DYNAMIC SPECTRUM MANAGEMENT ARCHITECTURE AND ALGORITHMS FOR THE FUTURE MOBILE INTERNET

BY AKASH BAID

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Professor Dipankar Raychaudhuri

and approved by

New Brunswick, New Jersey

May, 2014

ABSTRACT OF THE DISSERTATION

Dynamic spectrum management architecture and algorithms for the future mobile Internet

by Akash Baid

Dissertation Director: Professor Dipankar Raychaudhuri

This thesis presents an investigation of network assisted dynamic spectrum access techniques intended for use with emerging unlicensed, white space, and cognitive radio bands. Dynamic spectrum access (DSA) is motivated by the rapid proliferation of wireless devices which are expected to increase to the order to tens of billions by 2020. The dramatic increase in radio density by as much as 3-4 orders of magnitude relative to today's baseline implies the need for fundamentally new techniques that are both highly efficient and highly scalable. This thesis contributes towards that goal and studies wireless co-existence techniques in the 'age of the Internet' - leveraging ubiquitous network connectivity of wireless devices to enable spectrum co-existence through distributed collaboration.

The first part of the thesis describes the evolution of the mobile Internet, its relation to DSA techniques, and architectural solutions for better supporting current and future mobile Internet use-cases. Through this exercise, the need for network-level collaboration for improving the effectiveness of DSA techniques is shown. In the next two parts of the thesis, two specific applications of such an inter-network cooperation technique are presented - (i) Client-access point (AP) association optimization, and (ii) Channel selection. For the first application, the problem of connecting clients to APs is formulated as a non-linear integer program, and then the effect of inter-network cooperation is shown on the performance of the optimal solution. Large scale simulations with multiple overlapping networks, each consisting of 15-35 access points and 50-250 clients in a 0.5x0.5 sq.km show an average of 150% improvement in random deployments and up to 7x improvements in clustered deployments for the least-performing client throughputs.

In the channel selection application, a new scalable and accurate model for estimating the throughput of a Wi-Fi AP under arbitrary interference graphs is first shown. Based on this model, a graph based channel selection correction-phase is proposed, which can be appended to any centralized channel assignment scheme for performance improvement. Simulations with 100-500 APs/sq.km in homogeneous and mixed settings, corresponding to all APs adhering to same or different channel assignment schemes respectively, show $\sim$30% improvement in the number of starved APs. Further, in the case of mixed deployments, a key finding is made - as the percentage of centrally managed APs in a region is increased in comparison to simple residential APs, the performance of existing managed APs goes down due to decrease in the room for improvement. Results from a series of eight-node experiments on the ORBIT radio testbed are given for further validation of the simulation outcomes.

Having shown the potential gains from cooperation in terms of client-AP association and channel selection, the final part of the thesis outlines the techniques through which such forms of cooperation can be practically implemented. In particular, a specific set of software-defined network (SDN) extensions for wireless control are described in the context of a dense Wi-Fi scenario with multiple network operators. Experimental results from a real-time proof of concept prototype using radio nodes on the ORBIT testbed are given. The results for a small two network scenario validate the proposed inter-network coordination protocol and demonstrate useful performance gains as density increases.

Acknowledgements

The graduate school journey over the last six years has fundamentally changed the way I function and think; and except for the slight graying of my hair, the changes have all been for the better. The struggles of diving deep into technical problems and seeing the proverbial forest for the trees, have taught me that most problems in life, technical or otherwise, are just a thorough literature survey and a robust heuristic away from being solved. The best lessons however, have come from the interactions with my advisor and mentor Prof. Dipankar Raychaudhuri. I have always been inspired by the elegance of his logic, his acute engineering insights, and the relevance of his advice - both technical and on life. This thesis would not have been possible without his constant guidance and support.

I would also like to thank Prof. Roy Yates, Prof. Wade Trappe, and Prof. Jennifer Rexford for being on my thesis committee, and for their advice and suggestions regarding the thesis. Prof. Yates' knack of identifying the crux of any problem and his deep-thinking abilities have helped me immensely, especially at times when I was looking for research directions. I have taken courses under Prof. Yates and Prof. Trappe, which have both been very helpful in building my understanding of systems. Over the last two years, I have also had the pleasure of working with Prof. Rexford, whose clarity of thought is something I will always aspire to attain.

I have learnt a great deal about systems engineering from Ivan Seskar and I cannot thank him enough for answering my endless series of questions on ORBIT and for sharing his practical insights into every problem I have worked on. I am also indebted to several people who I have worked with during the course of this thesis - Dr. Tam Vu, Dr. Kiran Nagaraja, Prof. Yanyong Zhang, and Prof. Marco Gruteser. I was fortunate to work with some great minds during both of my internships; I would like to thank

Dr. Ritesh Madan and Dr. Ashwin Sampath for introducing me to the intricacies of industry research at Qualcomm, and Dr. Prabhakar Chitrapu for helping me navigate the jungle of telecom standards at InterDigital. I would also like to acknowledge the excellent academic and technical support provided by the Rutgers University staff, the staff at WINLAB, and Department of Electrical and Computer Engineering.

I am glad to have always been in the company of close friends both in and outside the lab. In the span of just a few years, I've gained a lifetime's worth of great memories with them - and for that I am very grateful. I will refrain from naming names lest I may forget someone, but simply put, life would have been a chore without my friends. A special thanks is due to Anwesha Chaudhury; it is said that life's journey is easier when you hear a friends footsteps beside you - she has been that constant companion in my PhD journey. Thank you for always being there. Finally, I would like to thank my family - my father Ramesh Kumar Baid and my mother Suchitra Baid, who taught me the importance of hard work and sincerity through their own example; my brother Anupam, who has been a guide and mentor from my first days in the world to the present day; my sister-in-law Aakanksha for being an epitome of positivity in our family; and my kid sister Anisha for never failing to amaze me with her ever-increasing skills and maturity.

The research presented in this thesis was supported by US National Science Foundation (NSF) Future Internet Architecture (FIA) grant CNS-1040735 and the NSF Enhancing Access to the Radio Spectrum (EARS) grant CNS-1247764. I gratefully acknowledge these funding sources that made my Ph.D. work possible.

Dedication

To my parents Ramesh Kumar Baid and Suchitra Baid

Everything I have and will ever accomplish is due to their love and encouragement

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

There are over 7 billion wireless devices in use today, a number that is expected to increase to 50 billion by the year 2020 [1]. This rapid growth in not just the number of wireless devices, but also the amount of communication traffic they generate, the shapes and sizes they come in, the applications they support, and the ways in which people have integrated them in every aspect of their lives is remarkable. Radio technology is, and will remain in the foreseeable future, at the very heart of the information and technology progress of our society.

With the growing adoption of various forms of long- and short-distance wireless technologies, the number of wireless devices that are required to co-exist in time, frequency, and spatial domains, is ever increasing. In order to support the projected rapid increase in deployment densities, there is a need to study spectrum access techniques that can ensure very high spectrum efficiency and yet not be limited in the number of radio devices and technologies it can support. The investigation done as part of this thesis is towards that goal of highly efficient and highly scalable dynamic spectrum access (DSA) techniques, and the approach taken is informed by ubiquitous network connectivity over the Internet.

Throughout the different specific problems addressed in this study, the overarching idea is that of creating a network service for dissemination of spectrum usage information between otherwise independent radio devices and systems, enabling them to implement decentralized spectrum coexistence policies that reduce interference and improve spectrum packing efficiency. This class of network-assisted DSA techniques is applicable to unlicensed band and shared spectrum systems in general (including femtocells), and is particularly relevant to emerging TV white spaces [2, 3] and cognitive radio systems which are still in need of scalable and accurate solutions for both primary-to-secondary and secondary-to-secondary coordination.

In the remaining parts of this chapter, we explain the fundamental motivation for increased visibility between wireless devices through a brief analogy, introduce the interactions between coordinated DSA techniques and the Internet, and highlight the key contributions made in this work.

1.1 Unlicensed band and Skating rink: An Analogy

Unlicensed bands are like skating rinks. If one has never seen or heard about a public skating rink, one could think the whole idea was crazy - toddlers, grandparents, teenagers, speed maniacs, most without helmets or knee-pads skating in a common area with no lanes, minimal rules and no guarantees about not getting hurt! But the fact is, it does work and it can be argued that part of the reason it works is the absence of specific rules.

Unlicensed wireless bands, such as the 2.4 GHz ISM band, too have a large common spectrum-area, different kinds of participants - high bandwidth Wi-Fi, narrow band but hopping natured Bluetooth, etc. Fig. 1.1 shows a simple illustration of this analogy. It has a few transmission rules and most importantly just like the skating rink, no guarantees about interference from other users. You can and probably will collide but overall the system seems to work fairly well.

The idea of spontaneous order: Daniel B. Klein speaks about the same skating rink analogy with respect to society and economy [4]. He argues that "intuition leads us to think that complex problems require complex, deliberate solutions. In a roller rink, the social good depends on getting the patterns to mesh. But no one is minding that good..... but in promoting my interest in avoiding collision with you, I also promote your interest in avoiding collision with me." And that is probably the key: coincidence of interest.

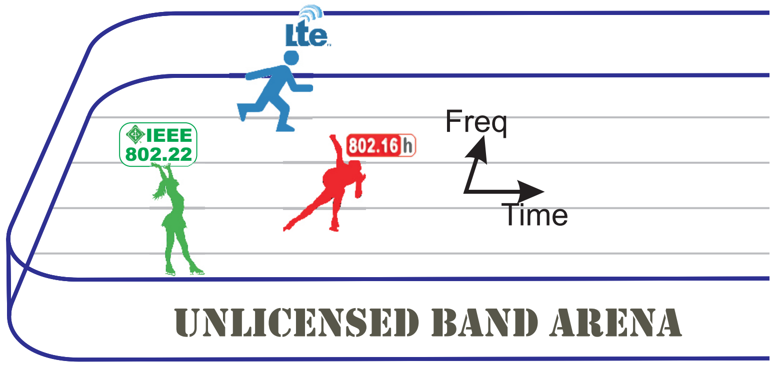

Figure 1.1: An analogy between skating rink and unlicensed wireless band: Similar to how skaters with widely different skill levels coexist on a public rink, wireless devices can seamlessly coexist if they have visibility over nearby devices.

Imagine a “rink master” in the rink who sits at the center and communicates instructions to individual skaters - ‘Move right in 2 seconds’, ‘Increase your speed by 5mph’ or ‘Shift lane in 5 seconds’. Such a system could only work if the rink master precisely knows the capabilities and desires of each participant, not to mention the unwieldy task of calculating the optimum decisions for hundreds of people and communicating it to them in time. This idea directly ties with the centralized versus distributed control debate in DSA. And from the skating rink analogy, we see that if the aim of the system is to accommodate a large number of devices with widely varying access capabilities, data rate, and delay requirements, an open free-for-all approach can work well with even very limited guidance. The coincidence of interest in avoiding collisions and the statistical multiplexing of available data help create spontaneous order in the wireless domain.

Essential requirements for co-existence: Taking the analogy further, there are two important things that minimizes accidents in the skating rink: (i) A basic set of rules: ‘No pushing or sudden stopping’, ‘No loose articles of clothing’, ‘Fixed direction of movement’, etc., and (ii) Visibility: Most importantly, the reason why we do not see constant collisions is that everyone can see and hear what others are doing and decide what is the best action to take in response.

In unlicensed band operation, we have an analogous set of rules that each device has to adhere to but one of the key arguments we make in this thesis is that the visibility of what other devices are doing is very limited. In particular we argue that what is required to promote co-existence is for each device or network to have a basic sense of the wireless environment - more than what the device can see itself, for example, the channel occupancy on other channels, number and type of devices operating nearby, and channel occupancy.

1.2 DSA and the Internet

In the traditional Open Systems Interconnection (OSI) model spectrum access techniques reside in the physical and data link layers while the Internet suite of protocols cover the network and transport layers. However, in a broader scope of the Internet as a global system of interconnected computer networks, the ubiquitous connectivity provided by the Internet can form the conduit for control communication between neighboring devices that employ DSA techniques for co-existence. Beyond this function as a medium for communication, there are also significant benefits of making the network an integral entity in the distributed decision making process. The distinct advantages of using network layer information about heterogeneous radio devices enables a range of feasible coexistence solutions that neither require common physical channels nor rely on sophisticated sensing architectures.

The first proposals for using a common communication channel for spectrum coordination came from the Common Spectrum Coordination Channel (CSCC) [5] and DIMSUMnet [6] projects. While both these schemes called for fine-grained coordination between wireless devices for building what was called as 'Coordinated Dynamic Spectrum Access Networks', CSCC relied on a dedicated radio channel for communication while DIMSUMnet banked on the Internet as the means for communication. Informed by the findings from these projects, the Dynamic Spectrum Access Protocol (DSAP) [7] argued for a Dynamic Host Configuration Protocol (DHCP) like network service for enabling lease-based DSA - an idea close to the approach we take in this work.

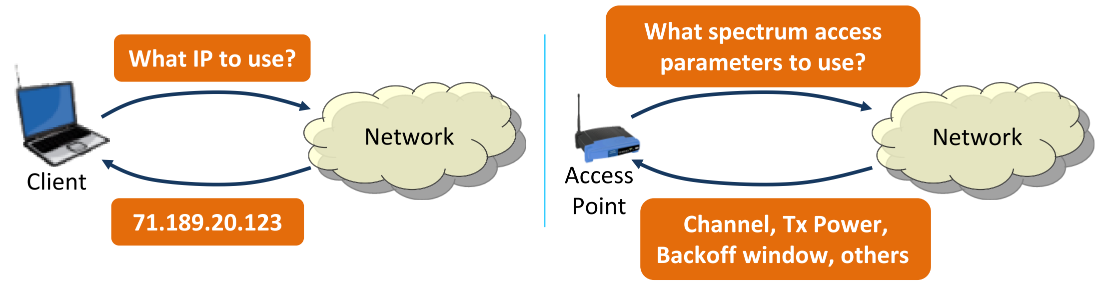

Figure 1.2: Spectrum management as a network service: Analogous to the way a client device's network configuration can be provisioned through DHCP, a network based service can provide spectrum access parameters

Fig. 1.2 shows a simplistic illustration of a DHCP like service for setting up the spectrum access configurations of a wireless device. The key argument for such an approach is that, just like in the case of DHCP, the network inherently has a global view of devices that are already connected to it and their parameter settings. Unlike DHCP however, building such a service for DSA is extremely complex. In order to determine the best set of spectrum access parameters for a given device, such a system of spectrum-management-as-a-network-service would need several inputs such as the geo-location, capabilities, and communication requirements of the device. It would require maintenance of much more state in terms of the transmit parameters of neighboring devices and their traffic pattern per device. And it would also need to be much more dynamic to respond to rapidly changing wireless conditions.

While certainly challenging, recent advancements from a number of other domains can be directly leveraged for network-based DSA. Immense progress has been made on indoor localization techniques to result in feasible, low-cost techniques that provide centimeter-level accuracy in most environments [8]. Techniques for low-latency distribution of messages can be used from past works on Geo-cast routing [9] and in-network multicast [10]. Lastly, the insights obtained from more than a decade of commercial centralized wireless local area network (WLAN) controllers can be leveraged for real-time adaptive spectrum parameter optimizations.

IETF initiatives: Realizing the complete vision of a network-assisted DSA framework still requires substantial progress on several sub-problems. However a few concrete plans have already been initiated through the Internet Engineering Task Force (IETF) towards this goal. The Control And Provisioning of Wireless Access Points (CAPWAP) (RFC 5415) [11] specifies a detailed protocol for standardized communications between a controller entity and wireless devices which accept their spectrum access parameters from that entity. While coming out of the requirements of WLAN deployments, CAPWAP is protocol-independent, and covers discovery of end-points, distribution of messages, and definitions of radio-related information message elements.

More recently, the IETF Protocol to Access White-Space (PAWS) [12], currently a draft proposal, targets the issue of spectrum coordination more directly. It aims to "open the door for innovations in spectrum management that can incorporate a variety of parameters, including user location ... user priority, time, signal type and power, spectrum supply and demand, payment or micro-auction bidding, and more." The specific model used so far for the definition of PAWS includes a database that can track available spectrum and which can inform wireless devices about the spectrum access parameters to use.

With increasing efforts towards the standardization of the communications between wireless devices for the purpose of spectrum management, the key focus now needs to shift towards designing algorithms and architectures that can make use of such communications. We aim to advance the state-of-the-art in that regard by identifying specific scenarios in which coordinated DSA can be most useful. This work can subsequently be useful in building a more seamless, standardized, network-assisted scheme of spectrum management.

1.3 Organization of the thesis

In the first part of the thesis, we study the similarities and differences between the mobile network architecture and that of the Internet, focusing on the placement of functionalities in different parts of the network. We then outline the requirements of mobile and vehicular nodes, and in that context, present a case for re-distribution of certain functions in the Internet. We introduce the key features of *MobilityFirst*, a clean-slate Internet architecture which considers the requirements of mobile nodes from the ground up. We then focus on one particular requirement in the next part of the thesis - the need to improve the first/last hop of the end-to-end Internet path of the mobile nodes in dense unlicensed-band wireless settings.

In such wireless access scenarios, we build upon a rich literature on dynamic spectrum access (DSA) techniques, and introduce an *operational* cooperation mode of DSA through which multiple networks asynchronously share certain key information about their networks with neighboring networks. This information is then used by each network while choosing the spectrum access parameters for elements in its own network. As concrete examples of such a mode of cooperation, we present two detailed case studies - one on client-AP association optimization, and one on channel assignment.

The concluding part of the thesis outlines the techniques through which this form of cooperative DSA can be implemented. In particular, we define a controller-to-controller API through which multiple wireless networks can take part in cooperative spectrum management techniques. This section also presents some results from a real-time proof-of-concept prototype which realizes inter-network coordination using the ORBIT testbed as experimental platform.

Chapter 2

Evolution of Mobile Internet Architecture

The architectural design of mobile networks and the Internet are, in many ways, opposite to one another - while the former entails centralized ownership and tight control, the latter is decentralized by design and involves cooperation between different entities. However, in spite of occupying two very different points in the design space, the two networks have an important common characteristic - both designs have proved to be incredibly resistant to architectural evolution. While the most prominent use-case of mobile networks has emphatically shifted from voice calls to data packets, the existing plumbing architecture is yet to evolve accordingly. Similarly, it is well recognized that the Internet architecture, largely designed for one-to-one interaction between fixed end hosts and servers, is ill-suited for the increasingly dominant use of content consumption by both fixed and mobile devices [13].

While efforts are underway to bring about change in both the architectures individually (most notably through Long Term Evolution (LTE) in mobile networks and Information Centric Networking (ICN) standards in the Internet), the interplay between these two networks is not very well understood. In particular, with Internet access becoming the dominant use-case of mobile networks, and mobile end-points projected to exceed fixed end-points in the Internet [14], the evolution of these two architectures should take into account the requirements and characteristics of one another explicitly. The broad aim of this thesis is to first understand and then leverage the convergence of mobile network and Internet architectures to enable decentralized dynamic spectrum coordination.

In this chapter, we explain the evolution of the mobile Internet architecture in brief, and through it, motivate the specific DSA problem that we concentrate on in the rest of the thesis.

2.1 Emergence of a Leaner Access Network

Fig. 2.1(a) shows the typical architecture of a mobile network, characterized by planned deployment of base stations, centralized control of both wireless and wired components, and tight enforcement of policy and QoS characteristics. While this managed architecture has served well for offering reliable services to end users, it has also led to extremely high levels of capital and operational expenses for the network operators. Moreover in the current data-dominated regime, it forms an inefficient path for most bytes flowing through the network, since user devices themselves are much closer to broadband-served Internet end-points than the mobile gateways.

New network models have started to challenge this centralized architecture in two prominent ways: (i) residential and enterprise femto cells, and (ii) Wi-Fi based Mobile Virtual Network Operators (MVNOs). Femto cells (or more generally small cells) are being promoted by the mobile network providers mainly to increase the capacity of their access networks. These small form-factor basestations connect directly to the on-premise broadband Internet connection and can either offload traffic directly to the Internet or establish a tunnel to the mobile network. While this breaks the centralization of the data path to some extent, it requires careful orchestration of the access frequency and power levels to avoid interference with the macro base stations. More importantly, being still operator owned and managed, the use of small cells is not expected to result in a substantial reduction in the cost of mobile Internet access to end-users.

A new set of MVNOs, on the other hand are directly aiming to bring low-cost alternatives to the users by directing all traffic, including voice calls, through available Wi-Fi networks and falling back on cellular coverage only in the absence of any Wi-Fi network. Republic Wireless is an example of such a MVNO and it currently offers a $20 unlimited voice, text, and data plan in the US [15]. As shown in Fig. 2.1(b), such an architecture results in a very lean access network which simply provides a connection path to the services which are hosted on remote servers.

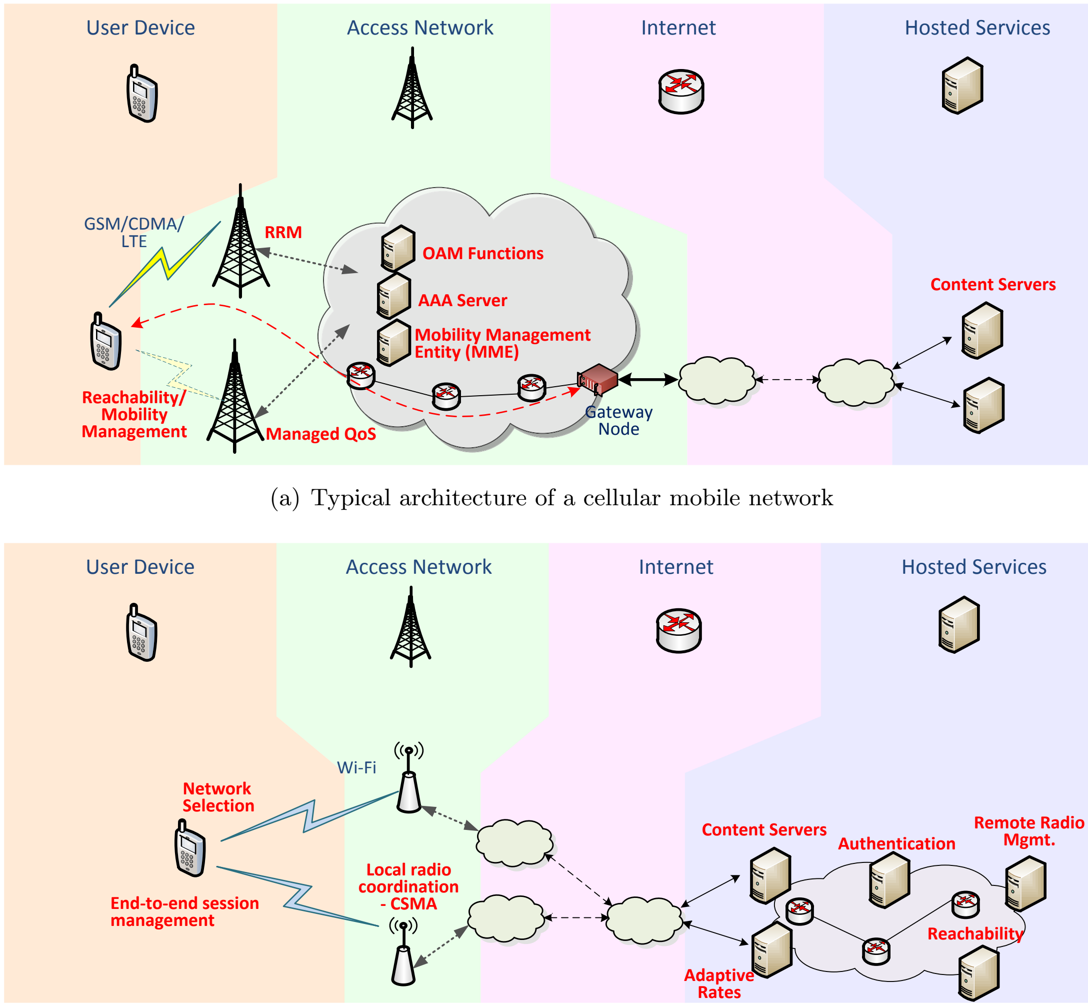

(a) Typical architecture of a cellular mobile network

(b) Architecture of a WLAN based mobile network

Figure 2.1: Location of different functionalities in a traditional mobile architecture and a WLAN-based mobile architecture

The key difference between the two architectures described above is the distribution of network functionalities between elements residing in the end-to-end path. In the traditional mobile network architecture, which was originally built to support voice calls only, all the critical tasks such as reachability and mobility management, radio resource management, authentication, and quality-of-service assurance, reside solely inside the access network. The transport network beyond the mobile operator's network just forms the pipeline between the operator-owned network and the hosted services. In contrast, the WLAN architecture, mainly in-order to reduce infrastructure and operating cost, opts for most functionality being implemented as cloud-based services.

The shift towards this approach has been guided by a number of factors among which the commoditization of computing elements, increase in the speed and reliability of the transport network, and cost-reduction of Wi-Fi access points, are the most prominent. The core Internet in this new architecture, however, just like the older mobile network design, serves the simple role of transporting packets between the access network and the cloud-based servers. The only difference being the replacement of a single gateway-type node of the mobile network architecture to arbitrarily many points-of-attachment to the network in the WLAN architecture.

Both the architectures shown in Fig. 2.1 are admittedly over-simplified - in reality, the boundaries between the four components of the network are often blurred and in general it is hard to label parts of the network as access, transport, and core. In addition several new and old attempts to optimize different parts of the end-to-end service break this simple view. For example, Content Delivery Networks (CDNs) aim to bring hosted services and content close to the users, often inside the access networks themselves, and the emerging idea of Cloud Radio Access Networks (Cloud RAN) attempts to thin the access network to just a collection of antennas with all the processing done inside remote data-centers [16]. Despite these shortcomings, the figures highlight the contrast in the placement of different functionalities in the two scenarios, which is critically important to the type and quality of the services that can be offered based on those architectures.

2.2 A Mobility-centric Internet Architecture

While this leaner-access-network mobile architecture is gradually maturing, the thrust of academic and industry research has now shifted to the re-design of another component of the end-to-end mobile Internet solution shown in Fig. 2.1 - that of the core Internet itself. The need for a new architecture for the Internet is primarily two-fold. First, the explosion in the number of mobile end-points connected to the Internet, and second, the shift of the dominant use-case of the Internet from communication between specific end-points to fetching of specific content pieces. This has motivated a number of "clean-slate" future Internet architecture projects aimed at investigating fundamentally new approaches for making the Internet better suited to these fundamental changes [17–21].

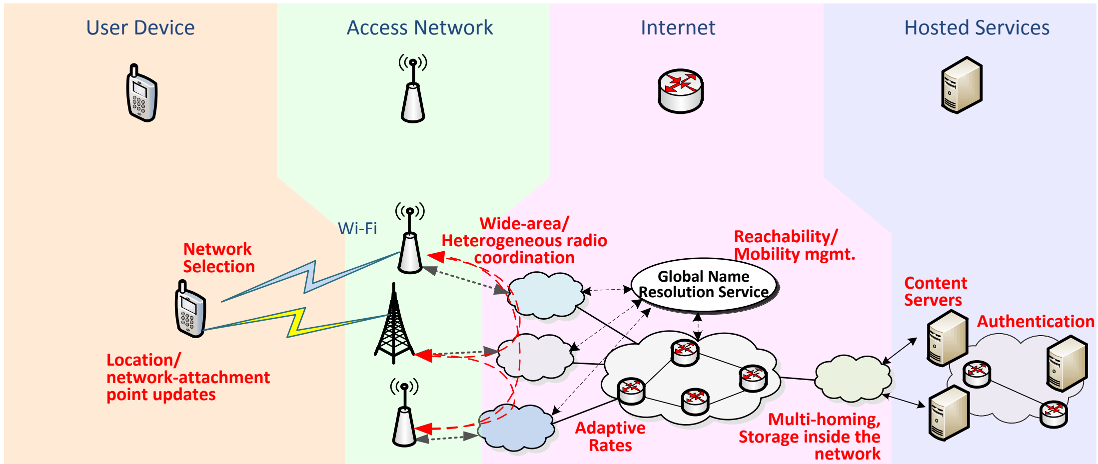

Figure 2.2: A mobility-centric Internet architecture

The work done as part of this thesis has contributed towards the basic design and development the MobilityFirst Internet Architecture, one such clean-slate effort with a particular focus on supporting large-scale, efficient and robust mobility services in the future Internet [22]. While the decisions on the placement of functionalities in current deployments have been driven more by business considerations than technical ones, the MobilityFirst project aims to reason about what functionality should be placed where depending on considerations for the requirements and the ubiquity of the service. This exercise has resulted in an architecture that is shown in a simplified fashion in Fig. 2.2 - note that in this architecture some of the functionalities reside in the transport network, i.e. the Internet, instead of its either ends - the access and core networks. In the remaining part of this chapter, we outline the wireless/mobile edge network perspective behind the MobilityFirst design to explain the rationale behind the architecture in Fig. 2.2, and then identify some of the resulting key protocol features of MobilityFirst.

2.2.1 Challenges and Requirements

To understand what types of services might be justified being placed in the network, we need to understand the specific requirements that current and future mobile nodes have from the Internet. In this section, we present the challenges and requirements of five key wireless access scenarios [23].

A. Host and Network Mobility The foremost characteristic of untethered nodes is that their points of attachment to the Internet can change easily and rapidly. The need for supporting mobility arises when an individual node or a group of nodes, for example a bus/train/plane network, moves and reconnects to the Internet. There has been extensive work on enhancing the Internet protocol suite to support mobility, most notably with mobility anchors as in Mobile IP [24]. These solutions are based on a set of implicit assumptions that users have an immutable "home" network, are connected to a single network at a time, and transitions across networks are infrequent. Consequently, packets in the current architecture are sourced from, and destined for, the network attachment point of end-hosts, i.e. their IP addresses. However, this network model has changed since Mobile IP was conceived. It is important to understand the simple but fundamental requirement for mobility support hosts need permanent names irrespective of their attachment points, and the network needs a packet transmission primitive that employs permanent names. This functional requirement can be translated to the following protocol design requirements (as outlined in Fig. 2.3):

A1. Disambiguation of the dual-roles of an IP address as both an identifier and a locator into two different primitives - a permanent name and a network-specific temporary locator.

A2. Dynamic binding of names to network addresses/locators.

B. Varying wireless link quality Fluctuations in access link quality are an intrinsic property of the wireless medium achievable bit rates in both Wi-Fi and 4G systems, can show large variations within a fraction of a second and disconnection due to mobility and/or insufficient signal strength is not uncommon. For example the sample trace of downlink throughput in an experimental 4G network shows bit-rate variations greater than 3:1 during just tens of seconds (Fig. 2.4). While these variations are usually handled at the PHY and MAC layers, they invalidate some implicit assumptions in the control algorithms used in the Internet. For example, it has been long known that TCP congestion control treats wireless link errors as congestion losses and performs poorly in high variation wireless channels [25]. Given the increasing dominance of the wireless

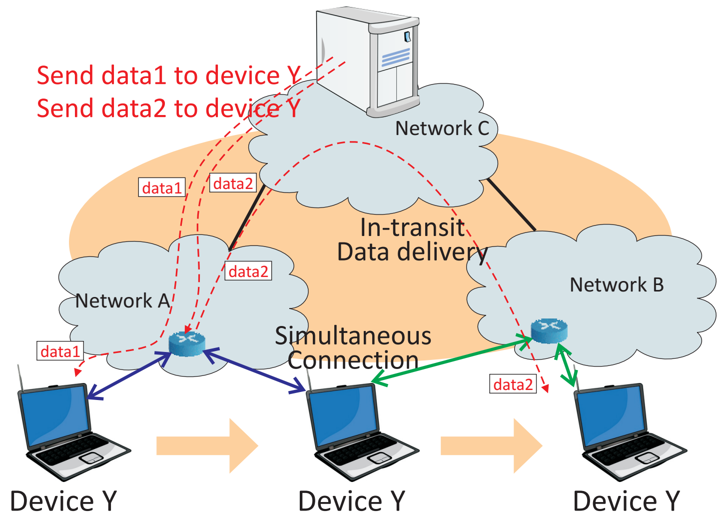

Figure 2.3: A server in Network C is sending packets to a host moving from Network A to B. Seamless delivery of packets can be achieved if packets are destined for the device rather than the current network address of the device.

last hop for Internet access, such link quality variations need to be natively supported at different layers of the Internet architecture. This leads to the following requirements:

B1. Link quality awareness at both the intra-domain and inter-domain routing layers to enable robust packet delivery strategies.

B2. Disconnection-tolerant routing and transport protocols that are capable of temporarily storing packets during disconnections and rerouting in-transit packets to new points of attachments.

C. Accessing multiple networks A typical wireless device in an urban area today might see 3-5 cellular networks and 10-20 Wi-Fi access points, but accesses only one of these due to both technical and business model constraints. Current techniques supporting simultaneous use of multiple interfaces rely on enhancements to the underlying end-to-end transport layer (see [26] and references therein). Specifically, these mechanisms require a multi-homed end-point to inform the sender about its multiple interfaces prior to the commencement of data-flow, and a data-striping algorithm on the sender stack that adapts the packet rate of each interface. This results in rigidity in two key aspects: (i) There is no mechanism by which users can specify under what conditions, and in what manner the interfaces are to be used; (ii) Since all decision

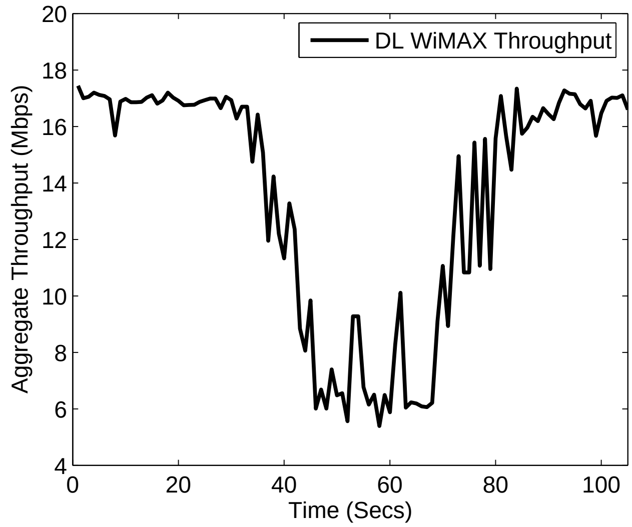

Figure 2.4: Variations in downlink throughput measured for a client connected to a GENI WiMax base station [10]

logic is implemented only at the end-nodes, in-network routers cannot adapt or buffer the flows in accordance with wireless channel quality variations. Thus efficient support for host multi-homing induces the following key requirements (see Fig. 2.5):

C1. Support for binding a single name to multiple addresses and interfaces.

C2. A routing plane capable of modifying the data-striping and storing decisions in accordance with the link quality at each interface.

C3. Service semantics to support interface selection and utilization (e.g. send to all interfaces, send to higher-throughput interface, send only to Wi-Fi, etc.).

D. Ad hoc networks Wireless ad hoc networks are important for infrastructure-less vehicle-to-vehicle (V2V) and sensor network scenarios, last-mile connectivity and applications such as photo/video sharing, local social networking, and multi-player gaming. One view of Internet design is that ad hoc networks are just a type of edge network; as long as they are connected to the Internet via a boundary IP router, the protocols used within the ad hoc network can be ignored. However, the ubiquity of non-specialized devices requiring support for ad hoc networking (e.g. phones, tablets, laptops, vehicular infotainment systems, etc.) forms a strong argument for an integrated design that avoids boundary translation solution. Integration of such networks within

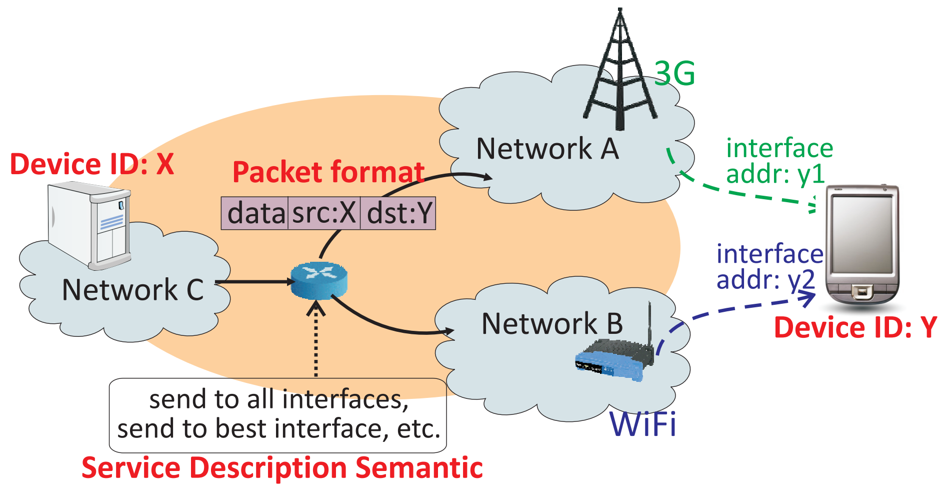

Figure 2.5: Qualitative illustration of multi-homing support in the Internet

the framework of a future Internet design results in the following distinct requirements:

D1. Critical network services such as authentication and dynamic binding of names to addresses should be capable of disconnected-mode operation.

D2. Routing and transport protocols should be robust to opportunistic association and changing network topologies.

E. Spectrum Access Coordination Finally, a critical challenge that differentiates wireless networks from wired networks, but which is common across all forms of wireless networks (such as cellular LTE, Wi-Fi, white-space networks, etc.) is the need for devices to coordinate their use of spectrum. These coordination schemes, whether centralized, distributed, or a hybrid, are typically implemented through overlay channels for example, the IETF PAWS protocol for accessing white space database uses an HTTPS overlay [12], and the X2 interface between LTE base stations uses SCTP over IP [27]. However supporting these wireless control plane functions at the scale of thousands of devices/km requires an integrated approach satisfying the following requirements:

E1. Support for a low-latency control plane that is unaffected by data plane congestion.

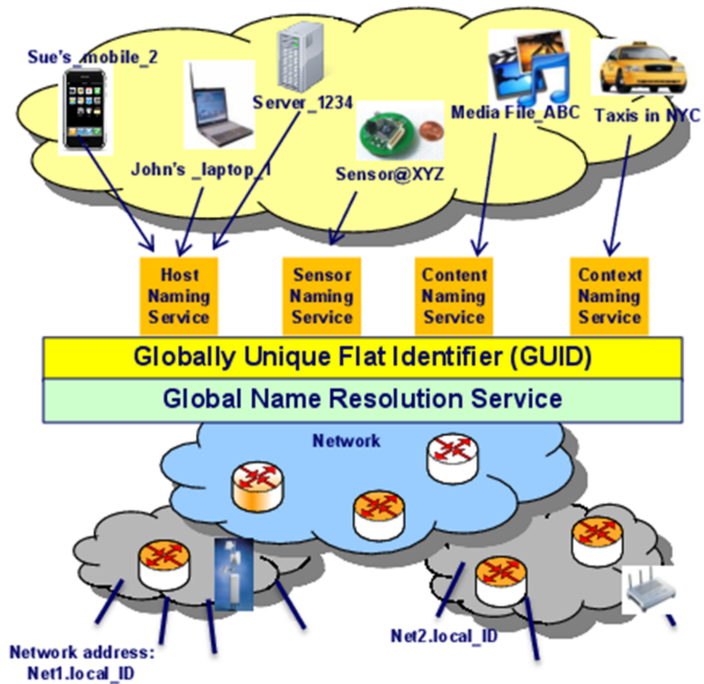

Figure 2.6: Separation of identification and network location in the MobilityFirst architecture

E2. Dynamic multicast of control messages, based on geographic location and radio-range of the sender, to enable efficient distributed coordination schemes.

2.2.2 Key Features of MobilityFirst Future Internet Architecture

The MobilityFirst architecture is built upon a new name-based service layer which serves as the "narrow-waist" of the protocol stack. The name-based service layer uses flat globally unique identifiers (GUIDs) for all network attached objects. GUIDs are different from the IP addresses of the current Internet architecture in two significant ways: (i) IP addresses are overloaded to signify both the identity and the location of an end-point, whereas GUIDs serve just as the long-lasting, consistent identifiers, (ii) IP addresses are typically assigned to net devices, but GUID is a single abstraction that covers a broad range of communicating objects - from a simple device such as a smartphone, a person, a vehicle, a group of vehicles, a piece of content, and even context, as shown in Fig. 2.6.

A GUID can be assigned to a network object by one of multiple name certification services (NCSs), and is derived through a cryptographic hash of the public key that corresponds to that object. The GUID being directly derived from the public key gives it a self-certifying property, i.e. authenticating a node does not require an external authority [28]. This feature solves an important problem in mobile environments where communication to a third-party server is often not possible or introduces substantial delay to critical applications. Identifying mobile nodes by long-lasting unique identifiers also helps in another fundamental challenge in such scenarios - that of mobility management. In MobilityFirst, GUIDs assigned to network objects are mapped to a set of network addresses (NAs) or locators corresponding to the current points of its attachment to the Internet. This enables a scalable name-based service API, i.e., packets can be sent based on the GUID of the destination, which is automatically resolved to the current NA or NAs based on where in the network the object is located. In the following subsections, we present further details on key architectural components.

Dynamic Name-Address Bindings

The GUID-based protocol stack described above handles host and network mobility through fast dynamic binding of identifiers to locators. That is, when a user sends packets directed to a particular identifier (GUID), the networking protocol must quickly ascertain the set of locators (NAs) attached to the GUID and route the packets correspondingly. We address the challenge of providing a fast global name resolution service at Internet scale through a router DHT-based *Direct Mapping (**DMap**)* scheme for achieving a good balance between scalability, low update/query latency, consistency, availability and incremental deployment [29]. In order to perform the name resolution for a given GUID, DMap distributes the GUID→NA mappings amongst Internet routers using an in-network single-hop hashing technique which derives the address of the mapping router directly from the GUID. Through a detailed simulation study described in [29], we have shown that DMap achieves a 95th percentile round trip query response time of below 100ms (Fig. 2.7 presents the key query response time result), considered more than adequate for current and future mobility services [30, 31]. The

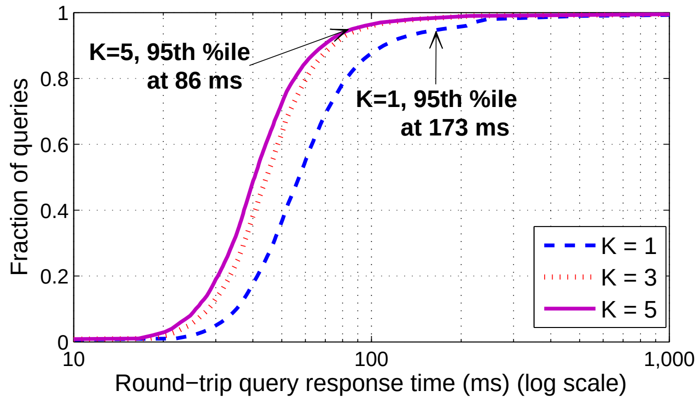

Figure 2.7: CDF of round trip query time from a measurement driven Internet scale simulation of 1 Million name resolution queries passing through a realistic Internet model. K represents the number of replicas of each mapping and provides a tradeoff between response time and storage load.

dynamic binding of GUIDs to network addresses thus helps meet mobility and multi-homing requirements.

Storage-aware and Delay Tolerant Routing

MobilityFirst uses a generalized storage-aware routing (GSTAR) algorithm to support delay and disruption tolerance in the routing layer. In GSTAR, each router employs in-network storage that facilitates store vs. forward decisions in response to varying link quality and disconnections [10]. These decisions are based on both short-term and long-term path quality metrics. In addition, packets along paths that become disconnected are handled by a disruption tolerant networking (DTN) mode of the protocol with delayed delivery and replication features. In particular, each router maintains two types of topology information:

1. An intra-partition graph is formed by collecting flooded link state advertisements which carry fine-grained, time-sensitive information about the intra-network links

2. A DTN graph is maintained via epidemically disseminated link-state advertisements which carry connection probabilities between all nodes in the network.

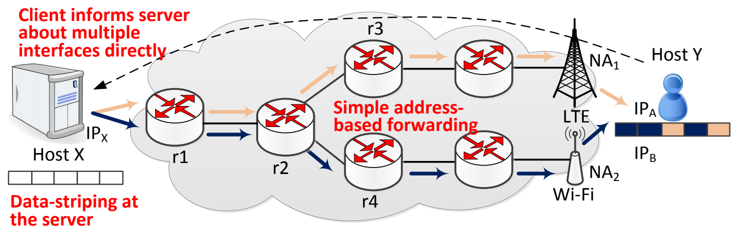

(a) End-to-end multihoming in IP based networks

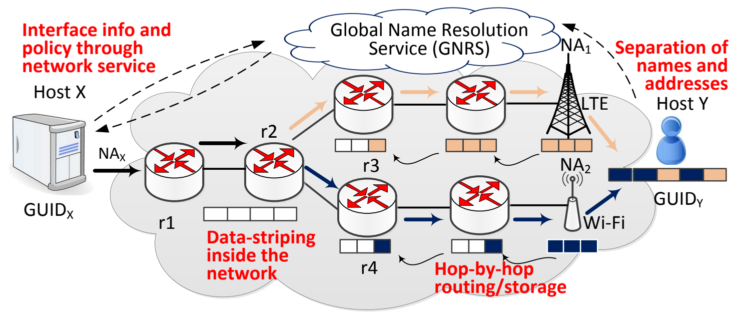

(b) Network-assisted multihoming in name-based networks

Figure 2.8: Key conceptual differences between conventional TCP/IP based multihoming and the proposed network-assisted approach

Recent results indicate that by intelligently utilizing in-network storage, GSTAR outperforms traditional and storage-augmented link-state protocols in both wired and wireless network environments [32].

In MobilityFirst, the requirements of multi-homing are met by incorporating support for multi-homed nodes directly in the routing layer (as opposed to the current end-to-end approach). As shown in Fig. 2.8, MF multi-homing makes use of network-assistance in two important aspects. First, the GNRS is used by multi-homed nodes to specify the availability of multiple interfaces and the corresponding interface preference policies. Second, the task of data-striping is shifted from the end-host stack to the in-network routers which have a better view of the end-to-end path quality through the underlying routing layer.

Spectrum Access Coordination

The MF management plane facilitates unlicensed band spectrum coordination through dissemination of spectrum usage information to networks within radio interference range of each other. In this architecture, routers which directly connect to the base stations/home APs run an evolved flavor of geocast routing [9] which stores the information about the region of operation of each network that they support. A simple example is illustrated in Fig. 2.9. The source X of any spectrum management message, signs it using $\{L_x, r_x\}$ where L_x is the geo-location of X and r_x is the radius of operation obtained by equating: $PL_x(r) = P_{x,max} + G_x - S_{x,min} - N$, where $PL_x(.)$ is the appropriate indoor/outdoor pathloss model used, G_x and $P_{x,max}$ are the antenna gain and maximum transmit power of X respectively, $S_{x,min}$ is the minimum received power required for operation and N is the noise floor. A router which receives the message tunnels it to a known radio resource management (RRM) entity which stores the list of $\{L_i, r_i\}$ pair for each of the networks that it supports either directly or through another server. The RRM entity, upon receiving this message, checks to see if the source region in the message overlaps with any of its networks and passes the message to all overlapping networks. It further routes the message to its connected RRM entities which then sends it to other routers connected to it using a similar overlap search. The key advantage of using this network service model for neighbor communication is the fact that devices or networks do not need to store the states of all interfering networks nor keep track of networks joining and leaving the neighborhood.

2.3 The Case for Network Assisted Services

Based on this description of these key features, we can see that MobilityFirst proposes moving three broad functionalities inside the network: (i) Reachability/Mobility management, (ii) Link quality based adaptations, and (iii) Support for spectrum coordination. The tradeoffs involved in these architectural choices were explored under the broader design of the MobilityFirst future Internet architecture project [22].

In order to understand what we mean by a ‘functionality residing in the Internet’,

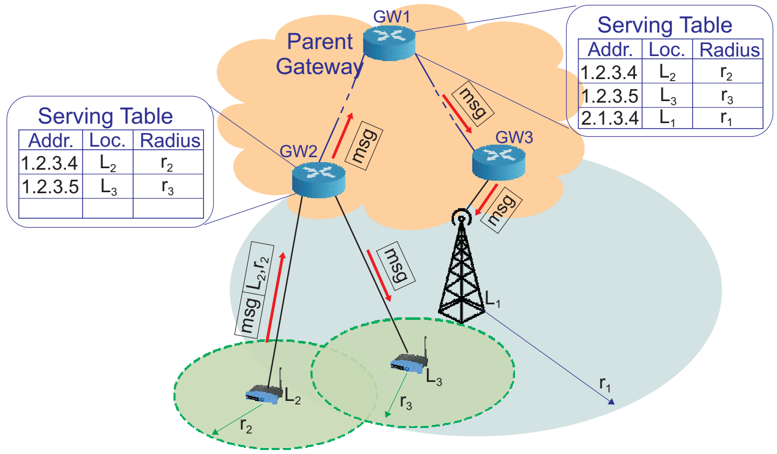

Figure 2.9: Network support for neighbor selection

consider the case of mobility management services. An example of mobility management being implemented in the access network is the paging process for completing a voice call - the mobile network keeps track of which basestation a user is connected to and another user willing to connect to this user implicitly makes use of the mobile number to basestation mapping service which is hosted completely inside the mobile network. On the other hand, some forms of mobility management are implemented 'over-the-top' in remote servers, for example, a message sent through Apple's iMessage service uses their mobility management service which keeps track of the IP address through which each user is connected to the network. In contrast to the examples above, when the mobility service is implemented in the Internet, the mapping of the user's telephone number or another form of unique identifier to the network attachment point is maintained by one or more networks that comprise the Internet. Thus a user can send messages into the Internet by specifying just the identifier of the desired destination instead of first using a directory service to determine where the user is and sending the message to the specific network location. Clearly this makes the task of the transport network more complex - in addition to providing routing services, it has to provide other types of services.

Here we reason about why the evolved mobile Internet should play a greater role in the support for these services, in the narrow context of the placement of functionalities discussion continued from Sec. 2.1.

- Reachability/Mobility management: The number of hosted services used by mobile users is continuously increasing. For any service which requires real-time reachability, there needs to be constant updates from the user device to the server as the user moves around. For example, a user having the popular Skype, Facetime, Google chat, Whatsapp, Facebook messenger, and Viber mobile applications running simultaneously on her phone will require application layer update messages to be sent to the directory service of each of these applications individually, whenever it changes its point of attachment to the network. Instead, if all applications use a single unique identifier for the user and the identifier to network locator mapping is done inside the Internet, the cost of the mapping infrastructure can be de-duplicated. In addition the placement of the mobility management functionality inside the network would enable dynamic, and when needed, late binding of the identifier to the network address of the mobile node. This would prevent loss of in-transit packets for fast moving nodes and allow for seamless support for multi-homing.

- Link quality based adaptations: The second type of network service or functionality which can greatly benefit from being implemented in-network rather than in end-device stacks, are adaptations in response to the changing link quality of the mobile nodes. Examples of such adaptations are: (i) Varying the sending rate in response to changing bandwidth of the wireless last-hop, (ii) Adaptation of the ratio of packets sent on each of multiple available wireless interfaces, (iii) Changing the routes of packets, especially in ad hoc scenarios, and (iv) Varying the amount of buffering of packets en-route to prevent packet-loss due to temporary disconnections. In each of these applications, the key benefit of an in-network implementation is low latency and better visibility of the path characteristic than what is available at the edge of the network. If the adaptation can be carried out at arbitrary, and when required, multiple points in the path, the respective parameters under consideration can be varied at the same, fast granularity as the wireless links vary.

- Support for spectrum coordination: While the requirements of spectrum coordination between networks is usually not considered in the purview of the transport network architecture, dynamic spectrum access (DSA) techniques can greatly benefit from the existence of a fast, ubiquitous communication channel for coordinating the use of the spectrum between heterogeneous devices. A key constraint in traditional DSA techniques is the lack of a common control channel between interfering nodes, since different nodes can use different radio technologies and the interference might be asymmetric. The most popular way around this constraint is the use of remotely hosted databases which can act as a common coordination point for disparate wireless nodes. However, if the underlying network can facilitate the communication between neighboring nodes, the scale and latency of the resulting DSA techniques can be vastly improved.

In our previous works, we describe the in-network reachability/mobility approach [29, 33], and the in-network support for link-quality based adaptations [34, 35]. A key focus of this thesis is on the design and analysis of in-network services of the dynamic spectrum coordination mentioned earlier. In particular, we explore the feasibility of collaboration between multiple wireless networks which have overlapping areas of coverage but which are interconnected through the Internet via direct or indirect paths.

Chapter 3

Inter-network Collaboration and its Application

The basic design goal of our network-assisted DSA scheme is to create network support for sharing of spectrum usage information between collocated wireless networks. The scope of such useful spectrum information includes transmitter and receiver locations, transmit power, bandwidth of operation, channels being used, radio sensitivity, SNR vs. bit-rate, MAC schemes being employed, antenna properties, etc. Each network can use these parameters in autonomous distributed algorithms for spectrum sharing (such as bandwidth/rate backoff similar in spirit to TCP congestion control). In this scheme, the sharing of spectrum use information, however, is just the first part of the solution that of increasing the visibility of each transceiver much beyond what it can sense on its own. The second part comes from the ability to instantiate a higher-layer negotiation protocol between neighboring networks to support joint assignment/management of spectrum resources, negotiations between heterogeneous entities, and controller delegation.

3.1 System Architecture

Fig. 3.1 shows the functionality of the proposed spectrum management solution along with the two levels of interactions between two adjacent networks. As shown in the figure, each network collects radio parameters from each transceiver in its domain and then summarizes them into an “aggregate radio map” which is shared with neighboring networks with one or more radios within the interference region. There is also a second control interface with higher-level semantics required to support policy expressiveness, global optimization algorithms, and controller delegation associated with management-level coordination of autonomous networks.

Fig. 3.2 shows a physical world view of the proposed system. Each wireless network

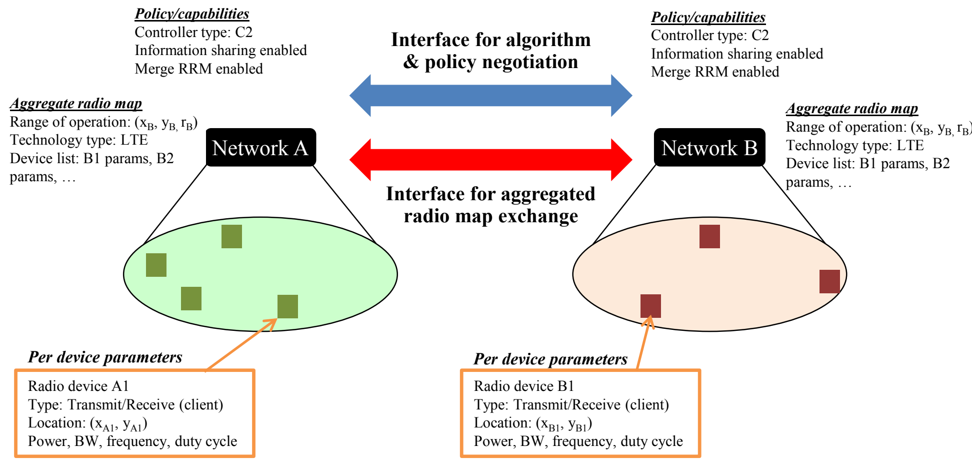

Figure 3.1: Functional outline of the proposed inter-network collaboration approach

has a local controller which collects radio device parameters as summarized in Fig. 3.1, along with an RRM control interface for setting parameters for operation. The local controllers communicate with each other over the control plane designed to have two specific services the first is a geographic multicast ("geocast") service which delivers the aggregate radio map to all networks in the region of interest (as calculated from the radio coverage parameters). This ensures that networks have information about spectrum use by other networks in the region, thus enabling each RRM controller to execute an appropriate distributed spectrum coordination algorithm to avoid excessive interference and achieve good spectrum use efficiency. With increasing spectrum packing, it may also be desirable for interfering networks to negotiate directly with each other using the management control interface shown for example two overlapping Wi-Fi networks in an urban area can use this interface to agree on a common radio resource management algorithm and merge their controllers to run a single more global scope algorithm this has the effect of creating a unified virtual wireless network with a single logical controller delegated to one of the networks involved. Merging of controllers via delegation can be realized with software-defined network technology, which we discuss further in Chapter 5.

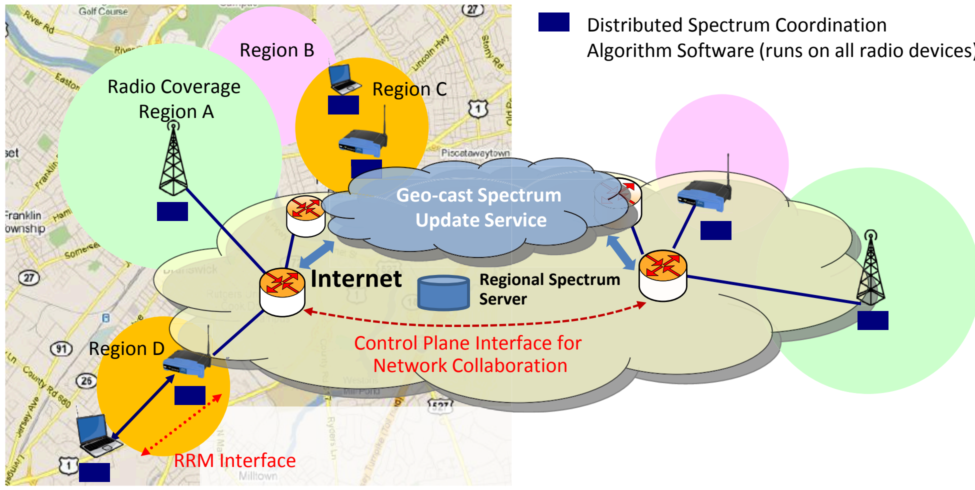

Figure 3.2: System architecture for network assisted spectrum coordination service

3.2 Background on Client-AP Association Problem

Next, we present a detailed use-case of inter-network cooperation for the optimization of client-AP associations in Wi-Fi networks. In a Wi-Fi deployment with multiple access points, optimizing the way each client selects an AP from amongst the available choices, has a significant impact on the realized performance. When two or more such multi-AP networks are deployed in the same region, APs from different networks can cause severe interference to one another. In order to show the use of network-assisted coordination described in Sec. 1, we study how *inter*-network interference affects the *intra*-network association optimization and propose a cooperative optimization scheme to mitigate the interference.

In order to alleviate inter-network interference, we propose back-end *operational cooperation* between the networks: each network periodically shares the information about the location and operating channels of its APs with all other networks operating in the same area. Note that clients belonging to one network cannot join other networks in this model. Within the scope of the traffic model described in Section 3.4, this form of information exchange followed by *intra*-network optimization is the same as a global optimization considering all APs of all networks as being controlled by a single entity. This follows from the fact that for certain problem formulations, the interference terms in the intra-network problem can be summarized and substituted using the information

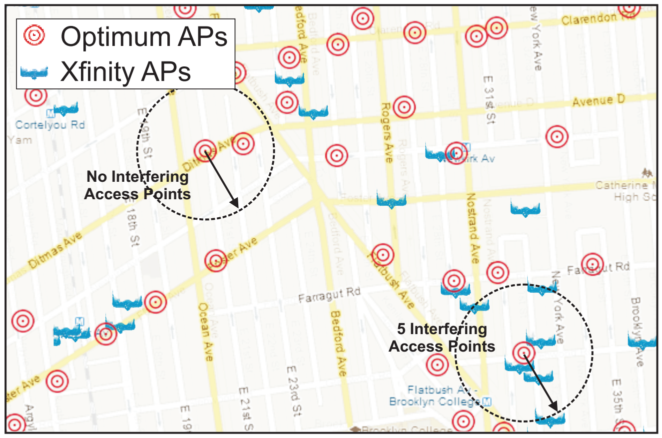

Figure 3.3: Wi-Fi AP locations of two providers in Brooklyn, NY

received from neighboring networks. To the best of our knowledge, such forms of cooperation between multiple managed Wi-Fi networks have received very little attention with only some recent works in the related area of cellular networks [36].

Figure 3.3 shows a real-world example of overlapping Wi-Fi AP deployments of two leading broadband Internet providers in a ∼1 sq. km. cross section of the Brooklyn area in New York, USA, compiled using their respective Wi-Fi location finder services [37, 38]. The exact nature of inter-network interference on client throughputs in such a scenario depends on the number of co-channel APs, their transmit powers, rate allocation algorithm, and MAC parameter selection. However, inference of these channel access parameters through passive observations is a hard problem and often requires active probing [39]. A key challenge in passive interference estimation is to incorporate the large variation in the number of interferers - for example the number of potential interferers, i.e. Xfinity APs surrounding an Optimum AP in Figure 3.3 varies from 0 to 5, even in this small scenario. Identification of active interferers becomes even more challenging considering the reality of tens of networks, non-beaconing APs, and dynamic channel selection.

In this part of the thesis, we do not focus on the messaging interfaces, which is taken up in Chapter 5. The coordination protocol itself can be implemented either in a distributed fashion where each AP sends a message to the neighboring APs of other networks or centrally where the aggregated information is passed through a single interface between networks. Rather, assuming the presence of such side-channel information, we show how each network can optimize client-AP associations to minimize the effects of inter-network interference. While this form of operational cooperation can be devised for optimizing the channel selection (as detailed in Chapter 4), rate allocation, power control, and back-off windows, we first focus on the more tractable case of optimizing client-AP associations. The client-AP association optimization problem can be stated as follows: Given a set of APs that a client can potentially connect to, select the best AP so as to maximize the sum utility of all the clients across the network. Due to its direct impact on both the client experience (in terms of throughput) as well as the network performance (in terms of traffic load), this problem has been approached through both centralized network utility maximization framework [40, 41] and game-theoretic formulations [42]. In particular, we follow the proportional fairness framework developed in [40] for the basic intra-network optimization and enhance it to incorporate inter-network interference.

Operational vs. Access Cooperation: While we propose the operational form of cooperation in this work, it is important to compare it with a simpler form of cooperation which can be termed *access cooperation*. Through access cooperation between two networks, unlike our assumption, clients of one network can join the other network. While this can increase the coverage area for both the networks, we show that unless the two networks also jointly manage their networks, i.e. solve a global optimization problem, network utility cannot be maximized only through access cooperation. In addition, operational cooperation has three distinct advantages over access cooperation: (i) a network does not have to handle authentication for clients from other networks, (ii) networks do not have to over-provision capacity since they do not have to cater to extra clients and (iii) networks can retain the control of sessions, policy, and billing of their clients.

3.3 Motivating Example

Figure 3.4 shows an illustrative example of cooperation gain. Client $C1$ is in communication range of three APs of the same network; and the default 802.11 rule as shown

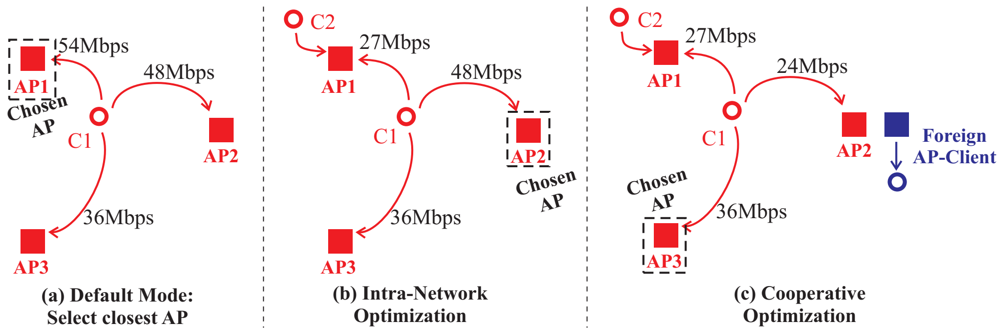

Figure 3.4: Qualitative example of why cooperation is needed: (a) Default mode chooses the least distance AP, (b) Intra-network optimization selects least crowded AP if their distance to the client are comparable, (c) Cooperative optimization selects AP with least internal and external interference

in Figure 3.4(a) is to choose the closest AP (here $AP1$), which gives the highest rate to the client. However, if there is another client $C2$ attached to $AP1$, $AP1$ has to divide its downlink transmission time between the two clients, as in Figure 3.4(b). Assuming proportional fair scheduling, the real throughput that $C1$ gets from $AP1$ is only 27 Mbps. Intra-network optimization through a central controller (e.g., Aruba WLAN controllers [43]) can identify this load imbalance and connect $C1$ to $AP2$ instead and allow the client to get a throughput of 48 Mbps. In doing so, the network controller assumes that $AP2$ has sole control of the channel. However in a multi-network setting, a foreign network may have a nearby AP that shares $AP2$'s channel. CSMA contention leads to approximately equal time share between the two APs, resulting in an actual throughput of only 24 Mbps for $C1$ if connected to $AP2$, as shown in Figure 3.4(c). Cooperative optimization incorporates the effect of APs of other networks and thus connects $C1$ to $AP3$ leading to a throughput of 36 Mbps.

3.4 System Model

We consider a system with N independently operated Wi-Fi networks with U_i and A_i denoting the set of clients and APs in the ith network respectively. Table 3.1 summarizes the notations we use in this work. Binary variables $x_{ij}(k)$ indicate the connection state between the jth client and kth AP of the ith network (1 is connected, 0 if not), while $p_{ij}(k)$ denote the fraction of time provided by the AP to the client. Similarly, $r_{ij}(k)$

Table 3.1: Definition of parameters

Symbol	Meaning
N	No. of Wi-Fi networks
U_i	Set of clients in network i
A_i	Set of access points in network i
R_{cs}	Carrier sense radius (equal for all APs)
R_{int}	Interference radius (equal for all APs)
B_{ik}	Set of co-channel foreign APs within R_{cs} of kth AP of ith network
C_{ik}	Set of co-channel foreign APs outside R_{cs} but within R_{int} of kth AP of ith network
η_{ik}	Number of clients connected to the kth AP of ith network
$r_{ij}(k)$	Wireless PHY rate obtained by the jth client of ith network when connected to the kth AP of that network
$x_{ij}(k)$	Association indicator between the jth client of ith network and its kth AP (value = 0 or 1)
$p_{ij}(k)$	Fraction of time the jth client of ith network gets from its kth AP

denotes the effective bit rate received by the client. Note that while the bit rate values primarily depend on the physical distance between the AP and the client, other factors such as collision induced retransmissions and nature of the rate selection algorithms also impact the bit rate values. In order to make the problem tractable, we only include the distance-dependent component, and in particular, assume $r_{ij}(k)$ to be a step-wise function of the distance between the client and the AP in our simulations. Since air time fraction and rate are relevant only for clients connected to an AP, $p_{ij}(k) = 0$ and $r_{ij}(k) = 0$ whenever the corresponding $x_{ij}(k) = 0$. Thus the jth client of the ith network has an effective downlink rate of $\sum_{k \in A_i} r_{ij}(k)x_{ij}(k)p_{ij}(k)$.

As is common in commercial WLAN controllers [43], each AP employs a proportional fairness policy. Ignoring the protocol overheads and assuming equal priorities for all clients, proportional fairness translates to equal time share between clients in multi-rate WLAN [44]. Thus for the kth AP of the ith network, each of its η_{ik} clients receive a fraction $1/\eta_{ik}$ the APs airtime. We focus on downlink traffic which forms the majority of Wi-Fi data transmission [45] and assume clients always have pending data requests at the AP. This assumption simplifies the estimation of the client rates significantly and is valid in hot-spot deployments where the number of clients is large enough that each client cannot receive its maximum desired data rate.

In order to account for the inter-network interference, we denote the set of co-channel foreign APs within carrier sense range of the kth AP of ith network as B_{ik} and those outside carrier sense but within interference range (potential hidden nodes) as C_{ik}. Each AP has to participate in CSMA and thus shares the channel with co-channel APs within its carrier sense radius. We assume that within each network, frequency planning is such that no two APs within carrier sense distance are assigned the same channel. Thus the kth AP of the ith network has to share its channel with $|B_{ik}|$ other APs, bringing its share of the channel access time fraction to approximately $1/(1 + |B_{ik}|)$ [46]. Further we model the hidden node interference (interference from APs outside the carrier sense range but with signals still strong enough to affect ongoing transmissions) by lowering the channel access time further. We introduce a parameter $\alpha \in [0, 1]$ which captures the average effect of hidden node interference per interferer. The channel access time fraction for the kth AP of the ith network is thus also reduced by a factor of $1/(1 + \alpha|C_{ik}|)$.

Note that an exact model of hidden node interference has been the subject of several past studies [47], [48], and usually requires aggregate interference power calculations which makes the resulting optimization problem extremely intractable. As such, we take a pragmatic approach towards capturing the effect of hidden terminals through the use of the parameter α - a value of 1 implies a hidden node has as much impact on the throughput of a given node as another node within carrier sense range, while a value of 0 implies that the hidden node has negligible impact. In practice, the choice of the α parameter can either be made through probe experiments during the deployment stage or be pre-set to the values derived through testbed measurements [47]. Values of α in the $(0.2, 0.6)$ range satisfy most of our past experiments on the ORBIT testbed [49].

The objective of the intra-network association optimization, given such a model, is to optimize the set of $x_{ij}(k)$ variables for maximum utility which we choose to be one which results in proportional fairness. The choice of *log(.)* or proportional fair utility function is a de facto standard in the current EV-DO, 3G cellular systems, as well as in emerging 4G systems based on LTE and WiMAX and has been shown to provide a good balance between resource utilization and fairness of allocation [40, 41, 50]. For cooperative optimization, each network first ascertains the values of $|B_{ik}|$ and $|C_{ik}|$ for each of its APs through periodic message exchange with other networks. This information is then used to formulate a similar optimization problem as in the case of intra-network optimization. Note however, that by including the hitherto unknown interference components, the cooperative problem formulation now matches the real interference scenario.

3.5 Problem Formulation and Solution

3.5.1 Individual Network Optimization

The intra-network non-cooperative optimization problem formulation is similar to the description in [40]. Since $x_{ij}(k)$ equals 1 only if client j is associated with AP k and channel access time is equally divided between clients connected to an AP, the association optimization within network i can be denoted by:

$$\begin{aligned}
\text{Maximize:} \quad & \sum_{j \in U_i} \log \left(\sum_{k \in A_i} r_{ij}(k) x_{ij}(k) p_{ij}(k) \right) \\
\text{subject to:} \quad & p_{ij}(k) = \frac{1}{\sum_{j' \in U_i} x_{ij'}(k)} \qquad \forall k \in A_i, j \in U_i \\
& \sum_{k \in A_i} x_{ij}(k) = 1 \qquad \forall j \in U_i \\
& x_{ij}(k) \in \{0, 1\} \qquad \forall k \in A_i, j \in U_i
\end{aligned} \tag{3.1}$$

Here the first constraint models the proportional fairness policy of each AP and makes the problem non-linear in $x_{ij}(k)$ while the second constraint along with the binary constraint restricts each client to connect to exactly one AP. Note that the $p_{ij}(k)$ in (3.1) is not the actual time fraction that the client would receive as it does not capture the effect of foreign APs. But without any cooperation, each network has no idea about the number/location of such APs and thus uses this value. Reference [40] shows an efficient approximation algorithm to solve this NP-hard non-linear integer problem for a slightly different problem formulation. This method first requires converting (3.1) to a relaxed discretized linear program without the integrality constraint on $x_{ij}(k)$, i.e., each client is allowed to connect to multiple APs simultaneously. Then the rounding process described by Shmoys and Tardos for the generalized assignment problem [51] is used to arrive at binary values. This polynomial time 2-approximate rounding algorithm thus results in a total utility bounded below by that of the optimal assignment scaled down by a factor of $2 + \epsilon$.

3.5.2 Cooperative Optimization

Extending the above formulation based on the assumptions of equal time sharing MAC and availability of $|B_{ik}|$ and $|C_{ik}|$ values, the global association optimization problem can be written as:

$$
\begin{aligned}
\text{Maximize:}\quad & \sum_{i=1}^{N} \sum_{j \in U_i} \log \left(\sum_{k \in A_i} r_{ij}(k) x_{ij}(k) p_{ij}(k) \right) \\
\text{subject to:}\quad & p_{ij}(k) = \frac{1}{\sum_{j' \in U_i} x_{ij'}(k)} \cdot \frac{1}{(1 + |B_{ik}|)(1 + \alpha |C_{ik}|)} \\
& \qquad\qquad\qquad \forall k \in A_i, j \in U_i, i \in [1, N] \\
& \sum_{k \in A_i} x_{ij}(k) = 1 \qquad \forall j \in U_i, i \in [1, N] \\
& x_{ij}(k) \in \{0, 1\} \qquad \forall k \in A_i, j \in U_i, i \in [1, N]
\end{aligned}
\tag{3.2}
$$

The constraints in (3.2) are a simple extension to those in (3.1) Note here that the first term in $p_{ij}(k)$ is directly dependent on the optimization variables $x_{ij}(k)$. However $|B_{ik}|$ and $|C_{ik}|$ are only dependent on the relative placement of co-channel APs of different networks and are thus constants given a certain topology. So once each network i knows about the $|B_{ik}|$ and $|C_{ik}|$ values for each of its AP k, it can individually solve the association problem. This joint problem can be solved using the same technique as the individual network optimization.

3.6 Simulation Results

Here we present results from detailed analytical simulations that show the benefit of inter-network cooperation for the specific use-case of client-AP association optimization. We compare three association schemes to quantify the gains of cooperation -

- Least Distance: Each client connects to the closest AP of the same network (benchmark case).
- Intra-Network Optimization: Each network optimizes the association pattern of its clients.
- Cooperative Optimization: All networks share information for optimizing the client association.

Note that in all the three cases we assume that the clients belonging to a network can only connect to APs from that network. The discretized linear program was solved using the open source *lpsolve* solver [52]. All the results presented are averaged over 10 simulation runs. We present results for two deployment scenarios: random deployment and clustered deployment as follows.

3.6.1 Random Deployment

Multiple overlapping networks are considered in a 0.5x0.5 sq. km area, which reflect deployment scenarios in urban hot-spot networks, multi-tenant buildings, or airports. Each network has a variable 15-25 APs placed at uniformly randomly selected points. While there is a minimum separation of 50 meters between two APs of the same network, there is no such restriction for APs of different networks. Reasonable frequency planning is assumed - each AP chooses one of the three orthogonal channels in the 2.4 GHz range to minimize the number of co-channel APs. However due to dense deployment of multiple overlapping networks, choosing a completely isolated channel is seldom possible. The carrier sense and interference range thresholds of all devices are set to 215 meters and 250 meters respectively as per the specifications in [53]. Within each network, a planned deployment model is assumed - under this assumption, it is ensured

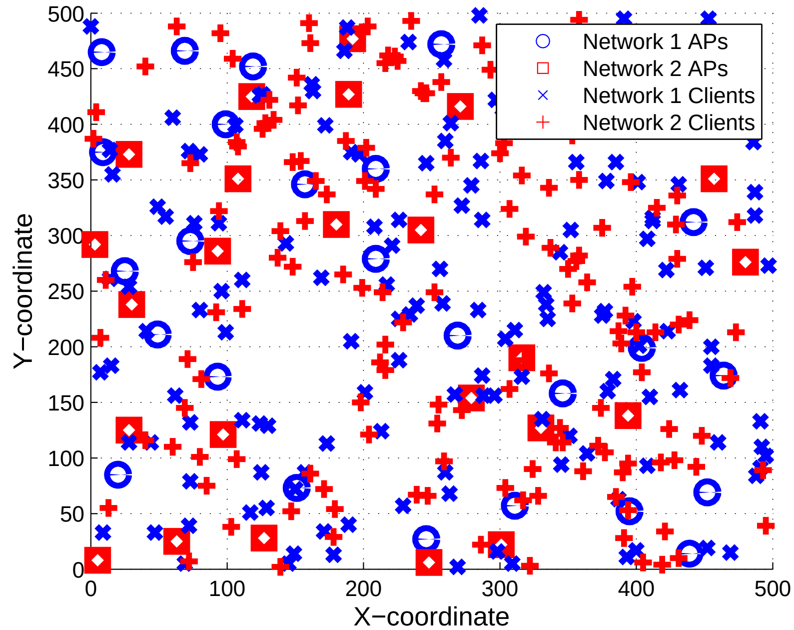

Figure 3.5: Topology showing an instance of 2 Networks, 25 APs per network and 150 clients per network at random points

that two APs from the same network which are within interference range are not on the same channel. Clients are placed at random within the area with the total number of clients of each network set as a parameter. The physical data rates $r_{ij}(k)$ are selected based on the distance between the client j and AP k, also from [53]. The value of the interference scaling parameter α is taken as 0.5. Figure 3.5 shows an instance of the random AP and client placement.

Figure 3.6 shows the cumulative distribution of the client throughputs for all the clients in the system for the topology shown in Figure 3.5. The plot shows that while intra-network optimization improves fairness in client throughput, its effect is limited due to the presence of APs of another network. Cooperative optimization more than doubles the 10 percentile throughput from 230 Kbps to 550 Kbps compared to least distance scheme and shows a 77% gain when compared to the same metric in intra-network optimization. Since the cooperative optimization problem (3.2) decouples into separate problems for each network, utility of each network is individually maximized.

Figure 3.7 further dissects the comparison between intra-network and cooperative

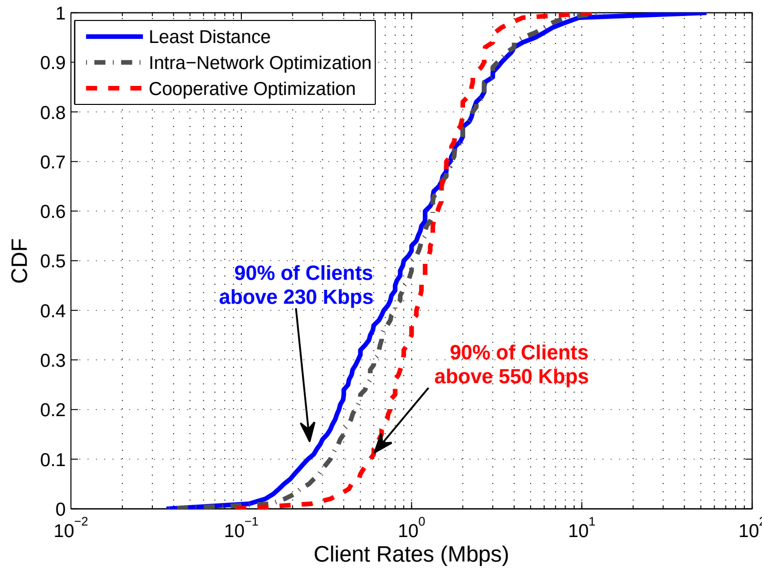

Figure 3.6: CDF of client throughputs for all clients. $N = 2, |A_i| = 25, |U_i| = 150$

optimization schemes. In this figure, clients are arranged in the increasing order of the throughput they get through intra-network optimization. The key observation here is that almost all lowest throughput clients are better off after cooperative optimization, while the accompanying loss in throughput is inflicted primarily on the clients with high throughputs.

Figure 3.8 shows the 10 percentile and mean throughput values for simulations with $N = \{2, 3, 4\}$, 25 APs, and 150 clients. We note that in each of the cases, the 10 percentile throughputs improve by 140-170% with a small 8-10% decrease in the mean throughput. The achievable mean throughput naturally goes down with increasing N due to sharing of the spectrum between a larger number of users. Table 3.2 shows the effect of variations in the number of APs and clients per network for the case of $N = 3$. The key observation here is that the percentage gain brought about due to cooperation increases with AP density, but decreases with client density. The insight from these trends suggests that higher AP densities lead to greater uncertainties that each network has to cope with and thus the information sharing becomes more valuable. However, under a capacity limited regime with large number of users, since all APs are heavily crowded, the relative gain of shifting clients from one AP to another reduces.

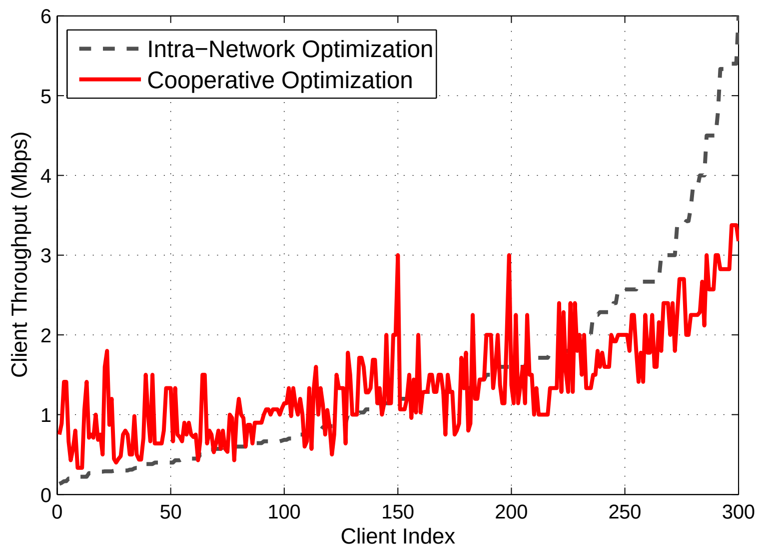

Figure 3.7: Per-client comparison between intra-network and cooperative optimization schemes

3.6.2 Clustered Deployment

Clustered deployments, characterized by a large number of APs placed in a targeted small region are commonly used to serve public places with very high number of peak users, e.g., waiting rooms, mall entrance, etc. In order to study the effects of such topology-specific interference patterns, we considered a clustered topology with two networks. APs of the first network are clustered in three rectangular regions of size 200x200 meters each, while the second network still has a random AP deployment. All other access parameters remain the same as in the random deployment case. Figure 3.9 shows the CDF of the client throughputs for each network. We observe that since network 1 APs are strongly clustered, the relative effect of network 2 APs on its performance is minimal. Hence cooperative optimization does not improve the client throughputs for this network. Conversely, network 1 clusters strongly effect the performance of network 2, thus cooperating between the two networks leads to large gains for network 2.

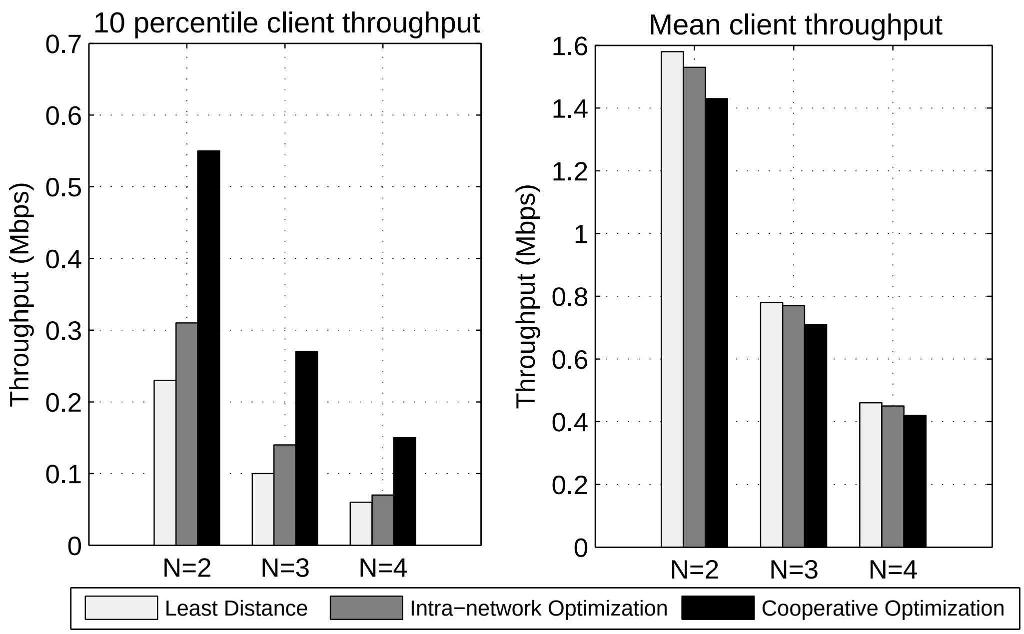

Figure 3.8: 10 percentile and mean client throughput values for varying N with $|A_i| = 25, |U_i| = 150$

Table 3.2: 10 percentile and mean client throughput values for varying number of APs and clients with $N = 3$

$\|A_i\|$	$\|U_i\|$	10 %ile throughput(Mbps)			Mean throughput(Mbps)		
		Least Dist.	Intra Optim.	Coop. Optim.	Least Dist.	Intra Optim.	Coop. Optim.
15	150	0.09	0.13	0.19	0.7	0.66	0.62
25	150	0.1	0.14	0.27	0.78	0.77	0.71
35	150	0.11	0.14	0.31	0.85	0.85	0.77
25	50	0.21	0.33	0.64	1.95	2.17	2
25	150	0.1	0.14	0.27	0.78	0.77	0.71
25	250	0.07	0.09	0.17	0.49	0.47	0.43

3.6.3 Comparison with *Access Coordination*

A simple alternative cooperation scheme in a multi-network scenario is *access coordination* in which two or more networks agree to allow each others' clients to access their networks. Each client can now connect to the nearest AP of any network. In order to compare the operational cooperation scheme proposed in this work with an access coordination scheme, we reuse the topology in Figure 3.5 but allow clients to connect to APs in either network. Figure 3.10 shows the throughput of each client under the three association schemes with client indices arranged in the order of increasing throughput.

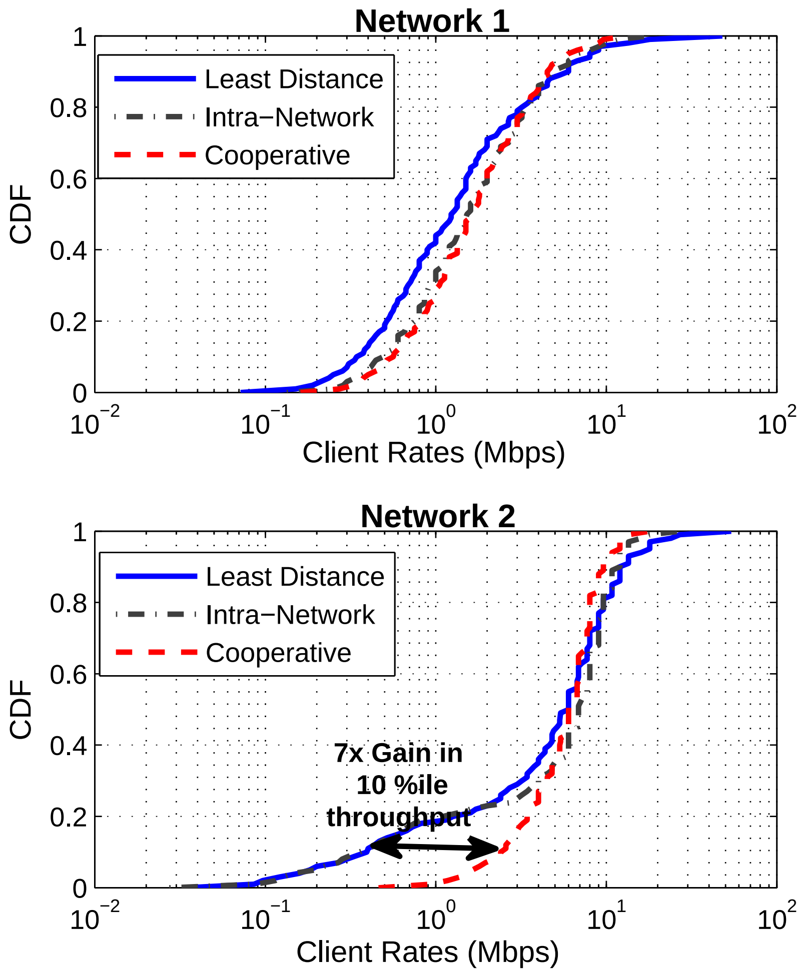

Figure 3.9: CDF of client throughputs for network 1 (clustered APs) and network 2 (randomly deployed APs). $|A_i| = 25, |U_i| = 150$

We note that, access cooperation leads to a decrease in the shortest distance between an AP and a client and thus gives higher throughput for almost all clients. However, since access to more APs does not solve the load balancing problem, operational cooperation results in better performance for more than 2/3rd of the lowest throughput clients.

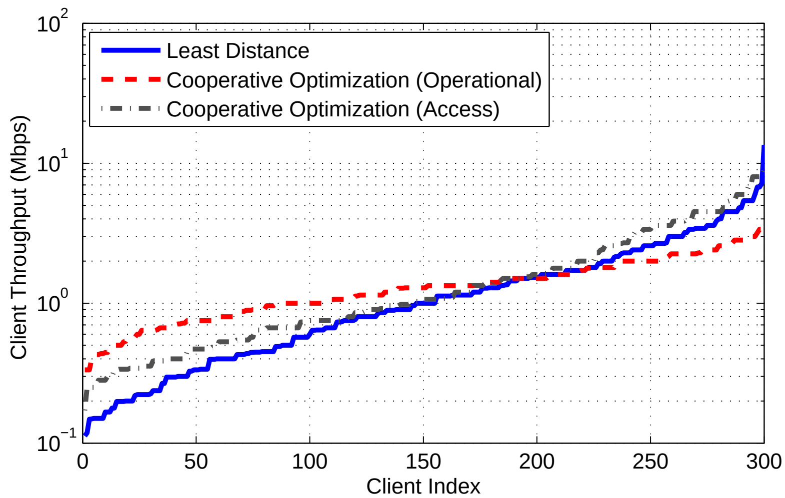

Figure 3.10: Throughput of each client for different association schemes (sorted)

Chapter 4

Network Cooperation for Channel Selection

In this part of the thesis, we focus on understanding the performance of different channel assignment schemes for Wi-Fi networks under assumptions of mixed deployment of cooperative and non-cooperative access points. The motivation behind this problem is the rapid rise of large-scale Wi-Fi deployments by mobile operators and broadband Internet providers, collectively termed as service provider (SP) Wi-Fi [54]. Combined with the already ubiquitous use of Wi-Fi in residential and enterprise settings, the emergence of SP Wi-Fi has led to an interesting mix of deployments where residential, enterprise, and SP APs operate on the same swath of ~80 MHz (and can thus interfere with each other), but enterprise and SP APs are usually centrally managed and can adapt to interference much better than residential APs due to better and more expensive hardware/software. Fig. 4.1 shows the combined percentage of enterprise and SP access points (of the total APs observed) in a 1 sq. km. area of four major US cities, as per the crowd-sourced WiGLE.net database [55]. A clear trend of an increasing fraction of 'managed WLANs' can be observed, especially since the beginning of 2012.

Channel selection is an immediate example of an essential functionality, the implementation of which vastly differs between residential and managed WLANs - most low-cost residential APs either operate on a fixed channel or change channels only upon power cycle, while most enterprise and SP APs incorporate centralized, adaptive channel assignment schemes. Thus in this part of the thesis we address the question: *What is the impact of the increasing density of managed enterprise/service provider APs on the performance of typical residential APs and whether cooperation between multiple groups of APs can help improve the performance?*

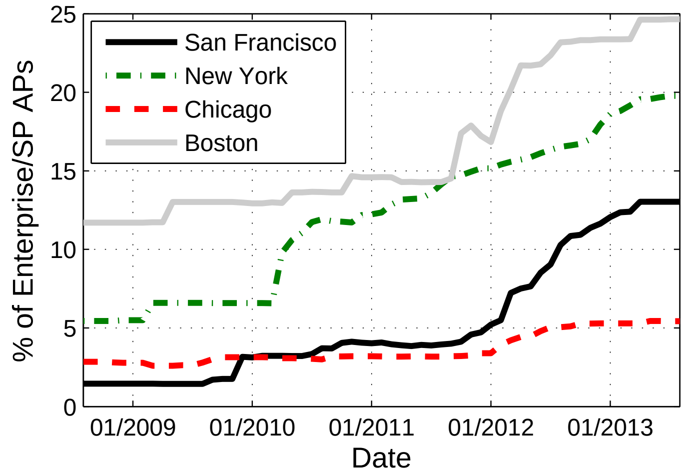

Figure 4.1: Percentage of APs from enterprise/service-provider WLAN vendors out of all observed APs from the WiGLE.net database [55]. An increasing trend can be observed in all regions.

In the process of measuring the performance impact, we show how the Least Congested Channel Search (LCCS) scheme which is most commonly used in low-cost APs [56] overestimates the channel share of an AP due to limited visibility of the neighboring APs. Given n other APs in carrier sense (CS) range on a given channel, an AP using the LCCS scheme implicitly assumes its channel share, if it chooses that channel, to be $1/(n+1)$. In their seminal paper [57] (which has since laid the foundation for throughput-optimal CSMA [58, 59]), Liew et al. proposed an approximate but highly-accurate graph-based technique to calculate the channel share of an AP. The key finding from that work is that the $1/(n+1)$ model is only applicable in 'all-inclusive' settings where all the n neighbors are also in range of each other, i.e. the contention graph has a clique form. Since this is not the case in general, the presence/absence of links between the neighbors has to be taken into account in addition to the number of neighbors, in order to accurately estimate the channel share.

A number of works have shown improvements over the LCCS scheme by utilizing the viewpoint of the complete contention graph and using well-known graph coloring heuristics for channel assignment [60]. In relatively low-density settings, when the number of available channels is enough to color the graph in a conflict-free manner, these heuristics result in optimal performance of all APs. However, when the average degree of the contention graph is much more than the number of channels, the objective of the channel assignment algorithms is usually set to minimize conflicts [60]. Using Liew's Maximum Independent Set (MIS) model for channel-share estimation [57], we show that minimizing the number of conflicts for each AP might result in inefficient assignments. As a solution, we propose an MIS-based 'correction-phase' which can be appended to any centralized channel assignment scheme to decrease the occurrence of starved nodes.

We then measure the performance of a typical centralized channel assignment algorithm in the presence of varying number of independent APs through dense-deployment simulations. The simulation scenarios are designed to reflect the current deployment mix in urban areas (5-25% managed and the rest residential) and also the possible continuation of the trends shown in Fig. 4.1, for example 50-75% managed APs. Different channel assignment schemes are assumed for the low-cost residential APs, in particular, static default, random, and least congested channel schemes. A key finding from the simulations is that, while the trend of increasing percentage of managed APs would improve the overall utilization of the ISM band, at high densities, the existing managed APs would perform worse and the existing residential APs would perform better. The intuition behind this result is that, to an extent, the better performance of the managed APs over residential APs is because of the non-optimal choices made by the latter; as more and more APs improve their resource-usage choices, the potential gains for managed APs is reduced due to the overall capacity of the spectrum being bounded.

4.1 Motivating Example

Fig. 4.2 shows the problem with using only locally observable information. Each node in the graphs shown depicts an AP and an associated client, both implemented using Linux-box nodes of the ORBIT radio testbed [49]. Each client is placed close to its connected AP, while the distance between different APs is changed as per the topology shown. All APs are set to the same channel for this experiment. Throughput tests are done by simultaneously running iperf with saturation UDP traffic between each pair of AP and client. The default ath5k driver with no modifications and auto-rate enabled

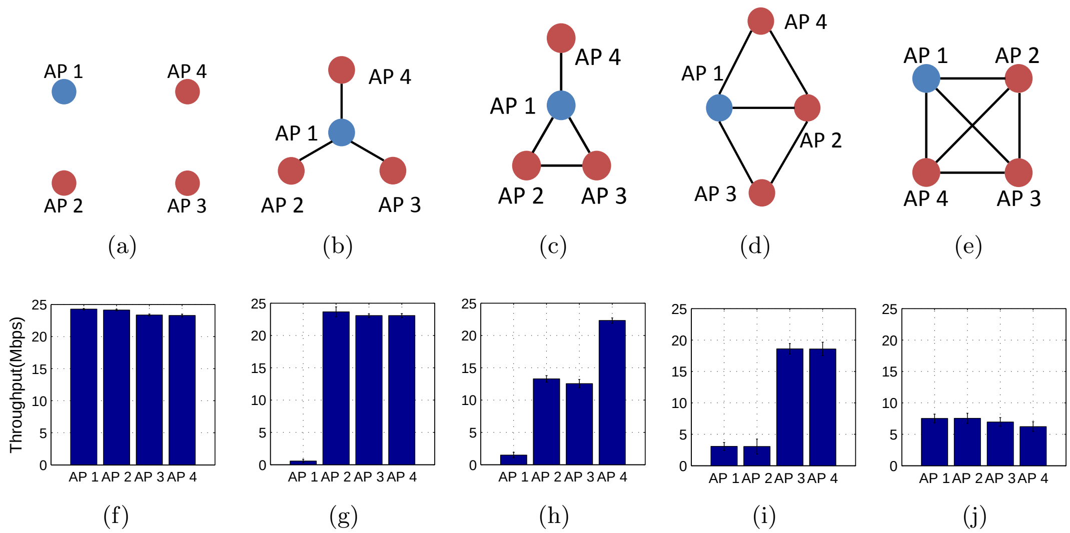

Figure 4.2: Experimental results for the throughput of a four node graph. (a)-(e) shows the interference graph and (f)-(j) shows the corresponding throughput for all nodes. Note that in topologies (b)-(e), AP 1 sees the same no. of co-channel APs yet its throughput varies widely.

is used for all nodes. Fig. 4.2(f) shows the throughput of the four links in isolation, i.e. no APs are in carrier sense range of any other AP. From the local observations that AP 1 can make, all other topologies would appear equivalent - in each case, it would log beacons from 3 APs. However, the channel share that AP 1 gets, and correspondingly the throughput that its connected client gets, widely varies based on the connectivity of its neighbors.

These experimental results show a key benefit of cooperation between networks. If AP 1 and the other APs are parts of two different networks, then without inter-network cooperation, the presence/absence of the links between APs 2, 3, and 4 cannot be ascertained by AP 1 acting alone. This can affect the spectrum access decisions made by this AP in several ways, but in this work, we focus on how this affects the channel assignment decisions.

4.2 Modeling the channel share of an AP

The number of available channels in Wi-Fi is substantially less than what is required to build a conflict-free graph. Hence all practical channel assignment schemes must assign the same channel to multiple APs in range of each other. A channel assignment scheme working with k available channels converts the distance-based graph, i.e. one in which an edge exists between two nodes if they are in carrier sense range of each other irrespective of the operating channel, to k derived-graphs. A node appears in derived-graph i if it has been assigned channel i and a link in the original distance-based graph is transferred to the derived-graph i only if both its end-points are in i. Given such derived-graphs, a general model for the channel share of each AP as per the underlying CSMA protocol has proven to be extremely elusive, except for the case of completely connected graph for which Bianchi's work provides an accurate model [61]. As shown in the motivational example above, the completely connected graph is only one of many possible topologies, and the channel share of an AP in other topologies can widely differ from the values obtained in the case of the clique topology.

4.2.1 Liew's MIS model

Liew et al. [57] proposed the following simple technique to calculate the approximate channel share of each node. Given a contention graph, first calculate its maximum independent sets (MISs) - an independent set is a set of vertices, no two of which are connected by a link in the graph, and the maximum independent sets are such sets with the highest number of elements. The normalized throughput of each node in the graph is then given by the ratio of the number of MISs that node appears in to the total number of MISs. For example, the MISs of the graphs shown in Fig. 4.2 are as follows - (b): {2,3,4}, (c): {2,4} and {3,4}, (d): {3,4}, (e): {1}, {2}, {3}, and {4}. Thus the estimated normalized throughput of the four nodes in sequence are - (b): [0,1,1,1], (c): [0,0.5,0.5,1], (d): [0,0,1,1], (e): [0.25,0.25,0.25,0.25]. While not exactly accurate in all cases, it can be seen that these estimated values closely match the experimental measurements shown in Fig. 4.2

While being derived from a theoretical analysis of the underlying CSMA networks, the key intuition behind the accuracy of the MIS model is that amongst the 2^N possible states comprised of each node of a N node graph being on or off, the CSMA protocol largely favors the 'greedy' states, i.e. the states which result in the maximum number of nodes transmitting simultaneously. Further, all such greedy states are equally probable and thus the throughput of each node is dependent on how many greedy states it appears in, relative to the total number of such states.

4.2.2 Parametric approximation of the MIS model

Although simple to reason with, the problem with utilizing this MIS model is that computing all maximum independent sets of a graph is a classical NP-hard problem with a long standing bound of exponential complexity [62]. As such, we propose the following approximation mechanism to parameterize the balance between computational complexity and desired accuracy.

Since computing the MISs of the complete graph is computationally expensive, we use the same MIS model per node over a neighborhood-graph centered around each node. The neighborhood-graph is defined by a parameter termed *span* which can range from 0 to the diameter of the graph. For a selected span s, the neighborhood-graph of a node i is formed from all the nodes at a graph-distance of less than or equal to s. For each node j at a distance exactly equal to s from node i, all directly connected nodes that are not already included in the neighborhood-graph of i are added to it but the connectivity between such nodes is assumed to be a clique. The process is illustrated in Fig. 4.3 which shows the process for building neighborhood-graphs of different spans around node 1. Note that for the span 0 graph, nodes 2, 3, and 4 are included but links 2-3 and 2-4 are added even though they are not present in the original graph.

The intuition behind the step of clique-formation at the edge of the span is to invoke the standard $1/(n+1)$ model beyond the point of the neighborhood-graph. This results in the computed channel share to be exactly equal to that found through the $1/(n+1)$ model for span 0 and equal to that derived from the MIS model for maximum span. Figs. 4.4(a) and 4.4(b) show the mean error compared to maximum span, and the time required for computation respectively when varying the span from 0 to 2 and the number of nodes in the graph from 20 to 50. All values are averaged over 100 random initiations of the graph. As is clear from these plots, while computing the span 0 or $1/(n+1)$ model is extremely fast, it can result in large errors; increasing the span decreases the error but results in a corresponding increase in the computation time.

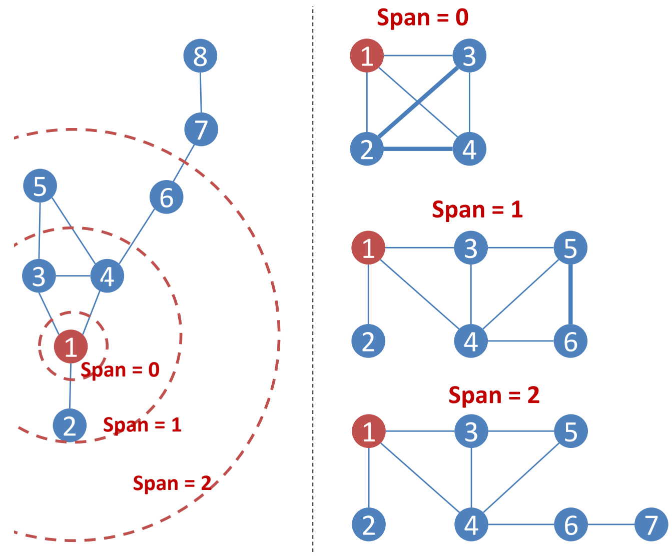

Figure 4.3: Formation of the neighborhood-graph for approximating the MIS model. Span signifies the no. of hops considered for building the graph around each node.

For a given application, the value of the span parameter should be chosen according to the requirements of accuracy and computation time.

4.2.3 MIS based correction phase for channel assignment

Any channel assignment algorithm that assigns channels to nodes sequentially can result in the formation of problematic graphs similar to the ones shown in Fig. 4.2, i.e. one or more nodes might get close to zero share of the channel. Devising a channel assignment algorithm that ensures that no nodes are starved is a difficult problem since each additional assignment can change the structure of the graph and can result in backtracking of the assignments. This would rule out all varieties of sequential greedy algorithms, which form the bulk of those proposed in the literature [60]. Instead, the MIS model can be used to detect and possibly correct the occurrence of starved nodes at the end of any channel assignment algorithm.

During this correction-phase of the algorithm, first the approximate MIS model described above can be used to estimate the channel share of each node. Then, for each starved node, all re-assignments of that node can be tested to check whether the re-assignment would result in the selected node remaining starved or additional nodes getting starved. If a re-assignment for that node which results in an overall improvement

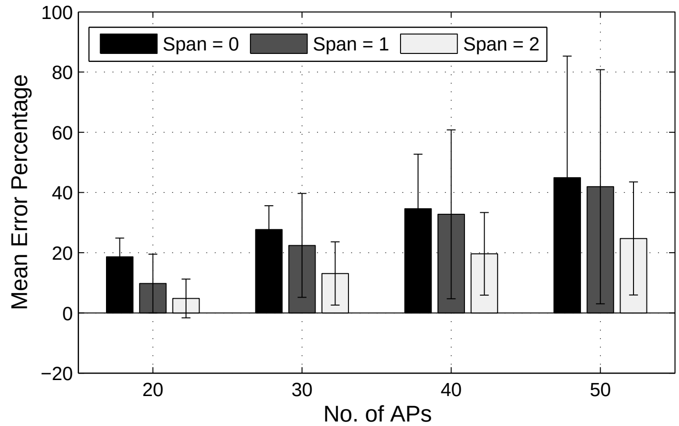

(a) Error percentage compared to max span

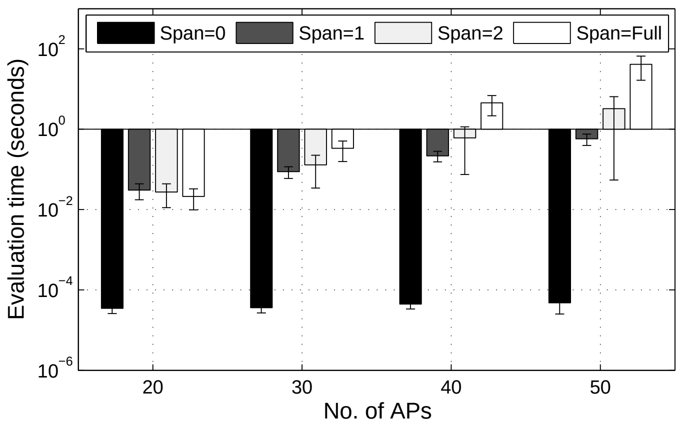

(b) Evaluation times for different spans

Figure 4.4: Performance of the approximate algorithm for graphs of sizes 20 to 50 APs: lower values of span lead to higher error in all cases, but are also computationally faster.

in performance is found, it can be used for the given node.

4.3 Analyzing channel assignments in mixed deployments

The approximate MIS model defined in Sec. 4.2 provides a scalable mechanism to estimate the saturation throughput of APs given the deployment topology and the channel assignment. In this section, we use that model to study the performance of different channel assignment mechanisms under different assumptions about the mix of residential vs. enterprise/hotspot APs. We first benchmark the performance of local and centralized channel assignment schemes in homogeneous settings, i.e. all nodes follow the same algorithm. Next, we consider more realistic settings where different APs might use different mechanisms to set their operating channels.

4.3.1 Simulation description

All results presented in this section are based on MATLAB simulations of dense AP deployments in a 1 sq. km. area. To exactly model the performance perceived by clients in a realistic deployment, the simulation must consider, at the least:

1. Environment-dependent pathloss, shadowing, and multipath, including wall losses
2. The number, placement, and capabilities of client devices
3. The offered load and its variation for each client
4. The policy of the AP for scheduling multiple backlogged clients (note that this is not specified by the 802.11 standard)
5. Capture effect, based on relative signal strength and timing of interfering signals.

Accounting for all these factors can make the simulations intractable, especially when simulating dense deployments. As such, we consider a much simplified simulation setting which retains the qualitative nature of the tradeoffs involved but admittedly misses some of the finer nuances involved in wireless communications in general.

In order to focus on node-starvation and similar network effects, we limit the granularity of simulations to APs, i.e. measuring the throughput achieved at each AP instead of each client. This relieves us from the task of modeling AP load-distribution policy, client locations and capabilities. We assume a downlink saturation scenario, which translates to the assumption of each AP always having one or more connected clients whose data demand is enough to prevent the AP from being idle when it gets the channel. We consider a purely distance-based interference model - if two APs are within carrier sense range (assumed 100 meters), there exists a link between them in the contention graph. Each channel assignment scheme is assumed to be working with three orthogonal channels, as in the 2.4 GHz band - similar results can be obtained for a regime with more number of channels or with non-orthogonal channels by considering a channel overlap dependent sharing model [63].

Metrics - Normalized throughput and starved nodes: We use two key performance metrics throughout this study. The first metric is mean normalized throughput received by an AP - as mentioned above, the throughput 'received by an AP' reflects the combined throughput that all clients connected to the AP would be expected to receive. This is calculated using the approximate MIS model described in Sec. 4.2. The other metric we focus on is the percentage of starved nodes, as estimated from the MIS model. A channel assignment scheme can result in a starved node if the neighborhood contention graph around the node is such that the node receives very less share of the channel. We want to emphasize that although the MIS model would estimate a zero throughput for such a node, it is only applicable in scenarios where all nodes have saturation traffic over a long period of time. Since in reality, some nodes may be intermittently idle, these starved nodes might get access to the channel during the idle-times of other nodes. Nonetheless from a deployment perspective, the starved nodes identified by the model would be topologically vulnerable to performance problems and would offer very low throughput to connected clients during times of peak traffic, i.e. near-saturation load.

4.3.2 Performance in homogeneous settings

We compared the following four different channel assignment schemes for varying density of deployments:

- Random Channel: Each AP independently selects one of the three available channels respectively.
- Local: APs are deployed sequentially and each AP selects the least congested channel from a local viewpoint.
- Centralized: A single entity assigns the channel for all APs using a commonly used, greedy graph coloring heuristic in which the sequence in which nodes are

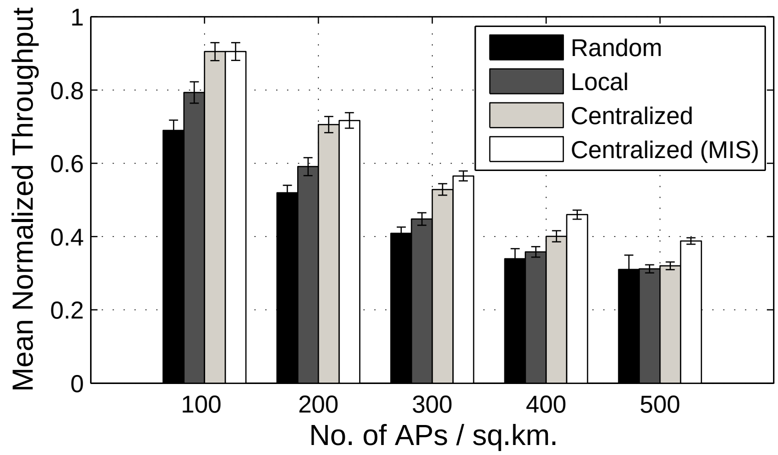

Figure 4.5: Throughput comparison of channel assignment schemes in homogeneous settings. A constant gain of ∼30% can be observed when using MIS compared to Random assignment. Benefit of MIS over Centralized without MIS increases with AP density.

colored is decided based on the number of already colored nodes surrounding each node [64].

- Centralized with MIS correction: The centralized algorithm is followed up with the MIS-based correction phase described in Sec. 4.2.3.

Fig. 4.5 shows the mean normalized throughput at an AP for the four channel assignment schemes listed above. Each point shown in the plots is the average of 1000 simulation runs with AP locations chosen from a uniform random distribution within the simulation area of 1 sq. km. for each run. An interesting insight from this result is that a simple random channel selection performs reasonably well, especially in extremely dense settings since the gains from an optimal choice of channel is vastly reduced if all channels are almost equally crowded. However, in moderate densities (100 - 200 APs/sq.km), the centralized algorithms result in sizable gains of up to 30%. Also interesting to note is that the gain from the MIS-based correction phase increases with density. This arises from the fact in higher density settings, more nodes are prone to being starved due to the network-graph resulting from the centralized channel assignment algorithm. The simple mechanism described in Sec. 4.2.3 for testing alternate channels per starved node thus results in about 20% gain above the centralized scheme in the 500 APs/sq.km scenario.

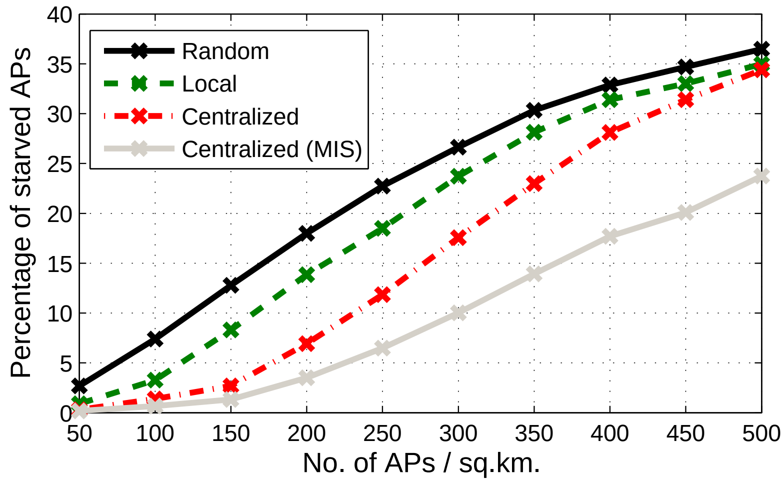

Figure 4.6: Mean percentage of starved nodes (nodes for which MIS estimates zero throughput) as a function of density in homogeneous settings: MIS leads to significant reduction in starved nodes.

The gains from the correction-phase process can be seen more prominently in terms of the starved node metric. Fig. 4.6 shows the mean percentage of starved nodes (out of all nodes in the simulation) for varying densities. The performance of the random, local, and centralized channel assignment schemes as per this metric generally follows the same trends as observed in Fig. 4.5. However the centralized-with-MIS-correction scheme results in substantially less number of starved nodes at all densities.

4.3.3 Mixed deployments with single centralized controller

Next we consider deployments where different APs in range of each other use different channel assignment schemes. In reality, the number of different channel selection algorithms is bounded only by the number of different vendors (we observed more than 500 different vendors in the WiGLE.net dataset used in Fig. 4.1, and hence a myriad of scenarios with various permutations of AP locations and channel assignment schemes can arise. To make this analysis tractable for simulations, we first compare the scenarios in which all APs under consideration either set channels based on a single centralized scheme or follow a different scheme independently. In practice, this assumption translates to the case of a single regional service assigning channels to all APs in the region, except for a varying number of non-subscribers. Alternatively, this is also applicable when the same fraction of APs implement a distributed version of a centralized algorithm by cooperating through a database-service such as the TV White Space database [2].

The ratio of the number of APs following an independent scheme to those under the centralized scheme is varied from all APs belonging to one camp to all APs belonging to the other camp in step size of 5% of the total APs in consideration. For each ratio, 1000 simulation runs are performed where APs are deployed randomly and the independent group is chosen at random after the deployment. The centralized algorithm with MIS based correction is used for the single centralized group, while three different assumptions are made for the independent group - random and local assignment schemes which were benchmarked in Sec. 4.3.2, and a third scheme in which all the independent APs select the same channel (similar to the case of all APs selecting channel 6 by default). Since the local assignment scheme involves scanning all channels locally for counting the number of neighboring APs on each channel, additional assumptions need to be made about the order in which the centralized and local assignments occur, for the simulations that involve local assignments. For this, we assume that the APs in the independent group are turned on sequentially after the centralized group has fixed its channels. However, since we want the independent and centralized groups to reflect the behavior of low-cost residential APs and actively-managed hotspot APs respectively, we assume that the local assignments, once made remain fixed, while the centralized assignments are re-computed after the deployment of the local group.

Figs. 4.7 and 4.8 show both the metrics described for the cases of random and local channel assignments for the independent group. For each plot in these figures, the averages are computed over all the APs in the simulation, i.e. APs from the independent group and the centralized group together, and the shaded regions show the standard deviation around the mean values. The trends across both the cases are similar - there is a gradual increase in performance, in terms of both throughput and number of starved nodes as the ratio of APs acting independently is decreased. In other words, when deployment scenarios evolve from completely independent operation to a completely cooperative regime, throughput gains of the order of 40% and 15%

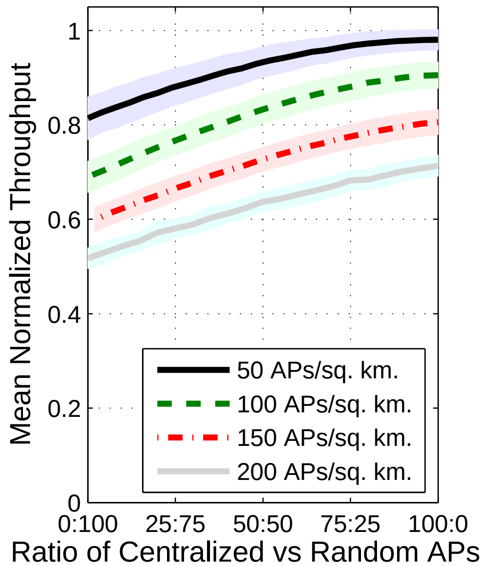

(a) Throughput

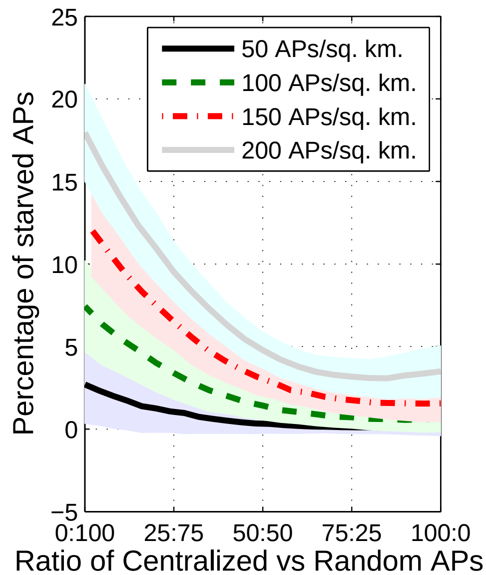

(b) Starved nodes

Figure 4.7: Performance under mixed deployment scenario: Varying ratios of APs following centralized and random assignments

are possible for the random and local assignments respectively. As observed in the case of homogeneous deployments, the gains in terms of alleviating starved nodes are more pronounced - approximately 4x and 3x respectively for the same scenarios as above. Another interesting point to note here is that the performance of the centralized algorithm falls very gracefully in presence of an increasing number of APs which are outside its control, as observed from the smooth nature of all curves.

The extreme scenario of all independent nodes choosing the exact same channel is shown in Fig. 4.9. As can be expected, the gains from all nodes using a centralized algorithm compared to individual operation are more here - around 2x in terms of mean throughput and 9x in terms of percentage of starved nodes.

The results above suggest that if the deployment trends shown in Fig. 4.1 continues, i.e. the percentage of more actively managed cooperating APs increases, the overall performance of APs will improve. However, when the same results are broken into the performance of the independent APs and the cooperating APs measured separately, a more nuanced view emerges. Fig. 4.10 shows the breakup of the mean throughput between the two groups for a particular simulation - mix of centralized and same channel

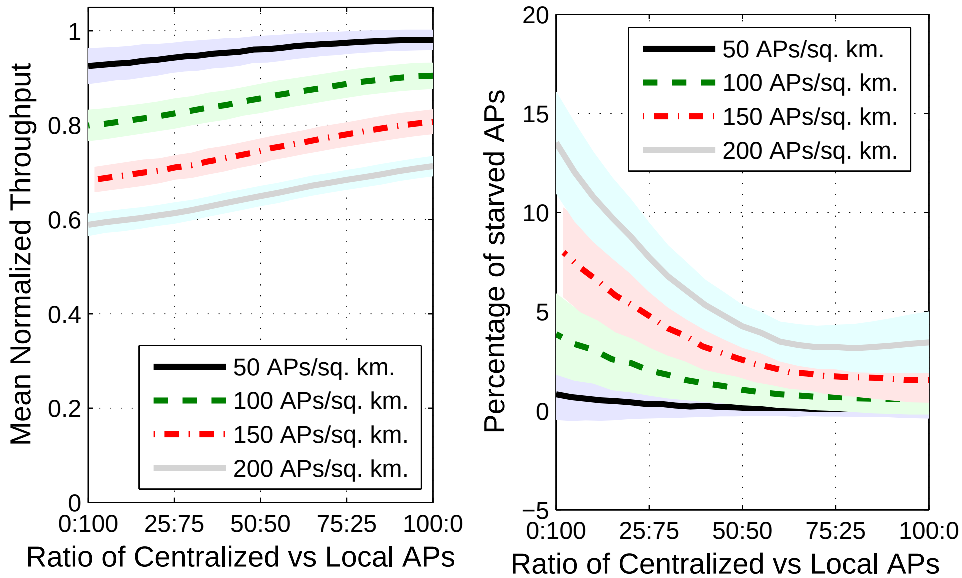

(a) Throughput

(b) Starved nodes

Figure 4.8: Performance under mixed deployment scenario: Varying ratios of APs following centralized and local assignments

APs with a density of 200 APs/sq. km. This shows that for a given density, as the percentage of independent APs decreases, the performance of the centralized APs also decreases whereas that of the independent APs increases. This somewhat counter-intuitive result arises from the fact that when only a few APs make a smart choice about the channel in the presence of many 'dumb' APs, they get more room to optimize the channel selection process.

4.4 Experimental validation of results

The simplified simulation setup presented above, with distance-based binary interference model and no clients, allowed us to analyze the behavior of different channel assignment schemes in dense settings. In this section, we show that the results obtained under these assumptions qualitatively match to those from similar real-world experiments.

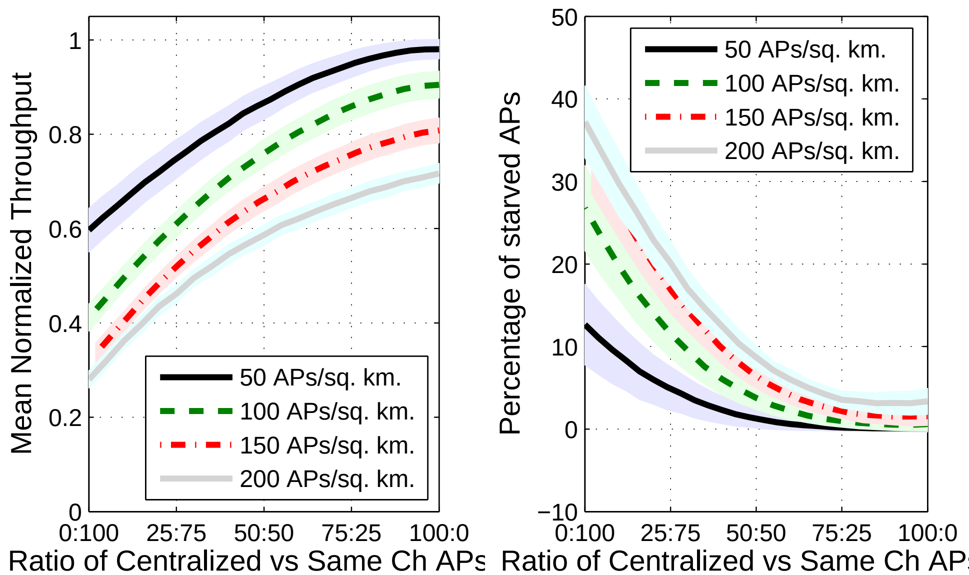

(a) Throughput (b) Starved nodes

Figure 4.9: Performance under mixed deployment scenario: Varying ratios of APs following centralized and same assignments

In the simulations, APs are dropped randomly in a given square area and the interference graph is then computed based on pair-wise distance calculations. Since randomly placing Wi-Fi devices in a physical space is not feasible, we create different interference topologies using an eight-node attenuator system described in the following section. This limits us to experiments with maximum eight APs, but given a certain number of nodes, it allows us to iterate over every possible interference graph that can arise from that many nodes. We limit the number of available channels to two since in an experiment with eight or less APs, having more channels would lead to mostly uninteresting results in which all APs would get an exclusive orthogonal channel to use. Other specific settings of the experiments are summarized in Table 4.1.

Using this setup, we perform two sets of experiments - one to compare the performance of the Random, Local, and Centralized MIS schemes for channel assignments, and the other to validate the observation of increasing fraction of centralized APs leading to decreasing performance for them, shown in Fig. 4.10.

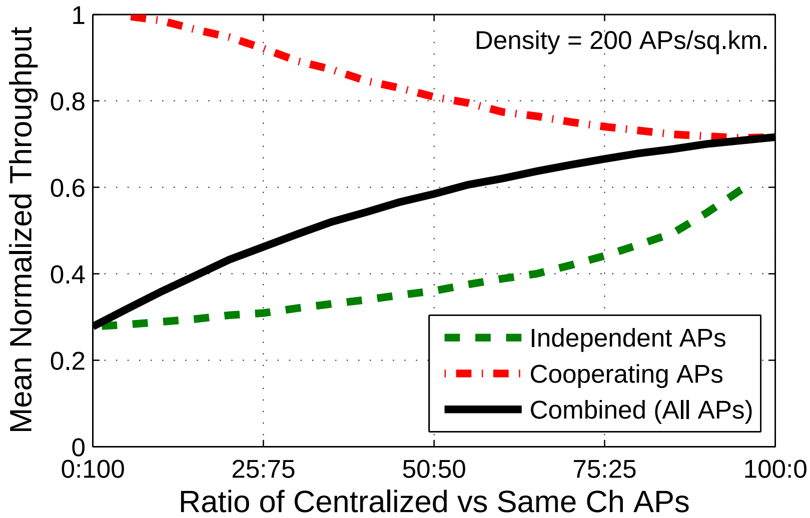

Figure 4.10: Breakdown of the performance gain: Increasing percentage of centralized APs leads to increasing performance of the independent APs.

Table 4.1: Experiment configurations

Attribute	**Value**
Node Hardware	Intel Core i7 CPU @2.93GHz
Operating System	Linux 3.2.0-34
Wireless Interfaces	Atheros 5212/5213 mini-PCI and Intel 6250 mini-PCIe Wi-Fi/WiMax
Wireless Driver	Ath5k (srcversion: FA2500C2D3C1A3B6094E37C) iwlwifi (srcversion: C9C876E115EE7BFFAFB2FA7)
Transmit Power	17 dBm
PHY/MAC/Freq. Used	IEEE 802.11g Channels 1,11
PHY Link Speed	Minstrel adaptive rate control algorithm [65], up to 54 Mbps
Traffic	Full-buffer TCP via Iperf [66]

4.4.1 Creating interference topologies

For all experiments reported in this section, we use an eight-node attenuator system available as a part of the ORBIT lab facility [49]. This measurement system consists of eight Linux boxes, each of which has an Atheros 5212/5213 mini-PCI card and an Intel 6250 mini-PCIe 802.11/802.16 card. The nodes are enclosed in a Ramsey Electronics RF enclosure [67] that provides 80 dB of isolation whereas all the input/output ports of the wireless cards are connected through a programmable attenuator as shown in Fig. 4.11. Further details about this setup is available at [68].

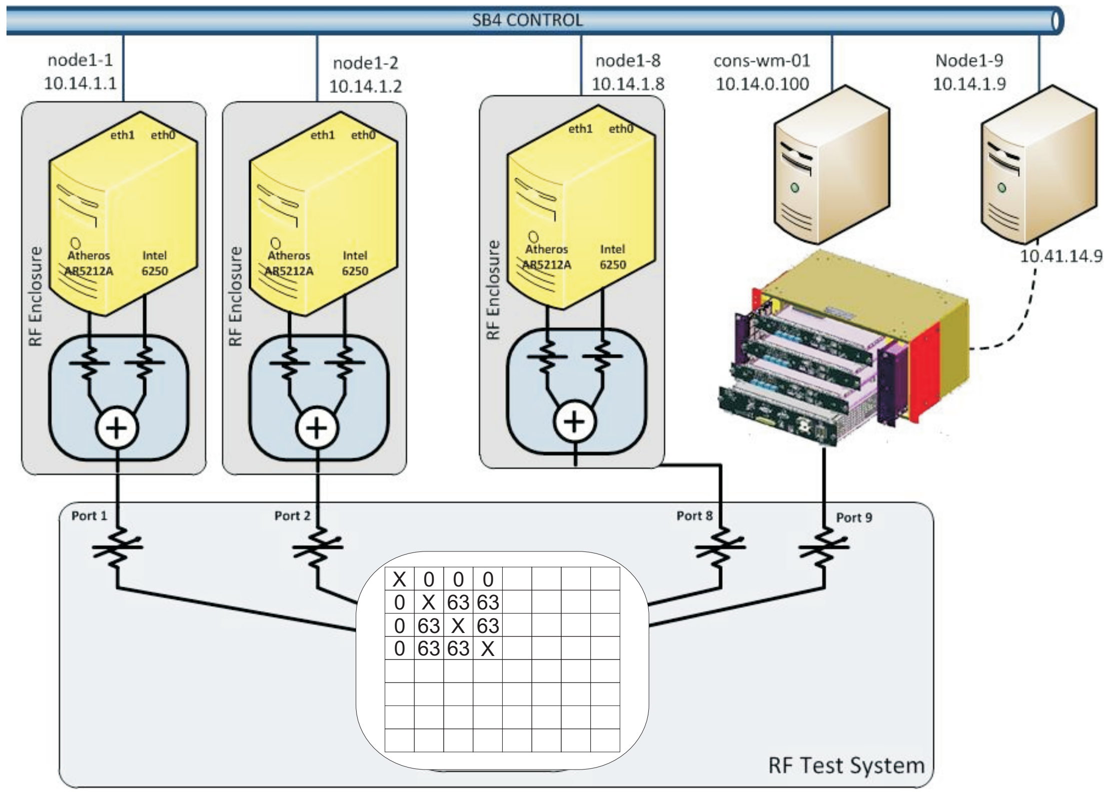

Figure 4.11: Eight node attenuator system

This setup provides a way to create arbitrary topologies (within the operating range of the attenuators) in a stable manner. The total attenuation between any two nodes consists of a fixed part and a programmable part. The minimum attenuation that can be set between two nodes is 70 dB when the programmable part is turned to 0. The upper limit of the programmable attenuators is 63 dB, which provides enough attenuation for isolating the two ends of the link. Thus for example to create the topology shown in Fig. 4.2(b), the attenuation between AP1 and each of APs 2, 3, and 4 is set to 0 dB while that between APs 2 and 3, APs 2 and 4, and APs 3 and 4, are set to 63 dB, as shown in the matrix in Fig. 4.11.

In all experiments, each AP is associated with a single client. As mentioned earlier, each Linux box has two 802.11 cards, and we co-locate the client and the AP on the same machine, each using a different wireless card. In doing so, we setup the NAT and IP settings on the Linux box in a way to ensure over-the-air transmission of packets between the client and the AP instead of direct kernel routing. This results in the client-AP communication being affected by the CSMA environment of the node, which is what we want in the experiments. The reasoning behind this single node client-AP setup is that since we are limited to an eight-node attenuator system, we can perform experiments with twice as many APs this way than when using separate nodes for clients and APs. In addition, it also makes the experimental setup comparable to the simulation setup in which we only measured the normalized throughput received by an AP, as it removes any throughput variations resulting from client-locations. The limitation of such a setup is that it reflects an extreme case of real-world deployment, i.e. clients located very close to the AP. However, this does allow us to understand one aspect of the problem (namely the interference graph between the APs) in greater detail by removing effects from other factors.

4.4.2 Comparison between channel assignment schemes

In order to validate the performance of different channel assignment schemes in homogeneous settings, as presented in Section 4.3.2, we undertook a series of five-node experiments. Using the attenuator system described above, we created all possible five-AP interference topologies. The number of such topologies is 21 under the condition that the resulting graph is connected [69]. The connectivity requirement makes the problem of determining the number of such topologies a combinatorial problem with no closed-form results. The solution forms series A001349 in Sloane's encyclopedia of integer sequences [70], with the first ten entries as follows: 1, 1, 2, 6, 21, 112, 853, 11117, 261080, 11716571.

Fig. 4.12 shows the graph of all topologies. Note that the placement of each node in the graphs shown in this figure is fixed; in a real-world scenario, the relative distance between the nodes will determine the resulting topology. For each topology, we ran 5 runs each of three different channel assignment scheme - Same, Local, and Centralized using MIS. Other details about the experiment setup are listed in Table 4.1.

Fig. 4.13 shows the mean throughput achieved at each node in each topology for all the three different channel assignment scheme. From the results, it can be observed that while MIS is an improvement over local and same channel assignments, the actual

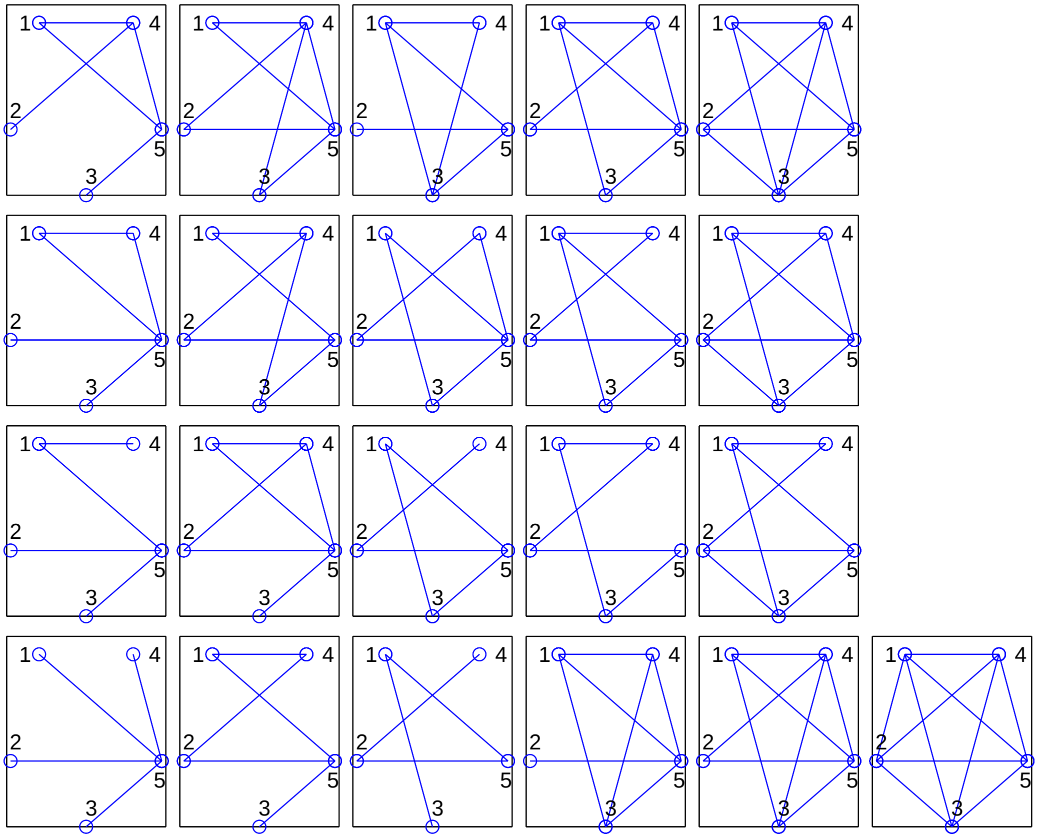

Figure 4.12: Graph of all 21 five-node connected topologies

improvement in performance is very topology-specific. To get a better understanding of the gains, the average throughput of all nodes in a topology is shown in Fig. 4.14. The figure also shows the averaged numbers over all topologies. As can be seen from the figure, in all topologies except topology 1, local channel assignment results in better average throughput than same channel setting. Also, in all topologies except topology 12, the centralized MIS provides further improvement over local assignment. Across all nodes and all topologies, local and centralized MIS assignments provides 30% and 60% improvements over same channel assignment respectively.

In order to further explore the relation between topologies and the potential of centralized channel assignment scheme, we show the degree of each node along with the percentage improvement in its throughput between the centralized and same channel

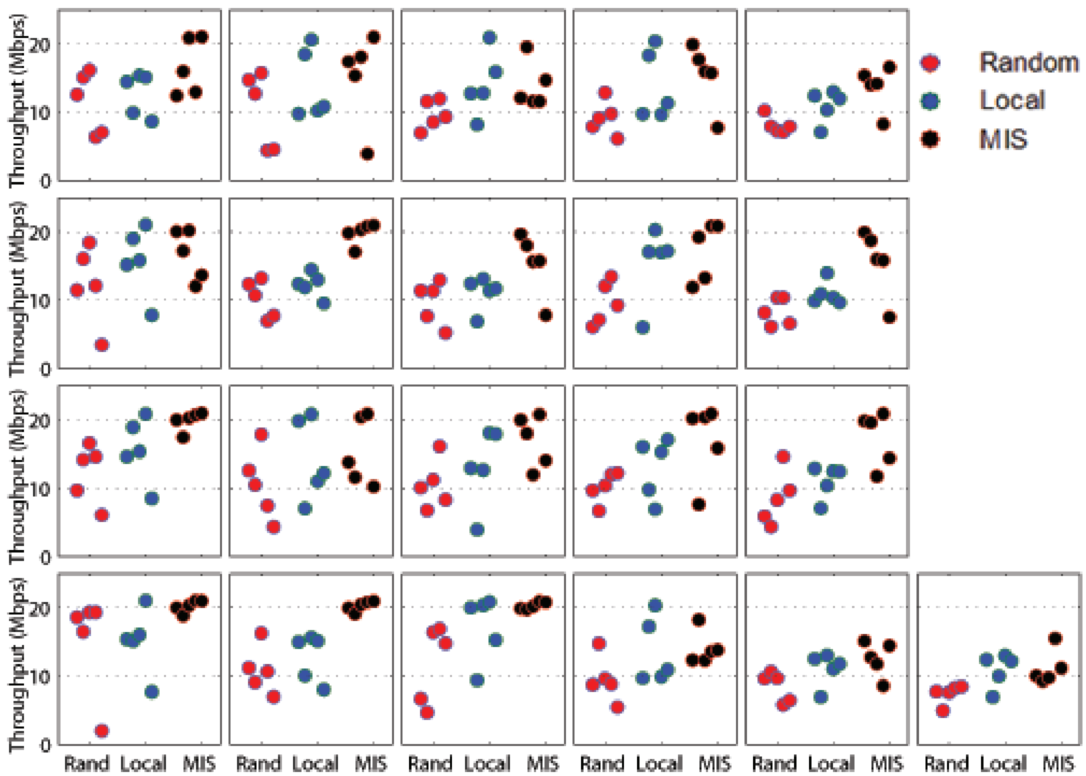

Figure 4.13: Throughput of each node in each topology for three cases - same, local, and centralized-MIS channel assignment schemes. Each plot shows one topology; horizontal axis only spreads out points and is not associated with any experimental parameter.

cases in Fig. 4.15. The inference from the figure is that the spread of the performance improvement is in general positively correlated with the node degree. However, the dense placement of points in the 0-100% region of the horizontal axis for nodes of all degrees indicates that a specific relation between the graph structure and performance seems unlikely. In the same lines as this analysis, Fig. 4.16 plots the total number of links in the topology and the overall percentage improvement for each of the 21 five-node topologies. The overall uncorrelated placement of node along the two axis reinforces the point that performance improvement is very topology-dependent.

4.4.3 Mix of centralized and independent APs

Next, we performed a different set of experiments with eight nodes in order to verify the relation between percentage of centralized nodes and their performance, observed

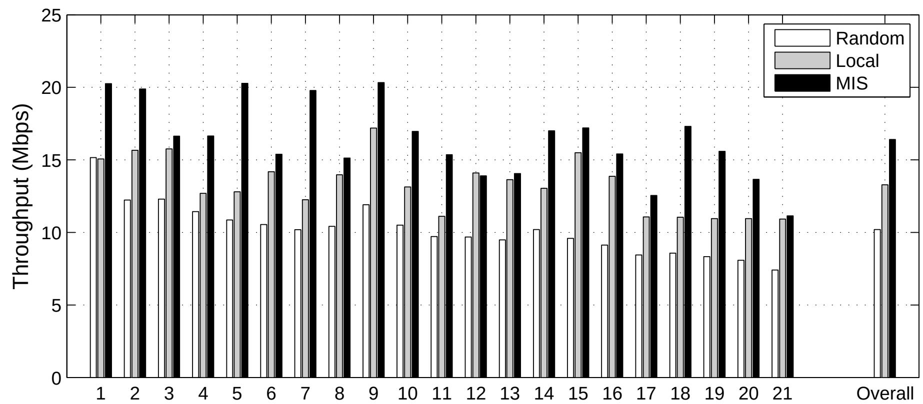

Figure 4.14: Average throughput across all nodes in each topology

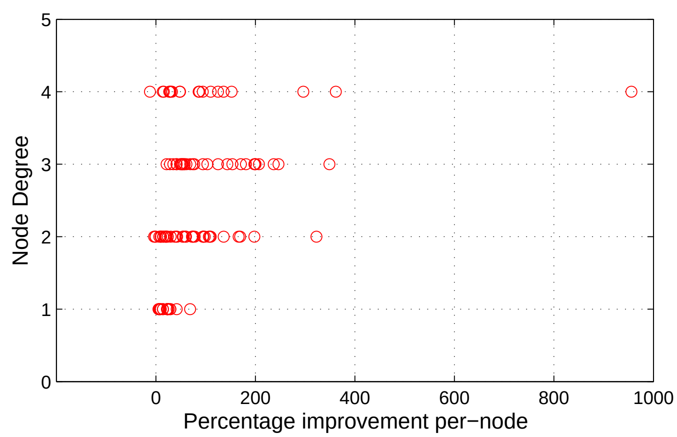

Figure 4.15: Scatter plot of performance improvement vs. node degree

in Fig. 4.10. For these experiments, we first randomly select 100 topologies out of the possible 11,117 connected topologies [69]. The reason for this sub-sampling is that each experiment takes a considerable amount of time and the error margins obtained by averaging over 100 topologies seemed within acceptable bounds. For each topology, we run 9 different experiments in which we vary the number of APs choosing a constant fixed channel from 0 to 8 in increments of 1. Channels for the remaining APs in each experiment are assigned using the centralized-MIS scheme.

Fig. 4.17 shows the performance of independent nodes, centralized nodes and that of all nodes combined. As observed in the simulation results, this figure indicates that as the fraction of nodes that are under centralized control increases, the room for improvement in their performance decreases. Specifically, across 100 different topologies,

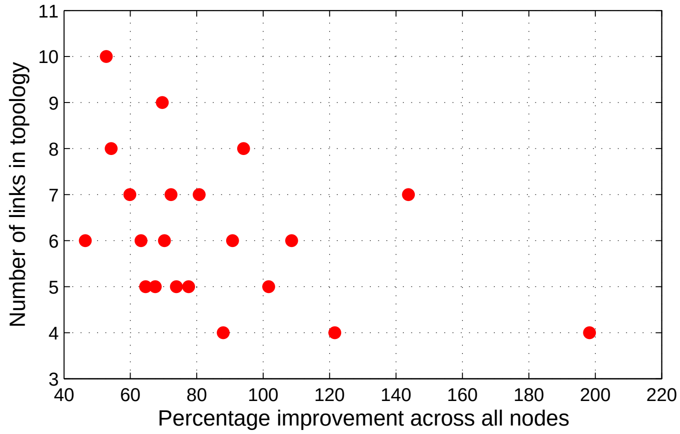

Figure 4.16: Scatter plot of number of links in topology vs. avg. performance gain

the throughput obtained by a single 'smart' AP in presence of 7 other APs all of which select the same channel, is on average 20 Mbps. Whereas, when all 8 APs are under the same centralized channel assignment scheme, the average throughput of each AP is about 15 Mbps, indicating a ∼25% drop in performance.

Figs. 4.18 and 4.19 show two example topologies from the 100 topologies selected along with the performance of independent APs (choosing the same channel), and the centralized APs for all possible mix of APs. In the first topology, when all nodes are under the centralized channel assignment scheme, the two channels available are sufficient to ensure orthogonal operation for each AP. Hence there is no drop in throughput as the fraction of centralized APs increase. On the other hand, the second example shows a case of when two channels are not adequate, and hence the average throughput of the centralized APs decreases as their fraction increases.

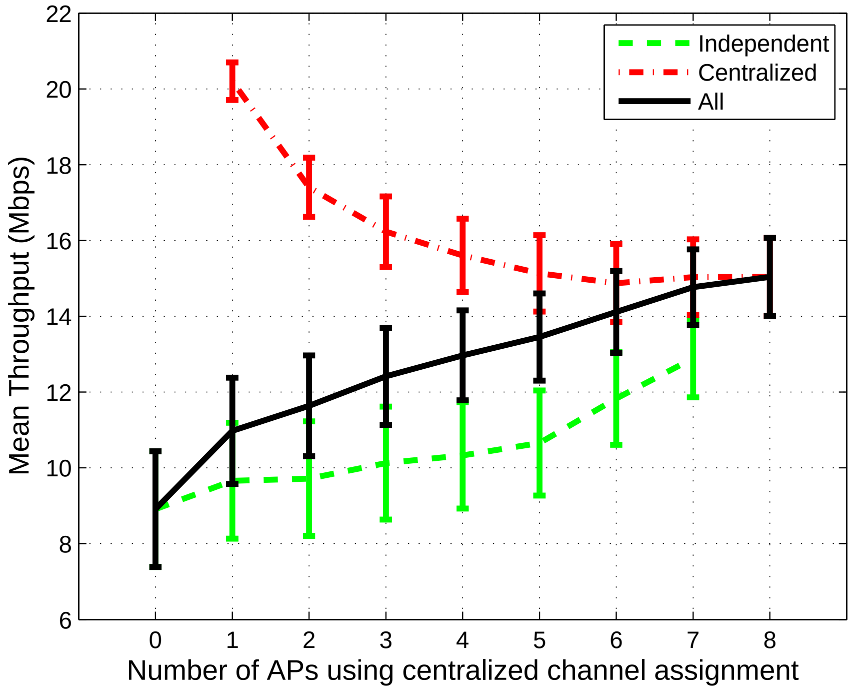

Figure 4.17: Performance of independent and centralized nodes in mixed-deployment scenarios

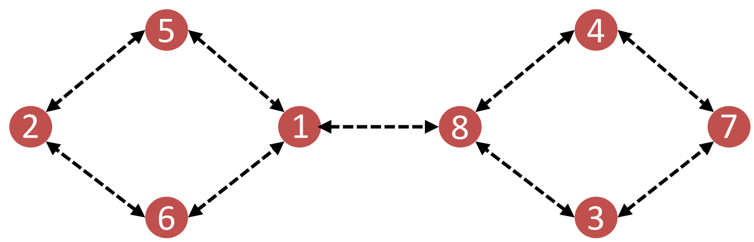

(a) Interference Graph

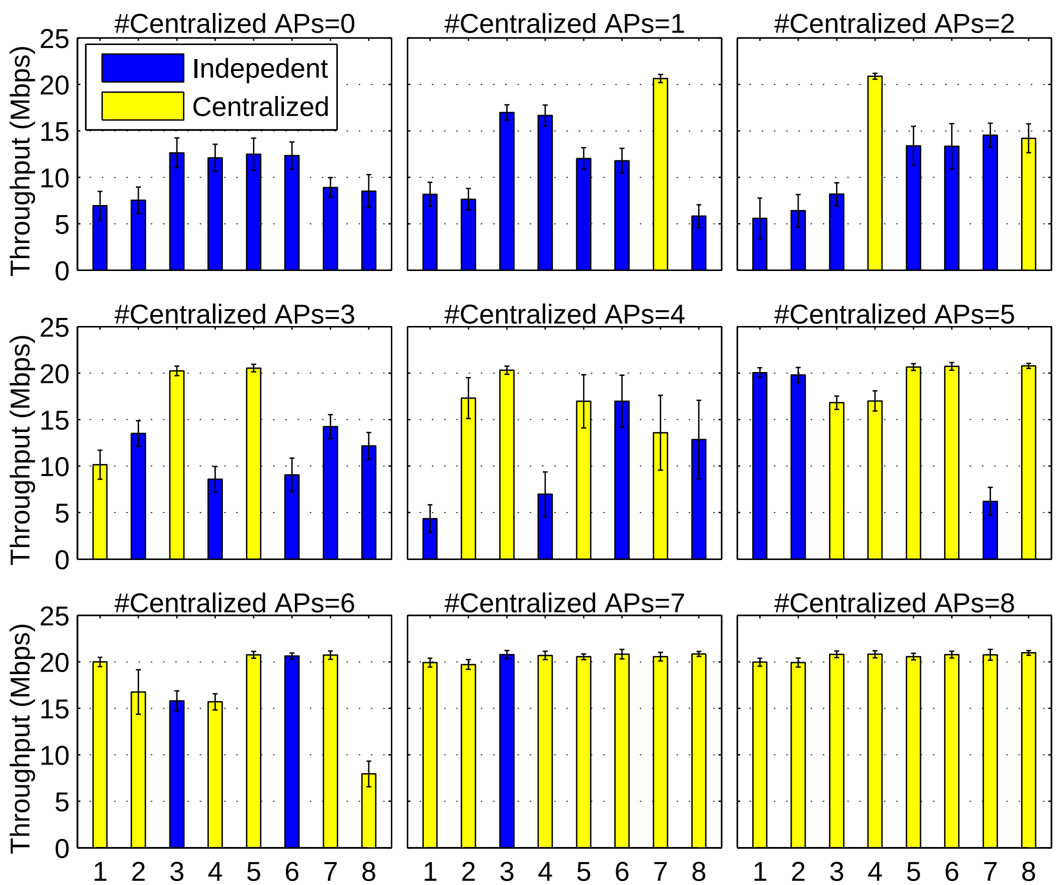

(b) Throughput for various mix of Centralized/Independent APs

Figure 4.18: Example 1 topology and the performance in all possible mixes - from 0 to all centralized APs. Increasing fraction of centralized APs do not result in reduced throughput.

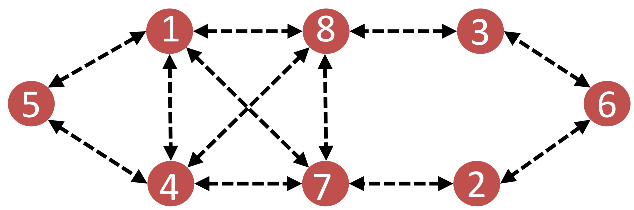

(a) Interference Graph

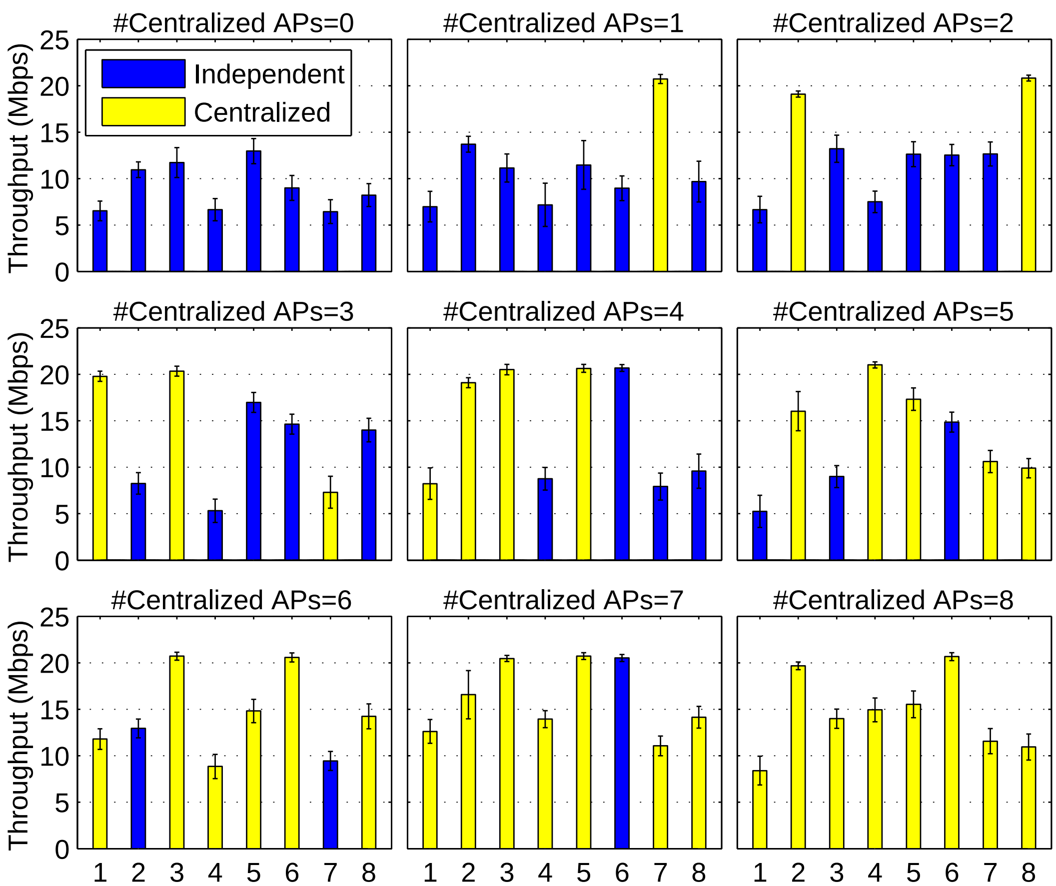

(b) Throughput for various mix of Centralized/Independent APs

Figure 4.19: Example 2 topology and the performance in all possible mixes - from 0 to all centralized APs. Increasing fraction of centralized APs result in lower throughput for them.

Chapter 5

Realizing Inter-network Cooperation

Having shown the benefits of cooperation between wireless networks through two separate explorations of the solution space in Chapters 3 and 4, we next focus on the protocol and software implementation aspects of such cooperation mechanisms. In this chapter we describe efforts towards defining an inter-network spectrum coordination application programming interface (API) and our ideas on leveraging the key principles of software defined networking (SDN) for building a flexible control plane in wireless networks.

The API defined in this work is a set of simple, radio technology neutral, interaction procedures through which controller-entities in different networks can communicate and cooperate. The proposed design is intended as a proof-of-concept baseline over which additional features can later be added to suit specific deployment requirements. A controller-entity, in this context, could be a hardware controller, which is often the case in large WLAN deployments, or software programs residing on either network devices or in a remote location, that controls the choice of spectrum access parameters used in the network. For example, in case of a residential AP, the controller-entity could be residing in the AP itself, and it controls the channel, transmit power, rate selection scheme, association scheme, and other MAC parameters of the AP. The API design is guided by the following example use cases:

- **Information sharing for radio control:** Two or more wireless networks with overlapping spectrum usage in terms of space, frequency, and time, inform each other about the radio parameters being used. The control of each network is retained by the network but is just done in cognizance of the information provided by neighboring networks. This use-case includes the scenarios described in

Chapters 3 and 4.

- **Delegation of radio control:** One network delegates the control functionalities of specific parameters to another network or to an aggregate or local area server. Such a service would be similar to the spectrum server and spectrum-management-as-a-network-service ideas discussed in Chapter 2.

- **Selectively turning off APs to reduce interference:** One network turns off some of its APs and another network hosts virtual APs acting on behalf of the first network. This situation might be desirable in high density areas where the most effective way of reducing interference is reducing the number of terminals contesting for the channel.

We adhere to two general principles that would be desirable in an inter-network spectrum coordination API. Firstly, the interactions should allow for flexibility in the extent to which each network desires or is capable of sharing information about. For example, due to competitive reasons, a commercial WLAN provider might not want to disclose fine-grained information about the number of APs or clients it has to other networks. Secondly, the decisions on the level of information sharing and transfer of control should be based on measurable parameters. In scenarios with greater interference, more information might be required for effective spectrum management. The issues of fairness, security, and privacy are also very important in such an inter-system interaction process. However, we do not consider the presence of malicious entities in this basic API design.

5.1 Key Steps

Fig. 5.1 shows the four basic phases of the cooperation API. The first step of the communication flow is the discovery of controller-entities of neighboring networks. As discussed in Chapter 2, there are several possible mechanisms through which neighboring controllers can discover each other, some already being implemented with support from the IETF PAWS specifications [12]. Web-based databases, and routing layer geographical flooding of special discovery packets are other alternative discovery mechanisms.

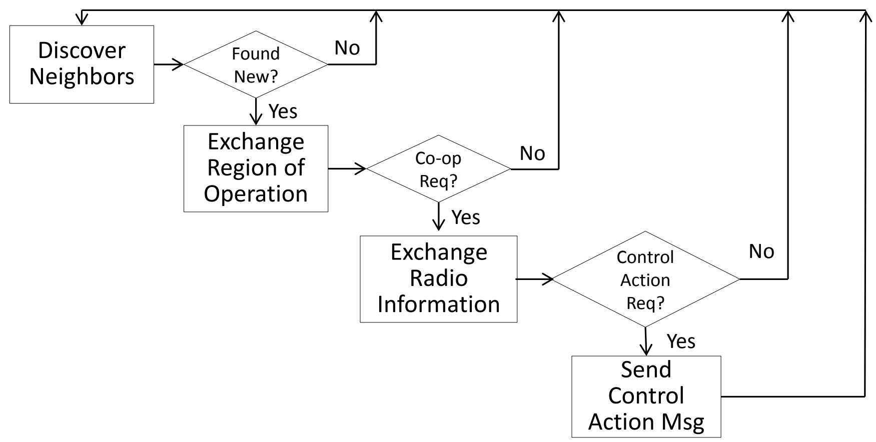

Figure 5.1: Key steps in inter-network spectrum coordination

Further, each network could also have prior off-line information about how to reach specific neighbors. It is unlikely that a single solution among these options would prove to be ideal in all deployment scenarios, and as such we consider discovery outside of the scope of this work.

Discovery is followed by the exchange of a coarse-grained region of operation message between the communicating neighbors. The structure of this exchange is shown in Fig. 5.2. It consists of the list of identifiers associated with the network (e.g. SSIDs in WLANs), and a tuple of geo-coordinates and operating radii that determines the region of operation of the network. The network can choose to specify a single region encompassing several APs or list the regions individually. The purpose of this exchange is to determine whether or not a network needs to coordinate with a discovered neighbor since most discovery mechanisms will not guarantee strict operating boundaries when determining neighbors. Once a network receives this information from a neighboring network, it can combine it with its own measurements (e.g. beacon reports from its own APs about visible SSIDs), and with its knowledge of its own region of operation, to determine the extent of cooperation required.

Based on the exchange of region of operation information, if two networks decide to cooperate, the next step involves the exchange of specific radio information. Since in general, different networks employ different control mechanisms which might require

RegionStruct:

Network Name	List of SSIDs
Region of Operation	List of < X, Y, Z, R > values

RadioInfoStruct :

Per-Transceiver Information	
Identification	Unique ID for the radio
Frequency Domain	Start-stop frequencies & corresponding Tx power level
Time Domain	Avg. Utilization Fraction
Spatial Domain	X, Y, Z Coordinates

Basic Message Types

```
1. Hello
2. Region Request
3. Region Response
   (List of RegionStruct)
4. Cooperation Request
   (Cooperation type)
5. Cooperation Response
   (Accept/Reject)
6. Radio Info Request
   (List of RegionStruct)
7. Radio Info Response
   (List of RadioInfoStruct)
8. Change Param Request
   (List:{RadioID,Param,Val})
9. Change Param Response
   (Accept/Reject)
```

Figure 5.2: Data structure and message types for spectrum API

different kinds of inputs, the radio information exchange is formatted in a request-response fashion. A requesting network specifies either a region of interest or specific ID of the radio transceiver it needs information about. Based on its policies for information sharing, the neighboring network responds with key parameters - an example of the radio information structure is shown in Fig. 5.2.

Finally, based on the information received and the type of coordination selected, a network can request a neighboring network about changing certain parameters. This is done using change parameter request messages which specifies the identity or region of the radio device, the parameter, and the value (e.g. radio1, channel, 1). A network receiving this message can decide whether or not to accept the requests based on its own measurements and policies. Fig. 5.2 also shows the nine basic message types needed for the steps described here.

5.2 Proof-of-concept validation

We performed a proof-of-concept validation of the API described above in a two-network Wi-Fi scenario, with each network consisting of 4 APs and 4 clients. The setup used for this experiment is the same as the one described in Sec. 4.4. It consists of 8 Linux boxes with two radio devices each that act as an AP and its connected client. The wireless interfaces of the nodes are connected through a programmable attenuator which enables

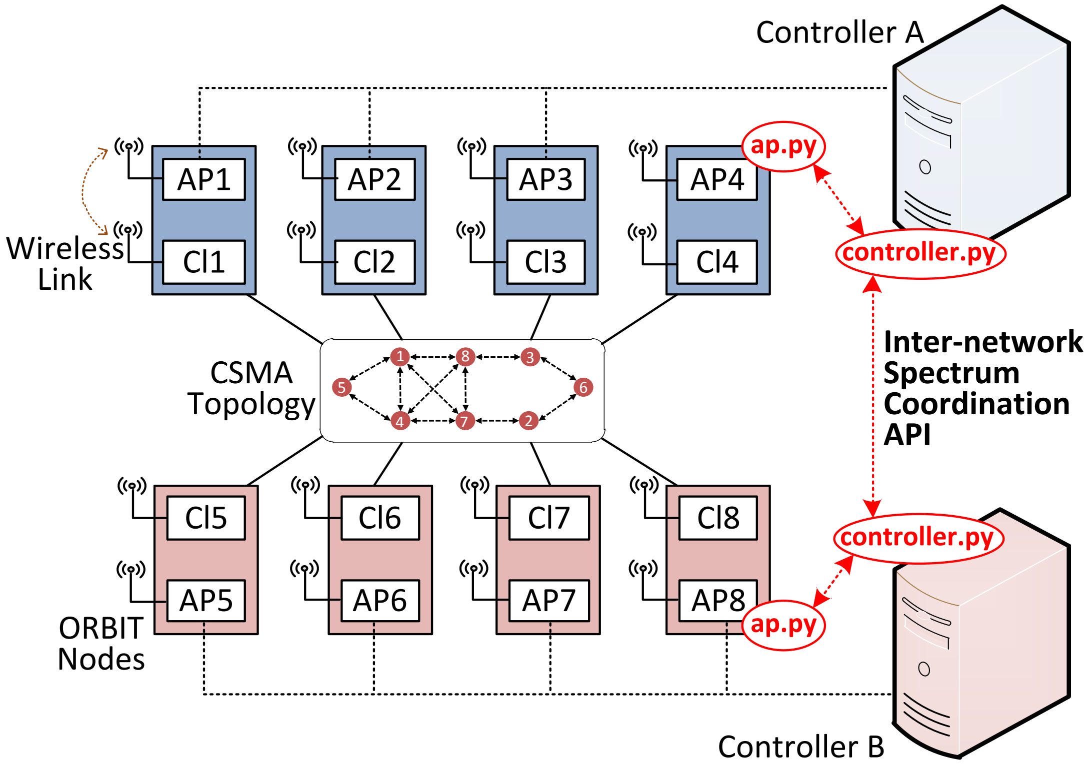

Figure 5.3: Experiment setting for proof-of-concept validation

us to create arbitrary topologies. Fig. 5.3 shows the topology used for this experiment along with the AP-to-Controller and Controller-to-Controller interactions. The payload used for the throughput tests in this case is downlink UDP flows with the offered load chosen uniformly randomly between 5-45 Mbps for each AP-Client pair. The offered load once chosen for each AP-client pair, remains constant throughout the experiment. As before, we assume only two channels are available for use by all the nodes.

We built simple AP-side and controller-side python applications which perform the basic tasks of beacon reporting, load reporting, setting/getting channels, and communicating using the API defined above. The 60-second experiments starts with both controllers collecting beacon reports from APs in its domain and setting the channels according to the centralized algorithm described in Chapter 4. Note that at this stage, there is no cooperation between the two controllers, yet an almost complete interference graph is visible to both networks since one or more of the four APs in each network can see all the four APs in the other network. In the topology shown in Fig. 5.3, all the links are visible to Network B while the only link that is not visible to Network A

Table 5.1: Load and channels for validation experiment

Attribute	**Values**
Offered Load (Mbps)	40, 41, 9, 7, 12, 40, 9, 22
Channels (before coordination)	1, 1, 1, 11, 1, 11, 11, 1
Channels (after coordination)	1, 1, 1, 11, 11, 11, 11, 11

is the one between nodes 7 and 8 since both its end-points lie in Network B.

At the 25 second mark, the two controllers initiate a cooperation handshake through the sequence of steps outlined in Fig. 5.1. In this example, Network B then hands over the control of the channels of its APs to Network A while also providing it information about the average load on each of its APs. Network A then uses the load information and uses a modified centralized algorithm which tries to minimize the load on each channel through a greedy search process. For the particular example we present next, Table 5.1 shows the selected offered load values, and the channels selected under the non-cooperation and cooperation regimes.

Fig. 5.4 shows the throughput of each AP-client pair for the entire duration of the experiment. Note that at time $t = 25$ seconds, the throughputs of nodes 5 and 8 go to zero since the controller sends a change channel commands to both these nodes. Immediate changes are also observed in the throughputs of nodes 1 and 4 - node 4's throughput goes to zero briefly due to the change in the topology when nodes 5 and 8 stop transmitting, but comes back up to its original level shortly afterwards. Fig. 5.5 summarizes the node throughputs during the before- and after-cooperation phases of the experiment. We can see that nodes 1 and 8 benefit from the change in channel assignment made possible by sharing of load information between the two networks.

Admittedly, this is a very simple, some-what contrived example which shows the benefit of sharing information about a single parameter in a single topology. The aim of this exercise, however has been to build a basic functioning code-framework which can prove the effectiveness of our inter-network spectrum coordination API in simple scenarios. Much work needs to be done to elevate this proof-of-concept code to a production code - a prime example being the task of reducing the down-time of an AP after a channel change as observed in the middle segment of Fig. 5.4.

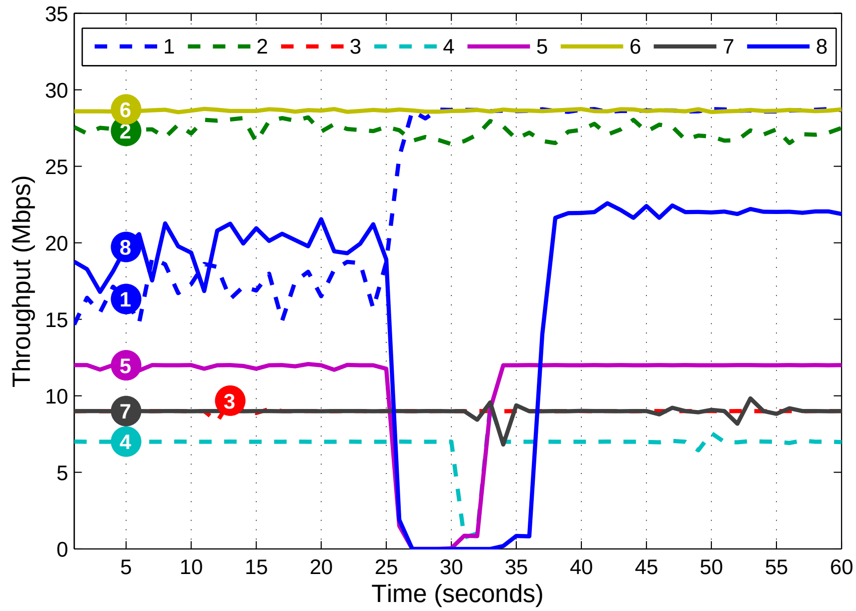

Figure 5.4: Measured throughput from proof-of-concept validation of spectrum coordination: At t=25 seconds, Controller B hands over control of channel selection to Controller A resulting in channel change for nodes 5 and 8.

5.3 A flexible control plane for wireless networks

While a well-defined and well-accepted API is a prerequisite for introducing cooperation between disparate wireless networks, just providing a means of communication is not enough - the information received from neighboring networks needs to be integrated holistically into the control plane implementation of wireless networks. Thus in our ongoing efforts, we are working on ways in which the wireless control plane in the different contexts of WLANs and LTE networks can be redesigned to be made more amenable to taking external inputs for spectrum management. In this section we describe our initial ideas towards that goal.

Fig. 5.6 shows the control architecture of typical enterprise/service-provider Wi-Fi networks - many of the control plane functions in such deployments are centralized and implemented through one of several vendors solutions such as Aruba, Cisco, and Ruckus [71]. The interface between this centralized wireless controller to the AP is partially standardized through the IETF CAPWAP protocol [11], but the algorithms used for dynamic spectrum management at the controller is almost always proprietary with

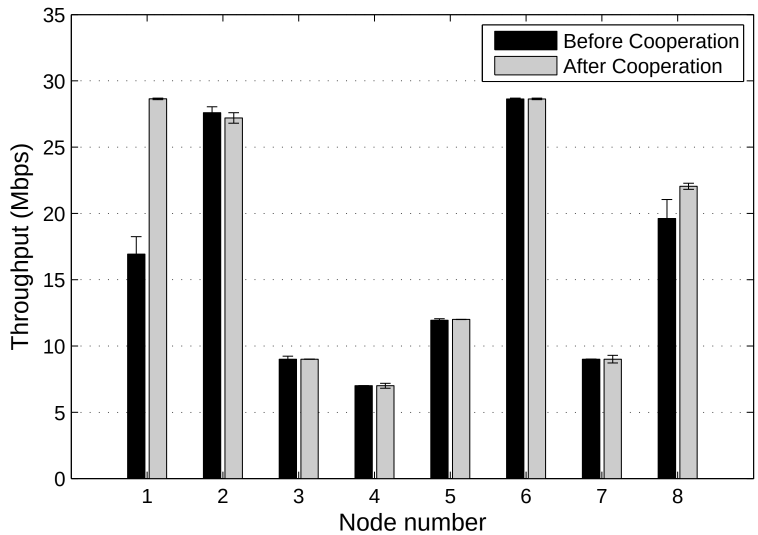

Figure 5.5: Per-node throughput before and after the two networks started cooperating. Nodes 1 and 8 see improvements while others suffer no losses.

no mechanisms for inserting external inputs. This closed-box approach leads to a lack of flexibility in the way in which inter-network cooperation can be implemented for spectrum management. More importantly, the control functions are currently distributed between the controller, AP, and the client devices. As such, even after receiving specific information from a neighboring network, a controller might be unable to change the pre-set algorithms running on the AP or client devices.

Our aim in this ongoing work is to design an open framework for implementing radio control algorithms in large wireless networks. The basic architecture, shown in Fig. 5.7, consists of a central controller which communicates to APs using a flexible switching mechanism, and allows different control plane algorithms, including inter-network coordination to be implemented through simple software applications. A key design principle in this work is to cleanly separate the decision logic behind the control of different parameters from their implementation parts, i.e. a separation between the decisions taken (analogous to control plane) and the actions taken (analogous to the data plane). The success of the SDN paradigm in the wired network is the key motivation behind this design goal and the rapid development of both expertise and code-base could be leveraged in our work. Centralization of *all* decision logic in wireless networks would enable external inputs to be used at that central point without the need

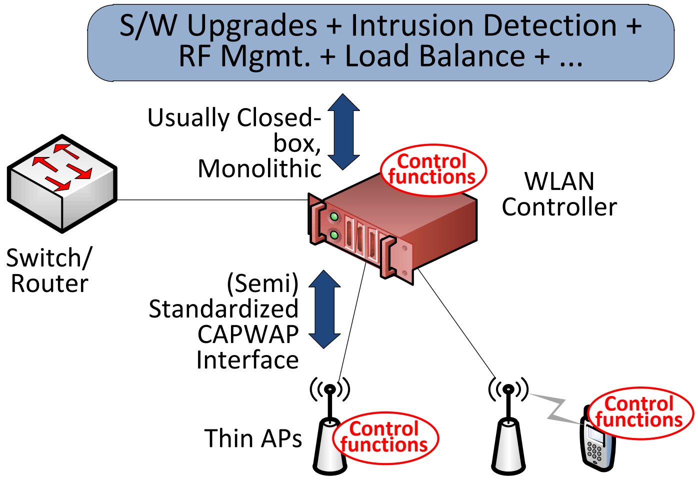

Figure 5.6: Typical WLAN control plane architecture

for implementing interactions between individual network elements and neighboring networks.

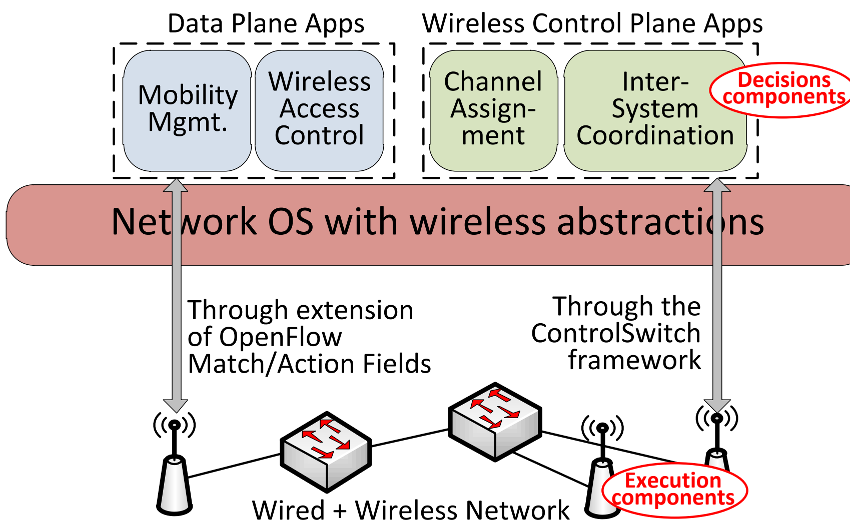

Figure 5.7: A flexible wireless control plane architecture for enabling inter-network collaboration for spectrum management

Chapter 6

Conclusions

The aim of this thesis has been to advance the state-of-art in coordinated dynamic spectrum access techniques for unlicensed band radio devices. Dense deployments of such radio devices are increasingly common in urban hot-spots, multi-tenant buildings, office complexes, and airports. The impending influx of large number of sensor and vehicular wireless devices in these already dense settings require highly efficient and scalable mechanisms for spectrum management. While optimizing a single network has been the subject of a number of studies in the last decade, the clear benefit, scope, and limitations of coordinated DSA are not well understood. In this work, we presented a first step towards that goal and proposed a network-assisted mechanism for increasing the spectrum-visibility of multiple networks operating in overlapping regions.

First, a brief overview of future Internet design considerations was presented, which are driven by emerging wireless access and mobility scenarios. Several key protocol requirements have been identified including name/address separation, robustness with respect to link quality variation and disconnection, multi-homing, ad hoc network formation, and spectrum coordination. Key design features of the MobilityFirst protocol stack have been outlined and shown to address some of these requirements. A novel network-assisted approach for dynamic spectrum coordination was proposed which leverages the management plane of this future Internet architecture. The proposed approach is intended for application to current unlicensed bands and for emerging white space and cognitive radio scenarios. Key design components necessary to implement the proposed coordination architecture have been described and validated. We note that the IETF has recently initiated standardization of an interface for access to a spectrum database (PAWS [12]), and it may be appropriate to consider further extensions to this or other networking standards to provide support for distributed inter-network spectrum coordination as well.

Detailed simulation results were presented for two specific dense Wi-Fi problems. The focus of the first problem was optimization of client to access point association in managed deployments. Since such planned AP deployments are designed to support a large number of users, balancing the number of clients associated to each AP is important. We show that ignoring the presence of other networks leads to significant throughput degradation, especially for clients at the edge of an APs coverage region. To alleviate this problem, an operational cooperation model was proposed, under which all networks share the information about the location and operating channel for all their APs. Results show that incorporating this information for client-association optimization within a network leads to 140-170% improvements in the 10 percentile client throughputs when clients and APs are randomly placed. Clustered AP deployments lead to a much higher gain of up to 7x since the value of the shared information increases significantly.

In the next problem, the importance of inter-network coordination was highlighted in the context of channel assignment for dense Wi-Fi networks. A mixed strategy of large-scale but simplified simulations combined with small-scale but detailed experimental validation was undertaken. Results show that in dense settings, a coordinated approach towards channel assignment can improve client throughputs by up to 40% and reduce the number of starved nodes by up to 4x compared to random channel assignments. Mixed scenarios were also studied, in which a varying fraction of the APs act independently while the other fraction are coordinated. In such settings, both simulation and experiments on the ORBIT testbed show that an increasing number of APs under coordination leads to improvements in throughput averaged over all the APs, and also for independent APs measured separately.

The ideas presented in this work can be implemented in real-world deployments through incremental changes in software architecture of the wireless devices and the management architecture of the network. While a more native Internet-layer support for spectrum coordination messages might require substantial business incentives, an over-the-top rendezvous system can be deployed to meet the broad functional requirements, if not the low-latency and simplified message distribution requirements. For a more broader deployment, further work is required in two key aspects. Firstly, a more evolved theoretical foundation needs to be established to understand the coordinated dynamic spectrum access regime, and its relation to sensing-based and database-type approaches. For example, given a sufficiently fast back-end coordination system, it might be possible to substantially reduce the amount of time a radio spends in sensing whether a channel is free or not. This is in comparison to the current CSMA regime in which carrier sensing has to be performed by every device before every transmission.

The second key aspect of the problem which requires further efforts is a re-design of the software architecture of unlicensed band wireless networks. Due to a combination of legacy-support, latency, and ease-of-deployment issues, the algorithms that control different spectrum-access parameters in most wireless networks are executed at different locations - some inside device drivers, some as user-plane configuration options, and some at an external controller entity. Taking a cue from the software defined networking (SDN) paradigm, the architecture can be made both more easier to manage and more flexible if all the decision logic behind the setting of different parameters can be separated from their implementation parts, so as to create a clean separation between the decisions taken (analogous to control plane) and the actions taken (analogous to the data plane). Such an architecture could prove much more amenable to the introduction of coordination since the inputs taken by the decision logic can be augmented by information provided by other networks in the neighborhood. As outlined earlier in the thesis, there are substantial challenges in the design and implementation of such a system, but if the engineering challenges prove solvable, this change in the architecture could provide significant improvements in the efficiency and flexibility of wireless networks.

References

[1] Ericsson White Paper, "More than 50 billion connected devices," http://www.ericsson.com/res/docs/whitepapers/wp-50-billions.pdf.

[2] Federal Communications Commission, "In the Matter of Unlicensed Operation in the TV Broadcast Bands: Third Memorandum Opinion and Order," Apr. 2012.

[3] ——, "FCC Notice of Proposed Rulemaking and Order 12-148, Enabling Innovative Small Cell Use In 3.5 GHz Band," Dec. 2012.

[4] D. B. Klein, "Rinkonomics: A Window on Spontaneous Order," http://www.econlib.org/library/Columns/y2006/Kleinorder.html.

[5] D. Raychaudhuri and X. Jing, "A spectrum etiquette protocol for efficient coordination of radio devices in unlicensed bands," in *Proc. 14th Intl. Symp. on Personal, Indoor and Mobile Radio Comm. (IEEE PIMRC)*, 2003, pp. 172–176.

[6] M. Buddhikot, P. Kolodzy, S. Miller, K. Ryan, and J. Evans, "DIMSUMnet: new directions in wireless networking using coordinated dynamic spectrum," in *World of Wireless Mobile and Multimedia Networks, Sixth IEEE International Symposium on a*, June 2005, pp. 78–85.

[7] V. Brik, E. Rozner, S. Banerjee, and P. Bahl, "DSAP: a protocol for coordinated spectrum access," in *New Frontiers in Dynamic Spectrum Access Networks, First IEEE International Symposium on*, Nov 2005, pp. 611–614.

[8] J. Xiong and K. Jamieson, "Towards Fine-grained Radio-based Indoor Location," in *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*, ser. HotMobile '12, 2012, pp. 13:1–13:6.

[9] M. Mauve, A. Widmer, and H. Hartenstein, "A survey on position-based routing in mobile ad hoc networks," *Proc. IEEE Network*, vol. 15, no. 6, pp. 30–39, Aug. 2002.

[10] S. C. Nelson, G. Bhanage, and D. Raychaudhuri, "GSTAR: generalized storage-aware routing for mobilityfirst in the future mobile internet," in *Proceedings of MobiArch, 2011*, pp. 19–24.

[11] P. Calhoun, M. Montemurro, and D. Stanley, "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification," Internet Requests for Comments, RFC 5415, March 2009.

[12] V. Chen, D. S., L. Zhu, J. Malyar, and P. McCann, *Protocol to Access White-Space Database*, IETF Internet Draft, draft-ietf-paws-protocol-11.txt, March 2014.

[13] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *Communications Magazine, IEEE*, vol. 50, no. 7, pp. 26–36, 2012.

[14] Cisco White Paper, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013-2018," Feb. 2014.

[15] Republic Wireless: What it costs, http://www.republicwireless.com/what-it-costs.

[16] S. Bhaumik, S. P. Chandrabose, M. K. Jataprolu, G. Kumar, A. Muralidhar, P. Polakos, V. Srinivasan, and T. Woo, "CloudIQ: a framework for processing base stations in a data center," in *Proc. 18th Annual Int'l Conf. on Mobile computing and networking (Mobicom)*, 2012, pp. 125–136.

[17] NSF Future Internet Architecture Project, http://www.nets-fia.net/.

[18] R. Moskowitz and P. Nikander, *Host Identity Protocol (HIP) Architecture*, IETF Internet Standard, RFC 4423, May 2006.

[19] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, I. Stoica, and S. Shenker, "ROFL: Routing on Flat Labels," in *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '06, 2006, pp. 363–374.

[20] D. Han, A. Anand, F. Dogar, B. Li, H. Lim, M. Machado, A. Mukundan, W. Wu, A. Akella, D. G. Andersen, J. W. Byers, S. Seshan, and P. Steenkiste, "Xia: efficient support for evolvable internetworking," in *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, ser. NSDI'12, 2012.

[21] "Named Data Networking (NDN) Project," http://www.named-data.net/ndn-proj.pdf, 2010.

[22] MobilityFirst: A Robust and Trustworthy Mobility-Centric Architecture for the Future Internet, http://mobilityfirst.winlab.rutgers.edu/.

[23] A. Baid and D. Raychaudhuri, "Wireless access considerations for the mobility-first future internet architecture," in *Sarnoff Symposium (SARNOFF), 2012 35th IEEE*. IEEE, 2012, pp. 1–5.

[24] C. Perkins, "Mobile IP," *Communications Magazine, IEEE*, vol. 35, no. 5, pp. 84–99, May 1997.

[25] G. Xylomenos, G. Polyzos, P. Mahonen, and M. Saaranen, "TCP performance issues over wireless links," *Communications Magazine, IEEE*, vol. 39, no. 4, pp. 52–58, April 2001.

[26] J. Iyengar, P. Amer, and R. Stewart, "Concurrent Multipath Transfer Using SCTP Multihoming Over Independent End-to-End Paths," *Networking, IEEE/ACM Transactions on*, vol. 14, no. 5, pp. 951–964, 2006.

[27] 3GPP TS 36.420, *X2 general aspects and principles*, http://www.3gpp.org.

[28] D. G. Andersen et al., "Accountable Internet Protocol (AIP)," in *Proc. ACM SIGCOMM'08*, Aug. 2008.

[29] T. Vu et al., "DMap: A Shared Hosting Scheme for Dynamic Identifier to Locator Mappings in the Global Internet," in *Proceedings of ICDCS, 2012.*

[30] V. Vassiliou and Z. Zinonos, "An Analysis of the Handover Latency Components in Mobile IPv6," *Journal of Internet Engineering*, vol. 3, no. 1, pp. 230–240, Dec 2009.

[31] A. Mishra, M. Shin, and W. Arbaugh, "An empirical analysis of the IEEE 802.11 MAC layer handoff process," *SIGCOMM Comput. Commun. Rev.*, vol. 33, pp. 93–102, April 2003.

[32] N. Somani, A. Chanda, S. C. Nelson, and D. Raychaudhuri, "Storage-Aware Routing for Robust and Efficient Services in the Future Mobile Internet," in *Proceedings of ICC FutureNet V, 2012.*

[33] A. Baid, T. Vu, and D. Raychaudhuri, "Comparing alternative approaches for networking of named objects in the future Internet," in *Proc. INFOCOM Workshop on Emerging Design Choices in Name-Oriented Networking (NOMEN)*, 2012, pp. 298–303.

[34] A. Baid, S. Mukherjee, T. Vu, S. Mudigonda, K. Nagaraja, J. Fukuyama, and D. Raychaudhuri, "Enabling vehicular networking in the MobilityFirst future internet architecture," in *Proc. of 14th Int'l Symposium and Workshops on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2013.

[35] S. Mukherjee, A. Baid, I. Seskar, and D. Raychaudhuri, "Network-Assisted Multihoming in the MobilityFirst Future Internet Architecture," in *WINLAB Technical Report WINLAB-TR-415*, 2013.

[36] S. Deb, K. Nagaraj, and V. Srinivasan, "MOTA: Engineering an operator agnostic mobile service," in *Proc. 17th Annual Intl. Conf. on Mobile Computing and Networking (ACM MobiCom)*, 2011, pp. 133–144.

[37] Optimum WiFi locations, http://www.optimum.net/WiFi/Find.

[38] Xfinity WiFi locations, http://www.comcast.com/wifi/hotspot.htm.

[39] N. Ahmed, U. Ismail, S. Keshav, and K. Papagiannaki, "Online estimation of RF interference," in *Proc. 2008 ACM CoNEXT Conf.*, 2008, pp. 4:1–4:12.

[40] L. Li, M. Pal, and Y. Yang, "Proportional Fairness in Multi-Rate Wireless LANs," in *Proc. 27th Conf. on Computer Comm. (IEEE INFOCOM)*, 2008, pp. 1004–1012.

[41] I.-H. Hou and P. Gupta, "Distributed Resource Allocation for Proportional Fairness in Multi-Band Wireless Systems," in *Proc. Intl. Symp. on Information Theory (IEEE ISIT)*, 2011, pp. 1975–1979.

[42] O. Ercetin, "Association games in IEEE 802.11 wireless local area networks," *Proc. IEEE Trans. Wireless Comm.*, vol. 7, no. 12, pp. 5136–5143, Dec. 2008.

[43] Aruba Networks White Paper, ARM Yourself to Increase Enterprise WLAN Data Capacity, http://bit.ly/5nrjZ.

[44] L. B. Jiang and S. C. Liew, "Proportional fairness in wireless LANs and ad hoc networks," in *Proc. Wireless Comm. and Networking Conf. (IEEE WCNC)*, vol. 3, 2005, pp. 1551–1556.

[45] K. Sundaresan and K. Papagiannaki, "The need for cross-layer information in access point selection algorithms," in *Proc. 6th ACM SIGCOMM Conf. on Internet Measurement (ACM IMC)*, 2006, pp. 257–262.

[46] S. Choi, K. Park, and C.-k. Kim, "On the performance characteristics of WLANs: revisited," in *Proc. SIGMETRICS Perform. Eval. Rev.*, vol. 33, no. 1, Jun. 2005, pp. 97–108.

[47] Z. W. Zhao and Y. C. Tay, "A model for calculating channel share of 802.11 access points with overlapping wireless cells," *Proc. IEEE Trans. Wireless Netw.*, vol. 17, pp. 1581–1593, Oct. 2011.

[48] B. Jang and M. Sichitiu, "IEEE 802.11 Saturation Throughput Analysis in the Presence of Hidden Terminals," *Proc. IEEE/ACM Trans. Netw.*, vol. 20, no. 2, pp. 557–570, 2012.

[49] D. Raychaudhuri *et al.*, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," in *Proc. Wireless Comm. and Networking Conf. (IEEE WCNC)*, vol. 3, 2005, pp. 1664–1669.

[50] Q. Wu and E. Esteves, "The CDMA2000 high rate packet data system," in *Advances in 3G Enhanced Technologies for Wireless Communications.* Artech House, 2002.

[51] D. B. Shmoys and E. Tardos, "An approximation algorithm for the generalized assignment problem," *Proc. Math. Program.*, vol. 62, pp. 461–474, December 1993.

[52] "lp_solve, a Mixed Integer Linear Programming (MILP) solver," http://lpsolve.sourceforge.net/.

[53] Cisco Aironet 1200 Data sheet, http://bit.ly/dv9MbI.

[54] S. Gundavelli, M. Grayson, P. Seite, and Y. Lee, *Service Provider Wi-Fi Services Over Residential Architectures*, IETF Internet Draft, draft-gundavelli-v6ops-community-wifi-svcs-06.txt, April 2013.

[55] WiGLE: Wireless geographic logging engine, http://www.wigle.net/.

[56] M. Achanta, "Method and apparatus for least congested channel scan for wireless access points," U.S. Patent 20 060 072 602, April, 2006.

[57] S.-C. Liew, C. H. Kai, H. C. Leung, and P. Wong, "Back-of-the-Envelope Computation of Throughput Distributions in CSMA Wireless Networks," *Mobile Computing, IEEE Transactions on*, vol. 9, no. 9, pp. 1319–1331, 2010.

[58] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 18, no. 3, pp. 960–972, Jun. 2010.

[59] B. Nardelli, J. Lee, K. Lee, Y. Yi, S. Chong, E. Knightly, and M. Chiang, "Experimental evaluation of optimal CSMA," in *INFOCOM, 2011 Proceedings IEEE*, 2011, pp. 1188–1196.

[60] S. Chieochan, E. Hossain, and J. Diamond, "Channel assignment schemes for infrastructure-based 802.11 WLANs: A survey," *Communications Surveys Tutorials, IEEE*, vol. 12, no. 1, pp. 124–136, 2010.

[61] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *Selected Areas in Communications, IEEE Journal on*, vol. 18, no. 3, pp. 535–547, 2000.

[62] I. M. Bomze, M. Budinich, P. M. Pardalos, and M. Pelillo, "The Maximum Clique Problem," in *Handbook of Combinatorial Optimization.* Kluwer Academic Publishers, 1999, pp. 1–74.

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

[64] J. Riihijarvi, M. Petrova, and P. Mahonen, "Frequency allocation for WLANs using graph colouring techniques," in *Proc. 2nd Annual Conference on Wireless On-demand Network Systems and Services (WONS)*, 2005, pp. 216–222.

[65] W. Yin, K. Bialkowski, J. Indulska, and P. Hu, "Evaluations of MadWifi MAC layer rate control mechanisms," in *Quality of Service (IWQoS), 2010 18th International Workshop on*, June 2010, pp. 1–9.

[66] Iperf network testing tool, http://www.noc.ucf.edu/Tools/Iperf/.

[67] Ramsey Electronics RF Enclosure, http://www.ramseytest.com/.

[68] Sandbox 4 at ORBIT, http://bit.ly/1eeYi2w.

[69] Information System on Graph Classes and their Inclusions, http://www.graphclasses.org/smallgraphs.html#nodes5.

[70] Number of connected graphs with n nodes. , http://oeis.org/A001349.

[71] School of Information Studies, Syracuse University, "Competitive Evaluation of Three-stream 802.11n Access Points," http://bit.ly/1pY4MrX.