user will transit from the AP he is associated with to an AP nearby. A system to prepare user handovers is possible to built: As soon as a user associates with an AP, the collector can notify registered nearby APs (which may or may not be in range with the user) that the user may soon attach to them. This can help reduce the time required to associate with the new AP in the event of user mobility. Note that the IEEE 802.11r [57] amendment deals with fast BSS transition mechanisms, specifying means to build and exchange topology information among neighbor APs and stations and techniques to reduce reassociation times. IEEE 802.11r assumes that topology discovery is carried out by stations either by active scanning or by requesting

IEEE 802.11k neighbor reports by the APs they are associated with. Obviously, if a central collector makes network-wide topology information available, a station would be able to map its neighborhood with more accuracy. In any case, potential security threats and their effects on fast handover performance are an issue for further study.

Chapter 7

Conclusion

This work approached wireless networking from a user-centric viewpoint, given recent technological advances of the last decades that made wireless equipment ubiquitous and enabled connectivity at low cost, as well as socioeconomic factors, that have lead to disruptive user empowerment. User-centric wireless networking is a broad research area, lying at the intersection of an interdisciplinary set of topics. Our study touches various aspects of wireless networking, including infrastructure, service and information provision, bringing out the central role that the user can play towards achieving ubiquitous, secure and optimized communication at low cost.

We studied issues pertaining to the emergence and operation of wireless community networks, a prominent case of autonomous communication built on user-provided network infrastructure. We provided insight on the social mechanisms and incentives that have led to the evolution and sustainability of WCNs and, based on a study of large operational WCNs, observed that some of their structural properties obey power laws; this observation is the first step towards deriving realistic WCN models and tools for the performance evaluation of services and protocols on top of them.

Wi-Fi sharing communities are a specific case of WCNs, with the distinct characteristic that community members do not necessarily know and trust each other. Rather, cooperation can be enforced by applying reciprocity-based service exchange mechanisms and without requiring centralized authorities to tackle registration and accounting. We proposed service architectures tailored to such communities, where secure communication is granted and user privacy is enhanced. Our particular focus was on the performance of peer-to-peer voice services, taking into account that they are designed to operate on low-cost home Wi-Fi equipment and resource-constrained mobile devices, while the decentralized nature of the underlying Wi-Fi sharing scheme should be respected. Our experimental evaluation revealed that even when high-overhead security mechanisms are in place, a few VoIP streams of acceptable quality can be sustained over commodity wireless equipment, a conclusion that supports our claim that, if Wi-Fi coverage is adequate, our solution can offer a secure, low-cost alternative to 2G/3G cellular services.

On the other hand, a basic premise of the user-centric networking scenario that we envisage is its operation in unlicensed spectrum, where anyone can become a micro-operator. A challenging aspect of this scenario is how to utilize the scarce spectrum efficiently, in view of interference. Many coexistence mechanisms and protocols have been proposed to optimize operation in unlicensed spectrum, aiming at efficient sharing across the space (power control, use of directional antennae), frequency (channel assignment), or time (TDMA schemes). These mechanisms heavily rely on accurate knowledge of the network topology and radio environment to operate. In this work, we focused on how to build such information in a robust, user-centric manner: Instead of relying on the infrastructure, we proposed that the task of monitoring and providing feedback about wireless topology and potential interference be crowdsourced to roaming users. We did not assume that users are trustworthy, though, and proposed reputation-based mechanisms to deal with fraudulent reporting. We implemented our crowdsourcing approach making use of the IEEE 802.11k standard for collecting topology information, for which we appropriately extended the functionality of the current Linux wireless networking stack. We derived analytic expressions on the performance of our scheme and have shown it to significantly outperform infrastructure-centric schemes, even in the presence of large numbers of users who follow specific attacker strategies. Research in the direction of combating more sophisticated attacks is ongoing.

Note that accurate wireless coverage maps based on user feedback are crucial not only as input to spectrum sharing mechanisms, but also for a diverse set of applications, such as Wi-Fi-based positioning systems or network planning tools.

The three research dimensions on which we organized our work need not be tackled in an isolated manner, though. We have discussed how user behavior at the reporting level is coupled with the Quality-of-Experience enjoyed at the service provision layer. The system designer, taking into account the fact that there may be disincentives for users to contribute honest feedback, could devise reward mechanisms, where reward for reputable honest users is expressed in QoS/QoE criteria. Future research should focus on the many facets of designing such systems, which range from economics to systems-level issues.

Bibliography

[1] B. Aboba, L. Blunk, J. Vollbrecht, J. Carlson, and H. Levkowetz, "Extensible Authentication Protocol (EAP)," IETF, RFC 3748, June 2004.

[2] B. Aboba and P. R. Calhoun, "RADIUS (remote authentication dial in user service) support for extensible authentication protocol (EAP)," IETF, RFC 3748, Sep. 2003.

[3] N. Ahmed and S. Keshav, "SMARTA: A self-managing architecture for thin access points," in *Proc. ACM CoNEXT '06*, December 2006, pp. 1–12.

[4] X. Ai, V. Srinivasan, and C.-K. Tham, "Wi-Sh: A simple, robust credit based Wi-Fi community network," in *Proc. IEEE INFOCOM 2009*, Rio de Janeiro, Brazil, April 2009.

[5] I. F. Akyildiz, B. F. Lo, and R. Balakrishnan, "Cooperative spectrum sensing in cognitive radio networks: A survey," *Physical Communication*, vol. 4, no. 1, pp. 40–62, March 2011.

[6] R. Albert and A. L. Barabasi, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, p. 47, 2002.

[7] J. G. Andrews, R. K. Ganti, M. Haenggi, N. Jindal, and S. Weber, "A primer on spatial modeling and analysis in wireless networks," *IEEE Communications Magazine*, November 2010.

[8] Apple Q&A on location data. [Online]. Available: http://www.apple.com/pr/library/2011/04/27Apple-Q-A-on-Location-Data.html

[9] Athens Wireless Metropolitan Network. [Online]. Available: http://www.awmn.net

[10] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan, "Interactive WiFi connectivity for moving vehicles," in *Proc. ACM SIGCOMM 2008*, 2008, pp. 427–438.

[11] R. Barbieri, D. Bruschi, and E. Rosti, "Voice over IPsec: Analysis and solutions," in *Proc.* 18th *Annual Computer Security Applications Conference (ACSAC '02)*, 2002, p. 261.

[12] N. Berger, B. Bollobas, C. Borgs, J. Chayes, and O. Riordan, "Degree distribution of the FKP model," *Theoretical Computer Science*, vol. 379, pp. 306– 316, 2007.

[13] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, "Architecture and evaluation of an unplanned 802.11b mesh network," in *Proc. 11th ACM Annual International Conference on Mobile Computing and Networking (Mobicom'05)*, Cologne, Germany, 2005.

[14] G. Biczók, L. Toka, A. Gulyás, T. A. Trinh, and A. Vidács, "Incentivizing the global wireless village," *Computer Networks*, vol. 55, no. 2, pp. 439–456, 2011.

[15] G. Biczók, L. Toka, A. Vidacs, and T. A. Trinh, "On incentives in global wireless communities," in *Proc. of the* 1st *ACM workshop on User-provided networking*, 2009, pp. 1–6.

[16] D. Blumenfeld, *Operations Research Calculations Handbook.* CRC Press, 2001.

[17] Boingo. [Online]. Available: http://www.boingo.com

[18] P. Brady, "A statistical analysis of on-off patterns in 16 conversations," *Bell System Technical Journal*, vol. 47, pp. 73–91, January 1968.

[19] S. Buchegger and J.-Y. L. Boudec, "Performance analysis of the CONFIDANT protocol," in *Proc. ACM MobiHoc*, 2002, pp. 226–236.

[20] R. Chen, J.-M. Park, Y. T. Hou, and J. H. Reed, "Toward secure distributed spectrum sensing in cognitive radio networks," *IEEE Communications Magazine*, April 2008.

[21] S. Cirani, R. Pecori, and L. Veltri, "A peer-to-peer secure VoIP architecture," in *Trustworthy Internet*, L. Salgarelli, G. Bianchi, and N. Blefari-Melazzi, Eds. Springer Milan, 2011, pp. 105–115. [Online]. Available: http://dx.doi.org/10.1007/978-88-470-1818-1_8

[22] A. Clauset, C. R. Shalizi, and M. E. J. Newman, "Power-law distributions in empirical data," *SIAM Review*, vol. 51, no. 4, pp. 661–703, 2009.

[23] R. G. Cole and J. H. Rosenbluth, "Voice over IP performance monitoring," *ACM Computer Communication Review*, vol. 31, no. 2, pp. 9–24, 2001.

[24] D. Cuff, M. Hansen, and J. Kang, "Urban sensing: Out of the woods," *Communications of the ACM*, vol. 51, pp. 24–33, March 2008.

[25] Cuwin - community wireless. [Online]. Available: http://www.cuwireless.net

[26] I. Dangerfield, D. Malone, and D. Leith, "Understanding 802.11e voice behaviour via testbed measurements and modeling," in *Proc. 5^{th} International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt'07) Workshops*, april 2007.

[27] J. R. Douceur, "The sybil attack," in *Proc. 1^{st} International Workshop on Peer-to-Peer Systems (IPTPS)*, 2002.

[28] eBay – All about feedback. [Online]. Available: http://pages.ebay.com/help/feedback/allaboutfeedback.html

[29] E. C. Efstathiou, "A peer-to-peer approach to sharing wireless local area networks," Ph.D. dissertation, Athens University of Economics and Business, 2006.

[30] E. C. Efstathiou, P. A. Frangoudis, and G. C. Polyzos, "Stimulating participation in wireless community networks," in *Proc. IEEE INFOCOM*, Barcelona, Spain, April 2006.

[31] ——, "Controlled Wi-Fi sharing in cities: A decentralized approach relying on indirect reciprocity," *IEEE Trans. Mobile Comput.*, vol. 9, no. 8, pp. 1147 – 1160, August 2010.

[32] Electronic Frontier Foundation – wireless-friendly ISPs. [Online]. Available: http://www.eff.org/pages/wireless-friendly-isps

[33] F. Elianos, G. Plakia, P. Frangoudis, and G. Polyzos, "Structure and evolution of a large-scale wireless community network," in *Proc. 10th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2009)*, Kos, Greece, 2009.

[34] F. Fabrikant, E. Koutsoupias, and C. Papadimitriou, "Heuristically optimized trade-offs: A new paradigm for power laws in the Internet," in *Proc. 29^{th} International Colloquium on Automata, Languages and Programming (ICALP 2002)*, 2002, pp. 110– 122.

[35] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power law relationships of the internet topology," *ACM/IEEE Trans. on Networking*, pp. 514–524, 2003.

[36] O. Fatemieh, R. Chandra, and C. Gunter, "Secure collaborative sensing for crowd sourcing spectrum data in white space networks," in *Proc. IEEE DySPAN 2010*, April 2010, pp. 1–12.

[37] FON. [Online]. Available: http://en.fon.com

[38] P. Frangoudis, V. Kemerlis, D. Paraskevaidis, E. Efstathiou, and G. Polyzos, "Experimental evaluation of community-based WLAN voice and data services," in *Proc. 3rd International Mobile Multimedia Communications Conference (MobiMedia 2007)*, 2007.

[39] P. Frangoudis and G. Polyzos, "Coupling QoS provision with interference reporting in WLAN sharing communities," in *Proc. IEEE PIMRC 2008 SocialNets Workshop*, September 2008, pp. 1–5.

[40] P. Frangoudis, G. Polyzos, and V. Kemerlis, "Wireless community networks: An alternative approach for broadband nomadic network access," *IEEE Communications Magazine*, vol. 49, no. 5, pp. 206–213, May 2011.

[41] P. A. Frangoudis, "The Peer-to-Peer Wireless Network Confederation protocol: Design specification and performance analysis," Master's thesis, Athens University of Economics and Business, 2005.

[42] P. A. Frangoudis and G. C. Polyzos, "Peer-to-peer secure and private community based multimedia communications," in *Proc. IEEE International Symposium on Multimedia (ISM'06) Workshops*, dec. 2006, pp. 1004 –1010.

[43] S. Garg and M. Kappes, "Can I add a VoIP call?" in *Proc. IEEE International Conference on Communications (ICC '03)*, vol. 2, 2003, pp. 779–783.

[44] Google latitude. [Online]. Available: http://latitude.google.com

[45] guifi.net - Open, libre and neutral telecommunications network. [Online]. Available: http://guifi.net

[46] T. Heer, S. Gotz, E. Weingartner, and K. Wehrle, "Secure Wi-Fi sharing at global scales," in *Proc. International Conference on Telecommunications (ICT 2008)*, june 2008, pp. 1–7.

[47] T. Heer, S. Li, and K. Wehrle, "PISA: P2P Wi-Fi Internet sharing architecture," in *Proc. 7th IEEE International Conference on Peer-to-Peer Computing (P2P 2007)*, 2007, pp. 251–252.

[48] *Market Review of Electronic Communications and Postal Services 2009*, Hellenic Telecommunications and Post Commission, 2010.

[49] S. D. Hermann, M. Emmelmann, O. Belaifa, and A. Wolisz, "Investigation of IEEE 802.11k-based access point coverage area and neighbor discovery," in *Proc. 32nd Annual IEEE Conference on Local Computer Networks (LCN 2007)*, 2007, pp. 949–954.

[50] D. Hole and F. Tobagi, “Capacity of an IEEE 802.11b wireless LAN supporting VoIP,” in *Proc. IEEE International Conference on Communications (ICC 2004)*, 2004, pp. 196–201.

[51] hostapd: IEEE 802.11 AP, IEEE 802.1X/WPA/WPA2/EAP/RADIUS Authenticator. [Online]. Available: http://w1.fi/hostapd/

[52] J. Howe, “The rise of crowdsourcing,” *Wired Magazine*, vol. 14, no. 6, Jun. 2006. [Online]. Available: http://www.wired.com/wired/archive/14.06/crowds.html

[53] IEEE 802.11 WG, *IEEE Standard for information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Medium Access Control (MAC) and Security Enhancements, IEEE 802.11i-2004*, The Institute of Electrical and Electronics Engineers, Inc., New York, USA, June 2004.

[54] ——, *IEEE Standard for information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements, IEEE 802.11e-2005*, The Institute of Electrical and Electronics Engineers, Inc., New York, USA, 2005.

[55] ——, *IEEE Standard for information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Std 802.11-2007*, The Institute of Electrical and Electronics Engineers, Inc., New York, USA, June 2007.

[56] ——, *IEEE Standard for information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Radio Resource Measurement of Wireless LANs, IEEE 802.11k-2008*, The Institute of Electrical and Electronics Engineers, Inc., New York, USA, June 2008.

[57] ——, *IEEE Standard for information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 2: Fast Basic Service Set (BSS) Transition, IEEE 802.11r-2008*, The Institute of Electrical and Electronics Engineers, Inc., New York, USA, July 2008.

[58] IEEE 802.22 Working Group on Wireless Regional Area Networks, http://www.ieee802.org/22/.

[59] "Recommendation P.59: Artificial Conversational Speech," International Telecommunication Union, Tech. Rep., March 1993. [Online]. Available: http://www.itu.int/rec/T-REC-P.59/

[60] "ITU-T Recommendation P.800: Methods for subjective determination of transmission quality," International Telecommunication Union, Tech. Rep., Aug. 1996. [Online]. Available: http://www.itu.int/rec/T-REC-P.800/

[61] "Recommendation G.107: The E-model, a computational model for use in transmission planning," International Telecommunication Union, Tech. Rep., Dec. 1998. [Online]. Available: http://www.itu.int/rec/T-REC-G.107-199812-S/en

[62] "ITU-T Recommendation P.862: Perceptual evaluation of speech quality (PESQ): An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codecs," International Telecommunication Union, Tech. Rep., Feb. 2001. [Online]. Available: http://http://www.itu.int/rec/T-REC-P.862/

[63] "Recommendation P.563: Single-ended method for objective speech quality assessment in narrow-band telephony applications," International Telecommunication Union, Tech. Rep., May 2004. [Online]. Available: http://www.itu.int/rec/T-REC-P.563/en

[64] "Recommendation G.113: Transmission impairments due to speech processing," International Telecommunication Union, Tech. Rep., Nov. 2007. [Online]. Available: http://www.itu.int/rec/T-REC-G.113

[65] ITU-T Recommendation H.323, "Packet-based multimedia communications systems," December 2009.

[66] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," *Wirel. Netw.*, vol. 11, no. 4, pp. 471–487, 2005.

[67] S. Jelassi, G. Rubino, H. Melvin, H. Youssef, and G. Pujolle, "Quality of experience of VoIP service: A survey of assessment approaches and open issues," *IEEE Communications Surveys Tutorials*, vol. PP, no. 99, pp. 1–23, 2012.

[68] K. Jones and L. Liu, "What Where Wi: An analysis of millions of Wi-Fi access points," in *Proc. IEEE PORTABLE 2007*, May 2007.

[69] G. Judd, X. Wang, and P. Steenkiste, “Efficient channel-aware rate adaptation in dynamic environments,” in *Proc. ACM MobiSys 2008*, June 2008, pp. 118–131.

[70] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, “Measurement-based self organization of interfering 802.11 wireless access networks,” in *Proc. IEEE INFOCOM 2007*, May 2007.

[71] S. Kent and R. Atkinson, “IP Encapsulating Security Payload (ESP),” IETF, RFC 2406, Nov. 1998.

[72] S. Kent and S. Keo, “Security architecture for the internet protocol,” IETF, RFC 4301, Dec. 2005.

[73] K. Kim, D. Niculescu, and S. Hong, “Coexistence of VoIP and TCP in wireless multihop networks,” *IEEE Communications Magazine*, vol. 47, no. 6, pp. 75–81, June 2009.

[74] N. Koblitz, A. Menezes, and S. Vanstone, “The state of elliptic curve cryptography,” *Design, Codes and Cryptography*, vol. 19, no. 2/3, pp. 173–193, 2000.

[75] R. D. J. Kramer, A. Lopez, and A. M. J. Koonen, “Municipal broadband access networks in the Netherlands - three successful cases, and how New Europe may benefit,” in *Proc. AccessNets ’06*, Athens, Greece, September 2006.

[76] Linux advanced routing and traffic control HOWTO. [Online]. Available: http://lartc.org

[77] Linux Wireless. [Online]. Available: http://linuxwireless.org

[78] Location and my privacy: Windows phone 7. [Online]. Available: http://www.microsoft.com/windowsphone/en-us/howto/wp7/web/location-and-my-privacy.aspx

[79] Managed DNS, Outsourced DNS & Anycast DNS. [Online]. Available: http://dyn.com/dns/

[80] M. H. Manshaei, J. Freudiger, M. Félegyházi, P. Marbach, and J.-P. Hubaux, “On wireless social community networks,” in *Proc. IEEE INFOCOM*, 2008, pp. 1552–1560.

[81] A. P. Markopoulou, F. A. Tobagi, and M. J. Karam, “Assessing the quality of voice communications over Internet backbones,” *IEEE/ACM Transactions on Networking*, vol. 11, pp. 747–760, October 2003.

[82] V. Mhatre, K. Papagiannaki, and F. Baccelli, “Interference mitigation through power control in high density 802.11 WLANs,” in *Proc. IEEE INFOCOM 2007*, May 2007.

[83] B. Milic and M. Malek, “Analyzing large scale real-world wireless multihop network,” *IEEE Communications Letters*, vol. 11, pp. 580–582, July 2007.

[84] S. Miltchev, S. Ioannidis, and A. Keromytis, “A study of the relative costs of network security protocols,” in *USENIX 2002 Annual Technical Conference*, June 2002, pp. 41–48.

[85] A. Mishra, S. Banerjee, and W. Arbaugh, “Weighted coloring based channel assignment for WLANs,” *ACM SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 9, no. 3, pp. 19–31, 2005.

[86] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. A. Arbaugh, “A client-driven approach for channel management in wireless LANs,” in *Proc. IEEE INFOCOM 2006*, April 2006.

[87] R. Murty, A. Wolman, J. Padhye, and M. Welsh, “An architecture for extensible wireless LANs,” in *Proc. HotNets VII*, 2008.

[88] A. Nascimento, A. Passito, E. Mota, E. Nascimento, and L. Carvalho, “Can I add a secure VoIP call?” in *Proc. 7th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM ’06)*, 2006, pp. 435–437.

[89] D. Niculescu, S. Ganguly, K. Kim, and R. Izmailov, “Performance of VoIP in a 802.11 wireless mesh network,” in *Proc. IEEE INFOCOM*, Barcelona, Spain, April 2006.

[90] No-IP – Dynamic DNS, Static DNS for Your Dynamic IP. [Online]. Available: http://www.no-ip.com

[91] NYCwireless. [Online]. Available: http://www.nycwireless.net

[92] OpenVPN: Open source VPN. [Online]. Available: http://openvpn.net

[93] OpenWRT Linux distribution. http://openwrt.org.

[94] P2PSIP IETF working group. [Online]. Available: http://tools.ietf.org/wg/p2psip/

[95] E. A. Panaousis, P. A. Frangoudis, C. N. Ververidis, and G. C. Polyzos, “Optimizing the channel load reporting process in IEEE 802.11k-enabled WLANs,” in *Proc. IEEE LANMAN 2008*, Cluj-Napoca, Romania, September 2008.

[96] J. Pang, B. Greenstein, M. Kaminsky, D. McCoy, and S. Seshan, "Wifi-Reports: Improving wireless network selection with collaboration," *IEEE Transactions on Mobile Computing*, vol. 9, no. 12, pp. 1713–1731, dec. 2010.

[97] T. Papaioannou and G. Stamoulis, "Effective use of reputation in peer-to-peer environments," in *Proc. 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2004)*, April 2004, pp. 259–268.

[98] ——, "Achieving honest ratings with reputation-based fines in electronic markets," in *Proc. IEEE INFOCOM*, 2008, pp. 1040–1048.

[99] ——, "A mechanism that provides incentives for truthful feedback in peer-to-peer systems," *Electronic Commerce Research*, vol. 10, no. 3-4, pp. 331–362, 2010.

[100] B. Patel, B. Aboba, W. Dixon, G. Zorn, and S. Booth, "Securing L2TP using IPsec," IETF, RFC 3193, Nov. 2001.

[101] K. Pelechrinis, M. Iliofotou, and S. V. Krishnamurthy, "Denial of service attacks in wireless networks: The case of jammers," *IEEE Communication Surveys and Tutorials*, vol. 13, no. 2, pp. 245–257, 2011.

[102] M. H. Pinson and S. Wolf, "A new standardized method for objectively measuring video quality," *IEEE Transactions on Broadcasting*, vol. 50, no. 3, pp. 312–322, September 2004.

[103] V. V. Prakash and A. O'Donnell, "Fighting spam with reputation systems," *ACM Queue*, vol. 3, pp. 36–41, November 2005.

[104] J. Rosenberg and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers," IETF, RFC 3263, Jun. 2002.

[105] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session Initiation Protocol," IETF, RFC 3261, Jun. 2002.

[106] N. Salem, J.-P. Hubaux, and M. Jakobsson, "Reputation-based Wi-Fi deployment," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 9, no. 3, pp. 69–81, 2005.

[107] J. Salim, H. Khosravi, A. Kleen, and A. Kuznetsov, "Linux Netlink as an IP Services Protocol," IETF, RFC 3549, Jul. 2003.

[108] N. Sastry, J. Crowcroft, and K. Sollins, "Architecting citywide ubiquitous Wi-Fi access," in *Proc. ACM HotNets VI*, 2007.

[109] A. Satsiou and L. Tassiulas, “Reputation-based resource allocation in p2p systems of rational users,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 21, no. 4, pp. 466–479, 2010.

[110] Seattle Wireless. [Online]. Available: http://www.seattlewireless.net

[111] N. Shahmehri, D. Byers, and R. Hiran, “TRAP: Open decentralized distributed spam filtering,” in *Proc. TrustBus*, 2011, pp. 86–97.

[112] Skyhook Wireless. [Online]. Available: http://www.skyhookwireless.com

[113] J. C. Snader, *VPNs Illustrated: Tunnels, VPNs, and IPsec.* Addison-Wesley Professional, 2005.

[114] *SEC1: Elliptic Curve Cryptography*, Standards for Efficient Cryptography Group, September 2000. [Online]. Available: http://www.secg.org

[115] *SEC2: Recommended Elliptic Curve Domain Parameters*, Standards for Efficient Cryptography Group, September 2000. [Online]. Available: http://www.secg.org

[116] D. M. S thigh, “Law in the last mile: Sharing internet access through WiFi,” *SCRIPTed*, vol. 6, no. 2, p. 355, 2009. [Online]. Available: http://www.law.ed.ac.uk/ahrc/script-ed/vol6-2/macsithigh.asp

[117] K. Stoeckigt and H. Vu, “VoIP capacity–analysis, improvements, and limits in IEEE 802.11 wireless LAN,” *IEEE Transactions on Vehicular Technology*, vol. 59, no. 9, pp. 4553 –4563, nov. 2010.

[118] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, “Chord: A scalable peer-to-peer lookup service for internet applications,” in *Proc. ACM SIGCOMM’01*, 2001, pp. 149–160.

[119] N. O. Tippenhauer, K. B. Rasmussen, C. Pöpper, and S. Čapkun, “Attacks on public WLAN-based positioning systems,” in *Proc. MobiSys’09*, 2009, pp. 29–40.

[120] D. Touceda, J. Sierra, A. Izquierdo, and H. Schulzrinne, “Survey of attacks and defenses on P2PSIP communications,” *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, pp. 1–34, 2011.

[121] J. Tourrilhes. Wireless extensions for Linux. [Online]. Available: http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/Linux.Wireless.Extensions.html

[122] A. Vasan, R. Ramjee, and T. Y. C. Woo, “ECHOS - enhanced capacity 802.11 hotspots,” in *Proc. IEEE INFOCOM 2005*, March 2005.

[123] P. Verkaik, Y. Agarwal, R. Gupta, and A. C. Snoeren, "Softspeak: Making VoIP play well in existing 802.11 deployments," in *Proc. 6th USENIX symposium on Networked systems design and implementation*, ser. NSDI'09, 2009, pp. 409–422.

[124] WiGLE – Wireless Geographic Logging Engine. [Online]. Available: http://www.wigle.net

[125] Wireless Philadelphia Executive Committee. [Online]. Available: http://www.phila.gov/wireless

[126] C. Xenakis, N. Laoutaris, L. Merakos, and I. Stavrakakis, "A generic characterization of the overheads imposed by IPsec and and associated cryptographic algorithms," *Computer Networks*, vol. 50, no. 17, pp. 3225–3241, 2006.

[127] J. Yao, S. Kanhere, and M. Hassan, "Improving qos in high-speed mobility using bandwidth maps," *IEEE Transactions on Mobile Computing*, vol. PP, no. 99, 2011.

[128] P. Zimmermann. PGPfone owner's manual. [Online]. Available: http://philzimmermann.com/docs/pgpfone10b7.pdf

[129] P. Zimmermann, A. Johnston, and J. Callas, "ZRTP: Media path key agreement for unicast secure RTP," Apr. 2011.